

**THE VOLGENAU SCHOOL OF ENGINEERING
DEPT OF ELECTRICAL AND COMPUTER ENGINEERING
SYSTEM ARCHITECTURES LABORATORY**

**Computational Modeling of Cultural Dimensions
in Adversary Organizations**

Grant Number: FA9550-05-1-0388

FINAL TECHNICAL REPORT

01 June 2005 to 30 November 2010

Submitted to:

Air Force Office of Scientific Research

One Liberty Center
875 North Randolph Street
Arlington, VA 22203-1995

Attn: **Dr. Terence Lyons**

AFOSR/RSL

(703) 696 9542

Fax: (703) 696 7360

Submitted by:

Alexander H. Levis

George Mason University
System Architectures Lab
ECE Dept., MS 1G5
Fairfax, VA 22030

(703) 993 1619

Fax: (703) 993 1601

email:alevis@gmu.edu



Carnegie Mellon

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2010		2. REPORT TYPE		3. DATES COVERED 00-00-2010 to 00-00-2010	
4. TITLE AND SUBTITLE Computational Modeling Of Cultural Dimensions In Adversary Organizations				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) George Mason University System Architectures Lab,ECE Dept., MS 1G5,Fairfax,VA,22030				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES Government or Federal Purpose Rights License.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 355	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

REPORT CONTRIBUTORS

George Mason University

Alexander H. Levis (PI)

Lee W. Wagenhals
Abbas K. Zaidi
Claudio Cioffi-Revilla
Robert J. Elder
Tod S. Levitt
Peter Pachowicz
Ashraf AbuSharekh
Smriti K. Kansal
A. Erkin Olmez
Faisal Mansoor
M. Faraz Rafi
Pedro Romero
John Pham

Carnegie Mellon University

Kathleen M. Carley (Co-PI)

Il-Chul Moon
Geoffrey Morgan
Jesse St. Charles
Brian Hirshman
Michael Lanham

University of Colorado - Denver

P. Papantoni-Kazakos

Table of Contents

	<u>Page</u>
List of Figures	vii
List of Tables	xiii
PART I: INTRODUCTION	1
Chapter 1: Introduction	3
PART II: TIMED INFLUENCE NETS: Theory and Applications	9
Chapter 2: Course of Action Analysis in a Cultural Landscape using Influence Nets	10
Chapter 3: Theory of Influence Networks	15
Chapter 4: Meta-model Driven Construction of Timed Influence Nets	41
Chapter 5: Adversary Modeling Applications	51
PART III: MODELS OF ORGANIZATIONS	95
Chapter 6: Computationally Derived Models of Adversary Organizations	97
Chapter 7: Extracting Adversarial Relationships from Texts	121
Chapter 8: Inferring and Assessing Informal Organizational Structures from an Observed Dynamic Network of an Organization	128
Chapter 9: Simulating the Adversary: Agent-based Dynamic Network Modeling	149
Chapter 10: Adversary Modeling – Applications of Dynamic Network Analysis	171
PART IV: META-MODELING AND MULTI-MODELING	203
Chapter 11: Introduction to Multi-modeling and Meta-modeling	205
Chapter 12: Meta-modeling for Multi-modeling Interoperation	217
PART V: COMPUTATIONAL EXPERIMENT	235
Chapter 13: Cyber Deterrence Policy and Strategy	237
Chapter 14: Application: The India-Pakistan Crisis Scenario	245
References	289
 Appendix A: Proof of Lemmas in Chapter 3	 305
Appendix B: Pythia, a Timed Influence Net Application	309
Appendix C: The C2 Wind Tunnel	315
Appendix D: Activation Timed Influence Nets	317
Appendix E: Modeling and Simulating Terrorist Networks in Social and Geospatial Dimensions	329

LIST OF FIGURES

Fig. 2.1	An example Timed Influence Net (TIN)	11
Fig 2.2	Probability profile for node C	13
Fig. 3.1	Cause-Effect relationships	17
Fig. 3.2	Example TIN	22
Fig. 3.3	Example TIN	35
Fig. 3.4.	Example TIN with COA and edge delays	36
Fig. 3.5.	Temporal model for the example TIN	37
Fig. 3.6.	Probability profile for the example COA	38
Fig. 3.7.	A Multi-node network	38
Fig. 4.1	Architecture of the approach	43
Fig. 4.2	An example mapping	44
Fig. 4.3	Architecture with respective applications	45
Fig. 4.4	Construction process	46
Fig. 4.5	Class hierarchy of the Kenya and Tanzania bombing ontology	47
Fig. 4.6	Template TIN used in the application	47
Fig. 4.7	Instantiated Timed Influence Net for Tanzania bombing	48
Fig. 5.1	Overall events data analysis process conducted in this study, starting with O’Grady’s [124], [125] data on attacks.	53
Fig. 5.2	Cumulative probability density for time between attacks T, Diyala Province, Iraq. March, 2003 - March, 2006	59
Fig. 5.3	Probability density for time between attacks T, Diyala Province, Iraq. March, 2003 - March, 2006	59
Fig. 5.4	Empirical survival function $\hat{S}(t)$, for time between attacks T, Kaplan-Meier estimate, Diyala Province, Iraq.	60
Fig. 5.5	Diyala Province, Iraq. March, 2003 - March, 2006	61
Fig. 5.6	The empirical complementary c.d.f. for time between attacks T in log-log space, Diyala Province, Iraq. March, 2003 - March, 2006	61
Fig. 5.7	Diyala Province, Iraq. Period 1, March, 2003 - June, 2004	62
Fig. 5.8	Diyala Province, Iraq. Period 2, July, 2004 - June, 2005	63
Fig. 5.9	Diyala Province, Iraq. Period 3, July, 2005 - March, 2006	63
Fig. 5.10	Diyala Province, Iraq. Period 1, March, 2003 - June, 2004	64
Fig. 5.11	Diyala Province, Iraq. Period 2, July, 2004 - June, 2005	65
Fig. 5.12	Diyala Province, Iraq. Period 3, July, 2005 - March, 2006	65
Fig. 5.13	Diyala Province, Iraq. March, 2003 - March, 2006	66

Fig. 5.14	Diyala Province, Iraq. March, 2003 - March, 2006	66
Fig. 5.15	Empirical complementary cumulative probability function for severity of attacks S (fatalities), Kaplan-Meier estimate, Diyala Province, Iraq. March, 2003 - March, 2006	67
Fig. 5.16	Diyala Province, Iraq. March, 2003 - March, 2006	68
Fig. 5.17	Empirical c.c.d.f. of severity S (fatalities) in log-log space, Diyala Province, Iraq. March, 2003 - March, 2006	68
Fig. 5.18	Diyala Province, Iraq. Period 1, March, 2003 - June, 2004	69
Fig. 5.19	Diyala Province, Iraq. Period 2, July, 2004 - June, 2005	70
Fig. 5.20	Diyala Province, Iraq. Period 3, July, 2005 - March, 2006	70
Fig. 5.21	Diyala Province, Iraq. March, 2003 - June, 2004	71
Fig. 5.22	Diyala Province, Iraq. July, 2004 - June, 2005	71
Fig. 5.23	Diyala Province, Iraq. July, 2005 - March, 2006	72
Fig. 5.24	Complete model of the case study TIN	81
Fig. 5.25	Static Quantitative COA Comparison	83
Fig. 5.26	Dynamic Temporal Analysis Input	85
Fig. 5.27	Probability Profiles of Scenario (COA) of Fig. 5.26	86
Fig. 5.28	Comparison of the Effect of Different Scenarios	87
Fig. 5.29	Timed Influence Net of East Timor Situation	90
Fig. 5.30	Sample TIN for Analysis	91
Fig. 5.31	Probability Profiles Generated by the CAST Logic Approach	92
Fig. 5.32	Probability Profiles for Case I	93
Fig. 5.33	Probability Profiles for Case II	94
Fig. 6.1	Model of the Five-Stage Decision Maker	98
Fig. 6.2	One-sided Interactions Between Decision Maker i and Decision Maker j	99
Fig. 6.3	Flowchart for culturally constrained solution space	104
Fig. 6.4	Command Relationship Chart for Red	105
Fig. 6.5	Block Diagram of the Organization as seen in the CAESAR III GUI	106
Fig. 6.6	Matrix representation of the design problem	107
Fig. 6.7	Universal Net	107
Fig. 6.8	Partially expanded solution space	108
Fig. 6.9	Culturally Constrained Solution Space for Red	108
Fig. 6.10	Expanded Lattice Structure from C-MINO(1) to CMAXO(1) for Red	109
Fig. 6.11	C-MINO(1) for Red	109
Fig. 6.12	C-MAXO(1) for Red	110

Fig. 6.13	C-MAXO(2) for Red	110
Fig. 6.14	C-MAXO(3) for Red	110
Fig. 6.15	Expanded Lattice Structure from C-MINO(1) to CMAXO(1) for Blue	111
Fig. 6.16	C-MAXO(1) for Blue	111
Fig. 6.17	Level-1 organizational block diagram	113
Fig. 6.18	Matrix Representation corresponding to Fig. 6.17	113
Fig. 6.19.	Solution space for Level-1 organization design as seen in CAESAR III	114
Fig. 6.20	MINO of Level-1 design	114
Fig. 6.21	MAXO of Level-1 design	114
Fig. 6.22	Block diagram and matrix representation for ACE	115
Fig. 6.23	Block diagram and matrix representation for GCE	115
Fig. 6.24	Block diagram and matrix representation for CSSE	116
Fig. 6.25	GCE structure selected for US	117
Fig. 6.26	GCE structure selected for Country A	117
Fig. 6.27	GCE structure selected for Country B	118
Fig. 6.28	Percent of tasks un-served for coalition options	118
Fig. 8.1	The visualization of the meta-matrix of the terrorist group responsible for the 1988 U.S. embassy bombing in Kenya	132
Fig. 8.2	The terrorist social network in the meta-matrix	132
Fig. 8.3	The task network in the meta-matrix	133
Fig. 8.4	The procedure of the introduced analysis framework	134
Fig. 8.5	The partial visualization of the task precedence network (task-to-task) and the task assignment network (terrorist-to-task).	136
Fig. 8.6a	A partial visualization explaining the formation of information sharing links: First step, <i>Ali Mohamed</i> is assigned to <i>surveillance of possible targets</i> .	137
Fig. 8.6b	Second step, <i>Ali Mohamed</i> requires <i>surveillance expertise</i> to perform his assigned task, but he does not have it.	137
Fig. 8.6c	Third step, the organization searches an agent with <i>surveillance expertise</i> from the agents near to <i>Ali Mohamed</i> . It finds an agent two social links away, <i>Anas Al-Liby</i> .	137
Fig. 8.6d	Fourth step, <i>Anas Al-Liby</i> has the required expertise and has to deliver the expertise through the social links.	138
Fig. 8.6e	Fifth step, there are three possible shortest paths from <i>Anas Al-Liby</i> to <i>Ali Mohamed</i> . These paths are information sharing links.	138
Fig. 8.7	A partial visualization of two tasks and ten assigned agents.	139
Fig. 8.8	Three extracted decision making structures. (Top) Information sharing,	

	(Middle) Result sharing, (Bottom) Command interpretation	141
Fig. 8.9	Charts displaying the difference of metrics between a meta-network and extracted structures	145
Fig. 8.10	Two projections of metrics of individuals using two principal components. The left is using only the original structure, and the right is from only the extracted structures.	147
Fig. 9.1	Cycle of Agent Activity	158
Fig. 10.1	The closeness CUSUM statistic graph over time for Al-Qaeda	172
Fig. 10.2	An overall simulation analysis procedure	174
Fig. 10.3	High level agent behavior log	177
Fig. 10.4	An example of agent behavior during the simulation from the Kenya data.	179
Fig. 10.5	A illustrative example of transactive memory transfer.	180
Fig. 10.5	Organizational performance over time, aggregated by the first factor	186
Fig. 10.6	Percentage of Task completion speed to the baseline, 64 virtual experiment cells	187
Fig. 10.7	Percentage of Mission completion speed to the baseline, 64 virtual experiment cells	188
Fig. 10.8	The estimated Gantt chart of the baseline case	189
Fig. 10.9	Collection of agent interaction and organizational transfer network over time, link thickness is adjusted to show the frequency of the link usage.	191
Fig. 10.10	Agent behavior logic. Compared to the previous behavior model, the geospatial relocation and the regional resource/expertise acquisitions are added.	192
Fig. 10.11	Annotated simulation procedure flow chart. The annotation specifies which items in the flow chart correspond to the pseudo code.	196
Fig. 10.12, a, b, c, d	Changes in task metric performance due to interventions.	200
Fig. 10.13	Agents gathered resources and skills and then moved to operational centers.	200
Fig. 11.1	The four layers of multi-modeling	206
Fig. 11.2	Influence Network meta-model	207
Fig. 11.3	Representation of knowledge and software	207
Fig. 11.4	Concatenation	209
Fig. 11.5	Amplification	210
Fig. 11.6	Parameter Discovery	210

Fig. 11.7	Model Construction	211
Fig. 11.8	Model Merging	211
Fig. 11.9	Modeling applications using different modeling languages	212
Fig. 11.10	Fragment of the concept map for Timed Influence Nets.	213
Fig. 11.11	Multiple types of model interoperation	214
Fig. 11.12	Large Screen Displays for C2WT Demonstration	216
Fig. 12.1	Model building overview	217
Fig. 12.2	A multi-modeling environment	219
Fig. 12.3	Overview of the meta-modeling approach	221
Fig. 12.4	Example Influence Net	222
Fig. 12.5	Example Social Network	223
Fig. 12.6	A sample Concept Map for constructs of Influence Net focus question	224
Fig. 12.7	Influence Net syntactic model	224
Fig. 12.8	Influence Net pseudo ontology snippet	225
Fig. 12.9	GraphViz Diagram - Influence Net inferred refactored ontology.	226
Fig. 12.10	GraphViz Diagram – Social Network inferred refactored ontology	227
Fig. 12.11	Enriched ontology classes	229
Fig. 12.12	Subject, Object classes mapped to Agent class	229
Fig. 12.13	Reasoner inferred equivalences	229
Fig. 12.14	Subject, Object, organization and Agent as equivalent classes	229
Fig. 12.15	Class hierarchy of the inferred enriched ontology	231
Fig. 14.1	Scenario timeline	246
Fig. 14.2:	Vignette A workflow	249
Fig. 14.3	Sphere of Influence Graphic for Indian Foreign Minister during Vignette A's time period. Note the presence of Deputy Prime Minister Advani, who was not in the first iteration of Pythia and CAESAR III models.	252
Fig. 14.4	Sphere of Influence Graphic for Pakistani National Security Advisor, for all time periods. There was complete overlap between CAESAR III and Pythia models with this model built through AutoMap and ORA.	253
Fig. 14.5	Sphere of Influence Graphic for Indian Prime Minister during Vignette A	253
Fig. 14.6	Sphere of Influence Graphic for Indian Prime Minister during Vignette B	254
Fig. 14.7	Sphere of Influence Graphic for Indian Prime Minister during Vignette C	254
Fig. 14.8	The base case presented from Vignette A.	255
Fig. 14.9	Pakistani Government organization model	256
Fig. 14.10	Indian Government organization model	258
Fig. 14.11	Sphere of influence of CENTCOM-J5	259

LIST OF TABLES

TABLE 3.1	Comparison of Influence Constants	33
TABLE 3.2	Conditional Probabilities	35
TABLE 3.3	Posterior Probabilities of B	37
TABLE 3.4	Probability Profile values	37
TABLE 5.1	Onset of attacks T (days between events)	58
TABLE 5.2	Shapiro-Wilk Test	60
TABLE 5.3	Severity of attacks S (fatalities data were either normally distributed or belonged to a lognormal distribution)	62
TABLE 5.4	Shapiro-Wilk Test	67
TABLE 5.5	The two Courses of Action	91
TABLE 6.1	Cultural Constraints	105
TABLE 6.2	Hofstede's scores for the three countries	116
TABLE 6.3	Cultural Constraints corresponding to ACE	116
TABLE 6.4	Cultural Constraints corresponding to GCE	117
TABLE 6.5	Cultural Constraints corresponding to CSSE	117
TABLE 8.1	The meta-network of the dataset, a terrorist group responsible for 1998 U.S. embassy bombing in Kenya. The numbers in the cells are the densities of the adjacency matrices.	131
TABLE 8.2	A table of descriptive statistics for the metrics. This table includes means, standard deviations, and a cross-correlation table.	133
TABLE 8.3	Three traditional centrality metrics and two dynamic network metrics used to assess the criticalities of individuals in the structure	139
TABLE 8.4	A table of QAP correlation and other distance metrics between the original structure and the extracted decision making structures.	142
TABLE 8.5	A table of MRQAP regression results.	142
TABLE 8.6	A table of top three individuals from five metrics and four structures	143
TABLE 8.7	I.D. assignments to individuals. I.D.s will be used to distinguish individuals in the later tables.	144
TABLE 8.8	Coefficients of two principal components from the original structure (top) and the extracted structures (bottom)	146
TABLE 9.1	A table illustrating how a user can characterize different classes of agents by specifying their number, activity, and message capabilities	162

TABLE 9.2	A table illustrating how a user can characterize a population by differentially distributing information and beliefs across classes of agents.	162
TABLE 9.3	A table illustrating how the user can differentiate agents by varying the socio-demographics.	164
TABLE 9.4	A table illustrating how the user can differentiate agents based on constraints.	165
TABLE 9.5	A table illustrating how to define agent classes by varying the information processing capabilities of the agents in that class.	166
TABLE 9.6	A table illustrating the way in which the user can adapt the agent classes by specifying the size of the sphere of influence per class.	167
TABLE 10.1	This table contains a summary the input and output variables, and the associated parameters, for the JDyNet simulation runs with associated names and description.	175
TABLE 10.2	A table describing the design of a virtual experiment assessing the impact of diverse courses of action for targeting difference adversaries. For each cell shown there would be 15 replications and 2500 simulation time steps.	181
TABLE 10.3	Dynamic network metrics used to determine the target agents to remove	182
TABLE 10.4	A table showing the standardized coefficients for regression to the six organizational performance metrics at the end time using the virtual experiment settings	184
TABLE 10.5	A table showing the standardized coefficients for regression to the six organizational performance metrics at the end time using the calculated metrics of removed agents (N=64 cases) (* for P<0.05)	185
TABLE 10.6	Geospatial simulation model main loop	192
TABLE 10.7	Geospatial simulation iteration for each time-step	193
TABLE 10.8	High level agent behavior	193
TABLE 10.9	Agent's social interaction implementation pseudo code	194
TABLE 10.10	Agent's transactive management pseudo code	195
TABLE 10.11	Agent's task execution implementation pseudo code	195
TABLE 10.12	A table describing the key parameters in the simulation and the implication of setting these parameters	197
TABLE 10.13	Virtual experiment design for simulation parameters (30 replications, 2500 simulation time-steps)	198
TABLE 10.14	A table of standardized coefficients for regression to the six Organizational performance metrics at the end time using the virtual experiment settings	198

TABLE 12.1	Influence Net refactored ontology elements (Concept Map Imports)	227
TABLE 12.2	Explicit Influence Net Concepts in Refactored Ontology	227
TABLE 12.3	Social Network refactored ontology elements (Concept Map Imports)	228
TABLE 12.4	Enriched ontology	230
TABLE 14.1	Scenario and Vignette timeline	247
TABLE 14.2	Vignette A, National Security Council only, CENTCOM & PACOM	251
TABLE 14.3	Vignette A, NSC and diplomats only, CENTCOM & PACOM	251
TABLE 14.4	Vignette A, all agents, CENTCOM & PACOM	251
TABLE 14.5	Construct experimental design, Vignette A	255
TABLE 14.6	CENTCOM sphere of influence report	259
TABLE 14.7	PACOM sphere of influence report	260
TABLE 14.8	Vignette A, National Security Council only, CENTCOM & PACOM	270
TABLE 14.9	Vignette B, NSC and Diplomats only, CENTCOM & PACOM	270
TABLE 14.10	Vignette B, all agents, CENTCOM & PACOM	270
TABLE 14.11	Measures reflected in Key Entity tables	272
TABLE 14.12	Construct experiment design, Vignette B	278
TABLE 14.13	CENTCOM sphere of influence report for new US lever	281
TABLE 14.14	PACOM sphere of influence report for new US levers	281
TABLE 14.15	Sphere of influence of common levers	281
TABLE 14.16	Final COA for combined CENTCOM PACOM actions	286

PART I: INTRODUCTION

Chapter 1: Introduction

Chapter 1

Introduction

Alexander H. Levis

The initial objectives of the “Computational Modeling of Cultural Dimensions in Adversary Organizations” were:

- (a) *To relate an adversary’s organizational structure to behavior when both structure and behavior are conditioned by cultural and social characteristics, as they always are in realistic settings.*
- (b) *To address basic research questions centered on locating the points of influence and characterizing the processes necessary to influence organizations in diverse cultures.*
- (c) *To explore, through a computational modeling framework, the nexus between data and models for individual adversaries (micro level) and data and models for organizations of adversaries (macro level).*

As the project evolved, additional objectives were introduced:

- (d) *To explore multi-modeling as a way to model adversary behaviors and research the underlying theory (meta-modeling)*
- (e) *Demonstrate the approach through a case study that addresses issues of deterrence*

A set of tasks was defined for achieving the these objectives. They were:

Task 1: Implement a testbed for computational modeling.

Task 2: Expand and enhance the existing models at George Mason University’s System Architectures Laboratory (GMU/SAL) and at Carnegie Mellon University’s Center for Computational Analysis of Social and Organizational Systems (CMU/CASOS)

Task 3: Conduct computational experiments to address the set of research hypotheses.

Task 4: Develop and transition theory-based tools to the Air Force

Task 5: Provide Education and Training

Task 6: Meta-Modeling for Multi-Modeling Integration

Task 7: Demonstration of Computational Experiment

Task 8: Management and Documentation

All tasks were carried out during the period of performance. In this report, the research approach taken and results obtained in Tasks 1, 2, 6, and 7 are presented. The many transitions of the tools that have taken place (Task 3) have been reported in detail in the annual productivity reports and in the annual program reviews. Similarly, a substantial education and training effort has been made by both collaborating organizations through the training on many graduate research assistants, the conduct of summer institutes (CMU), the offering of AFCEA sponsored short courses (GMU) to DOD personnel and staff of the Defense Industrial Base, as well as nu-

merous seminars and presentations to Air Force and other defense organizations. Much of the research material is now included in graduate level courses at both universities. Task 8 has also been reported annually to the Air Force office of Scientific Research in accordance with grant requirements.

Since 1992 the nature of military operations has changed. The type of objectives that the military has to address has expanded well beyond those of traditional major combat operations. As military operations became other than conventional war – whether against transnational terrorist threats or conducting stabilization operations – the need to broaden the focus of models that support effects based planning and operations became critical. One major weakness was the absence of socio-cultural attributes in the models used for Course of Action selection and effects based planning. Part II of this report illustrates an approach that enables analysts to evaluate complex situations such as those in which an adversary is embedded in a society from which he is receiving support. In Chapters 2 and 3, a modeling approach is described that enables analysts to examine and explain how actions of the military and other entities may result in desired or undesired effects, both on the adversary and on the population as a whole. First, Timed Influence Nets are described (Ch. 2) and then the theory that underlies them as well as some major extensions of the theory are presented in Chapter 3. A comprehensive theory of Influence Networks is presented that incorporates design constraints for consistency, temporal issues and a dynamic programming evolution of the Influence Constants. A software implementation of Timed Influence nets, a modeling and analysis tool called Pythia, is described in Appendix B. This tool has been distributed widely to military and intelligence organizations. One of the difficulties in using models for new situations is the challenge of starting with a blank screen. In Chapter 4 a novel approach for constructing Influence nets quickly is introduced. One of the main challenges in using TINs has been the difficulty in formulating them. Many Subject Matter Experts have difficulty in expressing their knowledge in the TIN representation. A methodology to develop domain specific Timed Influence Nets (TINs) via the use of an ontological representation of domain data is presented. The meta-model driven ontology based approach provides potential assistance to modelers by enabling them to create quickly new models for new situations through the use of Influence Net Templates. An extension of Timed Influence nets into Activation Timed Influence nets is presented in Appendix D.

In Chapter 5, several case studies are presented that use this approach. First, a power law approach for modeling uncertainty is described and used for analyzing adversary behavior. Data collected in the Diyala province in Iraq is used. Uncertainty is a hallmark of conflict behavior and low-intensity warfare, guerrilla, insurgency, and forms of violence that accompany civil war are no exception. In this case study, aspects of the theory of political uncertainty and complexity theory are applied to the analysis of conflict events during the first three years of the second Iraq war, 2003– 2006, limited to the Diyala province. Findings show that neither the time between attacks T or the severity of attacks S (fatalities) have a normal or log-normal distribution. Instead, both variables showed heavy tails, symptomatic of non-equilibrium dynamics, in some cases approximating a power law with critical or near-critical exponent value of 2. The empirical hazard force analysis in both cases showed that the intensity was high for the first occurrences in both variables, namely between March, 2003, and June, 2004, but even higher in a more recent period.

In the second case study, data from the same province are used to develop Courses of Action that would enable the suppression of IEDs. Two challenges are addressed: (a) the need to understand how actions taken by the military or other elements of national power may affect the behavior of a society that includes an adversary and non adversarial elements, and (b) the need to be able to capture and document data and knowledge about the cultural landscape of an area of operations that can be used to support the understanding of the key issues, beliefs, and reasoning concepts of the local culture so that individuals that are new to the region can quickly assimilate this knowledge and understanding. A Timed Influence Net was developed and analyzed.

The third case study illustrates the implementation of the theoretical developments presented in Chapter 3 to show how it is now possible to relax a number of limiting assumptions regarding causality (such as independence of causes) and include more realistic relationships between causes and effects. An East Timor scenario is used to illustrate the approach.

In Part III, methodologies for modeling adversary and coalition organizations are presented. In Chapter 6, a Petri Net based organization design approach is extended to include cultural constraints. The Lattice algorithm is used to design organizations subject to a number of structural and user defined constraints. These constraints are enhanced by introducing a set of cultural constraints based on Hofstede's dimensions. The approach is applied to an example where both Blue and Red organizations are modeled and the effect of cultural differences is highlighted. Finally, the approach is used to show how cultural attributes can be used in designing effective coalition organizations.

A key issue in modeling adversary organizations is the need to extract pertinent information about the adversary, such as interactions, activities, beliefs, and resources from a wide variety of unstructured textual data. In Chapter 7, a rapid ethnographic assessment procedure was used that moved from data to model through a semi-automated text analysis process. Central to this process is the AutoMap tool. AutoMap is based on network text analysis and so converts texts to networks of relations. Network Text Analysis is a set of methodologies for converting texts to graphs based on the theory that language and knowledge can be modeled as networks of words and relations such that meaning is inherent in the structure of that network. The semantic network is extracted first and then the meta-network composed of agents, resources, expertise, locations, activities, beliefs and organizations was obtained.

Understanding an organization's structure is critical when we attempt to understand, intervene in, or manage the organization. However, organizational structures in the real world often differ from their recognized formal structure, and sometimes its membership conceals the formal structure with various types of social interactions and communications. Furthermore, when the actual social interactions among the members of the group are observed, the observed social-network data are often noisy, and contain misleading and uncertain links. In Chapter 8, an approach for inferring the operational structure from the observed structure is proposed. The observed and the operational structure are likely to have distinct profiles, e.g., key personnel and clusters of individuals. This is because the operational is focused only on work related activities whereas the observed one is a concatenation of all activities, a snapshot of human endeavors. The approach is illustrated using data collected on a real-world, terrorist organization.

Social network simulation (SNS) is an emergent area of research that combines social network analysis and simulation, typically agent-based simulation. This area is often referred to as dynamic network analysis as much of the focus of the combined modeling approach is on how networks evolve, change, and adapt. Additionally SNS has a focus on how individual and group

learning and behavior is impacted by and impacts the changes in the networks in which the individuals are embedded. Frequently, in social network simulations, the social network and other networks, such as the knowledge network, and/or the individuals or “nodes” in the network are co-evolving as agents interact, learn, and engage in various activities. Cognitive and social factors combine to determine the level of information access that individuals/agents may have. Three different information access mechanisms: literacy, internet access, and newspaper readership were examined. In Construct, a dynamic network analysis tool, these access mechanisms affect whether agents can interact with a specific media and get information through a specific form. It is important to note that these mechanisms interact. For example, if an agent is illiterate and has a newspaper subscription, that agent may read the news articles but do so with error. On the other hand, if an agent is literate but does not have access to the internet, they still cannot read web-pages (and the literacy parameter has no effect). Construct and its application to simulating the adversary are described in Chapter 9.

Chapter 10 contains three applications of Dynamic Network Modeling. They illustrate that the key to reasoning about the adversary is taking social networks and embedding them within the spatio-temporal context. Organization theory and task processing analysis facilitate this embedding by providing the constraints and enablers on task-related activity.

In Part III of this report, recent research in multi-modeling and meta-modeling is described. No single model can capture the complexities of human behavior especially when interactions among groups with diverse social and cultural attributes are concerned. Each modeling language offers unique insights and makes specific assumptions about the domain being modeled. For example, social networks describe the interactions (and linkages) among group members but say little about the underlying organization and/or command structure. Similarly, organization models focus on the structure of the organization and the prescribed interactions but say little on the social/behavioral aspects of the members of the organization. Timed Influence net models describe cause-and-effect relationships among groups at a high level. In order to address the modeling and simulation issues that arise when multiple models are to interoperate, four layers need to be addressed. The first layer, Physical, i.e., Hardware and Software, is a platform that enables the concurrent execution of multiple models expressed in different modeling languages and provides the ability to exchange data and also to schedule the events across the different models. The second layer is the syntactic layer which ascertains that the right data are exchanged among the models. The Physical and Syntactic layers have been addressed through the development of two testbeds: C2 Wind Tunnel (C2WT) by Vanderbilt University in collaboration with UC-Berkeley and George Mason University (Appendix E) and SORASCS developed by CASOS at Carnegie Mellon University. Both have been used and developed further in this project.

Once the testbeds became available, a third problem needed to be addressed at the Semantic layer, where the interoperation of different models is examined to ensure that conflicting assumption in different modeling languages are recognized and form constraints to the exchange of data. In the fourth layer, the Workflow layer, valid combinations of interoperating models are considered to address specific applications. Different applications require different workflows. The use of multiple interoperating models is referred to as *multi-modeling* while the analysis of the validity of model interoperation is referred to as *meta-modeling*. Such an approach has been used in simulation mode or to explore the possible outcomes of proposed courses of action; it has not been used to predict outcomes.

In Chapter 11, the focus is on issues relating to the syntactic and semantic layers. In Chapter 12, an ontology based approach is used to analyze (deconstruct) modeling languages and identify common concepts, unique concepts, and contradictory concepts. An enriched ontology is obtained that then guides the interoperation of models by shedding light on which questions can be answered via a valid interoperation of two models and which questions would trigger the use of contradictory concepts. This type of result is key to developing valid workflows for using multiple models in addressing adversary modeling and complex policy issues. This work was not included in the original scope of work; it became apparent in the third year of the research effort that the simulation technology had reached a stage where multi-modeling became practical.

In Part IV, most of the research results were integrated by conducting a complex computational experiment. The issue addressed was deterrence – specifically determining Courses of Action for the US in encouraging de-escalation of an evolving crisis between two states that have strong ties to the US. In Chapter 13, the concept of deterrence, as it is evolving beyond nuclear deterrence between two peer states, is discussed with emphasis on cyber deterrence policy and strategy. Then in Chapter 14, a detailed case study based on an India-Pakistan crisis scenario is described. Multi-modeling was used extensively to represent India, Pakistan, the US central Command, and the US Pacific Command. Other state actors were also included. The results, presented in a day-long workshop, showed that the approaches taken to adversary modeling have promise and are implementable.

PART II: TIMED INFLUENCE NETS

Theory and Applications

Chapter 2: Course of Action Analysis in a Cultural Landscape
using Influence Nets

Chapter 3: Theory of Influence Networks

Chapter 4: Meta-model Driven Construction of Timed
Influence Nets

Chapter 5: Adversary Modeling Applications

Chapter 2

Course of Action Analysis in a Cultural Landscape Using Influence Nets

Lee W. Wagenhals and Alexander H. Levis

2.1 Introduction

In this chapter, two challenges are addressed: (a) the need to understand how actions taken by the military or other elements of national power may affect the behavior of a society that includes an adversary and non adversarial elements, and (b) the need to be able to capture and document data and knowledge about the cultural landscape of an area of operations that can be used to support the understanding of the key issues, beliefs, and reasoning concepts of the local culture so that individuals that are new to the region can quickly assimilate this knowledge and understanding.

The first challenge relates to capabilities that enable the analysis needed to conduct focused effects based planning and effects based operations. Models to support Effects Based Operations developed to date relate actions to effects on the adversary [1]. Such models can be quite effective in informing the comparison of alternative courses of action provided the relationships between potential actions and the effects are well understood. This depends on the ability to model an adversary's intent and his reactions and identifying his vulnerable points of influence. But as the nature of Blue's military operations goes well beyond the traditional major combat operations, there is the need to anticipate the effects of actions not only on the adversary (Red), but also on the local population which may support or oppose that adversary. Such support may depend in part on the actions taken by Blue.

The second challenge involves the need for new personnel to rapidly assimilate the local knowledge needed to analyze the local situation and to analyze and formulate the effects based plans and operations. Data about a culture exists in many forms and from many sources including historical reference documents, observations and reports by intelligence analysts, and unclassified (and unverified) sources such as the internet. The data is often incomplete and partially incorrect and includes contradictions and inconsistencies. Analysts, particularly those new to an area of operation who are responsible for formulating courses of action, are hard pressed to quickly develop the necessary understanding of the cultural factors that will affect the behavior of the adversary and the society in which it is embedded.

2.2 Timed Influence Nets

Several modeling techniques are used to relate actions to effects. With respect to effects on physical systems, engineering or physics based models have been developed that can predict the impact of various actions on systems and assess their vulnerabilities. When it comes to the cognitive belief and reasoning domain, engineering models are much less appropriate. The purpose of affecting the physical systems is to convince the leadership of an adversary to change its behavior, that is, to make decisions that it would not otherwise make. However, when an adversary is imbedded within a culture and depends upon elements of that culture for support, the effects of physical actions may influence not only the adversary, but the individuals and organizations

within the culture that can choose to support, be neutral, or oppose the adversary. Thus, the effects on the physical systems influence the beliefs and the decision making of the adversary and the cultural environment in which the adversary operates. Because of the subjective nature of belief and reasoning, probabilistic modeling techniques such as Bayesian Nets and their influence net cousin have been applied to these types of problems. Models created using these techniques can relate actions to effects through probabilistic cause and effect relationships. Such probabilistic modeling techniques can be used to analyze how the actions affect the beliefs and thus the support to and decisions by the adversary.

Influence Nets (IN) and their Timed Influence Nets (TIN) extension are abstractions of Probabilistic Belief Nets also called Bayesian Networks (BN) [2, 3], the popular tool among the Artificial Intelligence community for modeling uncertainty. BNs and TINs use a graph theoretic representation that shows the relationships between random variables. These random variables can represent various elements of a situation that can be described in a declarative statement, e.g., X happened, Y likes Z, etc.

Influence Nets are Directed Acyclic Graphs where nodes in the graph represent random variables, while the edges between pairs of variables represent causal relationships. While mathematically Influence Nets are similar to Bayesian Networks, there are some key differences. BNs suffer from the often intractable task of knowledge elicitation of conditional probabilities. To overcome this limitation, INs use CAST Logic [4, 5], a variant of Noisy-OR [6, 7], as a knowledge acquisition interface for eliciting conditional probability tables.

The modeling of the causal relationships in TINs is accomplished by creating a series of cause and effect relationships between some desired effects and the set of actions that might impact their occurrence in the form of an acyclic graph. The actionable events in a TIN are drawn as root nodes (nodes without incoming edges). Generally, desired effects, or objectives the decision maker is interested in, are modeled as leaf nodes (nodes without outgoing edges). In some cases, internal nodes are also effects of interest. Typically, the root nodes are drawn as rectangles while the non-root nodes are drawn as rounded rectangles. Figure 2.1 shows a partially specified TIN. Nodes B and E represent the actionable events (root nodes) while node C represents the objective node (leaf node). The directed edge with an arrowhead between two nodes shows the parent node promoting the chances of a child node being true, while the roundhead edge shows the parent node inhibiting the chances of a child node being true. The inscription associated with each arc shows the corresponding time delay it takes for a parent node to influence a child node. For instance, event B, in Fig. 2.1, influences the occurrence of event A after 5 time units.




Fig. 2.1 An Example Timed Influence Net (TIN)

Formally, a TIN is described by the following definition.

Definition 2.1: Timed Influence Net (TIN)

A TIN is a tuple $(V, E, C, B, D_E, D_V, A)$ where

V : set of Nodes,

E : set of Edges,

C represents causal strengths:

$E \rightarrow \{ (h, g) \text{ such that } -1 < h, g < 1 \},$

B represents Baseline / Prior probability: $V \rightarrow [0,1],$

D_E represents Delays on Edges: $E \rightarrow Z^+$

(where Z^+ represent the set of positive integers),

D_V represents Delays on Nodes: $V \rightarrow Z^+,$ and

A (input scenario) represents the probabilities associated with the state of actions and the time associated with them.

$A: R \rightarrow \{ ([p_1, p_2, \dots, p_n], [t_{11}, t_{12}], [t_{21}, t_{22}], \dots, [t_{n1}, t_{n2}]) \}$

such that $p_i = [0, 1], t_{ij} \rightarrow Z^*$ and $t_{i1} \leq t_{i2},$

$\forall i = 1, 2, \dots, n$ and $j = 1, 2$ where $R \subset V \}$

(where Z^* represent the set of nonzero positive integers)

The purpose of building a TIN is to evaluate and compare the performance of alternative courses of actions. The impact of a selected course of action on the desired effects is analyzed with the help of a probability profile. Consider the TIN shown in Fig. 2.1. Suppose the following input scenario is decided: actions B and E are taken at times 1 and 7, respectively. Because of the propagation delay associated with each arc, the influences of these actions impact event C over a period of time. As a result, the probability of C changes at different time instants. A probability profile draws these probabilities against the corresponding time line. The probability profile of event C is shown in Fig. 2.2.

To construct and use a TIN to support effects based operations, the following process has been defined.

1. Determine the set of desired and undesired effects expressing each as declarative statement that can be either true or false. For each effect, define one or more observable indicators that the effect has or has not occurred.
2. Build an IN that links, through cause and effect relationships, potential actions to the desired and undesired effects.




Fig 2.2 Probability Profile for Node C

Note that this may require defining additional intermediate effects and their indicators.

3. Use the IN to compare different sets of actions in terms of the probability of achieving the desired effects and not causing the undesired effects.
4. Transform the IN to a TIN by incorporating temporal information about the time the potential actions will occur and the delays associated with each of the arcs and nodes.
5. Use the TIN to experiment with different timings for the actions to identify the “best” COA based on the probability profiles that each candidate generates. Determine the time windows when observation assets may be able to observe key indicators so that assessment of progress can be made during COA execution.
6. Create a detailed execution plan to use the resources needed to carry out the COA and collect the information on the indicators.
7. Use the indicator data to assess progress toward achieving the desired effects.
8. Repeat steps 2 (or in some cases 1) through 7 as new understanding of the situation is obtained.

In building the IN, the modeler must assign values to the pair of parameters that show the causal strength (usually denoted as g and h values) for each directed link that connects pairs of nodes,. Each non-root node has an associated baseline probability that must be assigned by the modeler (or left at the default value of 0.5). It represents the probability that the random variable will be true in the absence of all modeled influences or causes. Each root node is given a prior probability, which is the initial probability that the random variable associated with the node (usually a potential action) is true.

When the modeler converts the IN into a TIN (step 4), each link is assigned a corresponding delay d (where $d \geq 0$) that represents the communication delay. Each node has a corresponding delay e (where $e \geq 0$) that represents the information processing delay. A pair (p, t) is assigned to each root node, where p is a list of real numbers representing probability values. For each probability value, a corresponding time interval is defined in t . In general, (p, t) is defined as

$$([p_1, p_2, \dots, p_n], [[t_{11}, t_{12}], [t_{21}, t_{22}], \dots, [t_{n1}, t_{n2}]]),$$

where $t_{i1} < t_{i2}$ and $t_{ij} > 0 \forall i = 1, 2, \dots, n$ and $j = 1, 2$

The last item is referred to as an input scenario, or sometimes (informally) as course of action.

To analyze the TIN (Step 5), the analyst selects the nodes that represent the effects of interest and generates probability profiles for these nodes. The probability profiles for different courses of action can then be compared.

Chapter 3

Theory of Influence Networks

Abbas K. Zaidi, Faisal Mansoor, P. Papantoni-Kazakos, Alexander H. Levis

3.1 Introduction

The easy access to domain-specific information and cost-effective availability of high computational power have changed the way people think about complex decision problems in almost all areas of application, ranging from financial markets to regional and global politics. These decision problems often require modeling of informal, uncertain and unstructured domains, to allow the evaluation of alternatives and available courses of actions by a decision maker. The past decade has witnessed an emergence of several modeling and analysis formalisms that target this need, the most popular one being represented by Probabilistic Belief Networks [3, 8], most commonly known as Bayesian Networks (BNs).

BNs model uncertain domains probabilistically, by presenting the network nodes as random variables. The arcs (or directed edges) in the network represent the direct dependency relationships between the random variables. The arrows on the edges depict the *direction* of the dependencies. The strengths of these dependencies are captured as conditional probabilities associated with the connected nodes in a network. A complete BN model requires specification of all conditional probabilities prior to its use. The number of conditional probabilities on a node in a BN grows exponentially with the number of inputs to the node, which presents a computational challenge, at times. A major problem in BNs is the specification of the required conditional probabilities, especially when either objective values of these probabilities cannot be provided by experts or there exist insufficient empirical data to allow for their reliable estimation, or when newly obtain information may change the structural topology of the network. Although a pair-wise cause and effect relationship between two variables of a domain is easier to establish or extract from a domain expert, a BN of the domain requires prior knowledge of all the influencing causes to an effect as well as their aggregate influence on the effect variable, where the measures of influences are conditional probability values. To demonstrate cases where BN modeling may be problematic, we identify the following situations of practical significance: (1) When new, previously unknown, affecting variables to some effect event arise, there are no algorithms allowing easy pertinent adaptation of conditional probabilities. (2) When the need arises to develop a consolidated BN from partial fragments of separate BNs, there are no algorithms that utilize the parameters of the fragments to calculate the parameters of the consolidated structure.

Recognizing the problems in the construction of BNs, especially regarding the specification of the involved conditional probabilities, Chang et al. [4] developed a formalism at George Mason University named Causal Strength (CAST) logic, as an intuitive and approximate language. The logic utilizes a pair of parameter values to represent conditional dependency between a pair of random variables, where these parameter values model assessed (by experts) mutual influences between an affecting and an affected event. The CAST logic approximates conditional probabilities via influence relationships by employing an influence

aggregation function. The approach provides the elicitation, update, reuse, and merge interface to an underlying BN, or multiple fragments of a BN, that only requires specification of individual influences between each pair of an affecting and an affected variables. The approach then combines these individual influences to calculate the aggregate effect of multiple affecting variables on an effect variable in terms of conditional probability values of a resulting BN. This pair-wise specification of influences provides us with the, albeit approximate, means to solve the three problems discussed earlier.

The CAST logic approach was later extended to represent relationships between events involved in network interconnections, as in BNs. The extension is basically a BN with conditional probabilities approximated via the use of influence parameters and was named Influence Nets (INs) [5, 9, 10, 11]. INs require an expert who specifies the influence parameter values and their interrelationships, as well as some a priori probabilities, all needed for the approximation of the pertinent conditional probabilities. As basically modified BNs, the objective of INs is to compute the probabilities of occurrence of sequential dependent events, and do not provide recommendations for actions. However, the probabilities of occurrence computed by the INs may be utilized by activation networks towards the evaluation and recommendation of actions [12].

BNs and INs are designed to capture *static* interdependencies among variables in a system. A situation where the impact of a variable takes some *time* to reach the affected variable(s) cannot be modeled by either one. In the last several years, efforts have been made to integrate the notion of time and uncertainty. Wagenhals et al. [12, 13, 14] have added a special set of temporal constructs to the basic formalism of INs. The INs with these additional temporal constructs are called Timed Influence Nets (TINs). TINs have been experimentally used in the area of Effects Based Operations (EBOs) for evaluating alternate courses of actions and their effectiveness to mission objectives in a variety of domains, e.g., war games [1, 15, 16, 17], and coalition peace operations [18], modeling adversarial behaviors [35], to name a few. The provision of time allows for the construction of alternate courses of action as timed sequences of actions or actionable events represented by nodes in a TIN [13, 15, 17]. A number of analysis tools have been developed over the years for TIN models, to help an analyst update beliefs [19, 20, 21, 22, 23] represented as nodes in a TIN, to map a TIN model to a Time Sliced Bayesian Network for incorporating feedback evidence, to determine best course of actions for both timed and un-timed versions of Influence Nets [24, 25] and to assess temporal aspects of the influences on objective nodes [26, 27].

The existing developments of INs and TINs suffer from a number of deficiencies: they do not represent scenarios encompassing dependent conditioning events and they utilize a priori probabilities inconsistently, in violation of the Bayes Rule and the Theory of Total Probability. The motivation behind the work presented in this paper is to address these shortcomings of INs and TINs by developing a correct analytical framework for the design and analysis of influences on some critical effects due to a set of external affecting events. We present a comprehensive theory of Influence Networks, which is free of restrictive independence assumptions, which is consistently observing the Bayes Rule and the Theorem of Total Probability. In this theory, we are concerned with the evaluation of cause-effect relationships between interconnected events. In particular, if the status of some event B is affected by the status of a set of events, A_1 to A_n , we are interested in a qualification and quantification of this effect. We first graph the relationships between events B and A_1 to A_n in a

network format, as in Fig. 3.1 below, with each event being a node, with arcs indicating relationships and with arrows representing the cause-effect directions. This graphical representation is identical to that used in BNs.




Fig. 3.1 Cause-Effect Relationships

Given the graph of Fig. 3.1, we next decide the metric to be used for the quantification of the effects of events A_1 to A_n on event B . As in BNs, modeling each of the involved events as binary random variables, we use conditional probabilities as effect metrics: in particular, we use the probabilities that event B occurs, given each of the 2^n scenarios regarding the occurrence or nonoccurrence of each one of the events A_1 to A_n .

Upon the decision to use conditional probabilities as the effect metrics, the issue of their computation arises. In most realistic scenarios, there exist insufficient amount of data for the reliable estimation of these probabilities. Instead, some influence indicators may be provided by experts. In the example of Fig. 3.1, for instance, for each one of the 2^n scenarios regarding the occurrence or nonoccurrence of each one of the events A_1 to A_n , an expert may provide a number between -1 and 1 , to reflect his assessment as to the effect of the above scenario on the occurrence of event B . The latter number is named **influence constant**. The objective at this point is to utilize the so provided influence constants for the approximate evaluation and computation of the required conditional probabilities, in a mathematically correct and consistent fashion. These conditional probabilities are subsequently utilized for the probabilistic evaluation of event occurrences in a network of events, giving rise to an Influence Network (IN). In different terms, a IN is a BN whose conditional probabilities are computed via the use of influence constants. The term IN should not be confused with a similarly named formalism called Influence Diagrams [28, 29, 30, 31]. Unlike INs, an Influence Diagram (ID) has different types of nodes (i.e., decision nodes, chance nodes, and utility nodes) and different types of influences (i.e., arcs between the nodes); and the decisions in an ID are assumed to have a certain precedence relationship among them. The IDs can be considered a BN extended with a utility function, while a IN, as noted above, is a special instance of a BN whose conditional probabilities are computed via the use of influence constants and which uses a set of special purpose algorithms for calculating the impact of a set of external affecting events on some desired effect/objective node.

Frequently, in several realistic scenarios, assessments of event occurrences may be needed at times when the status of all affecting events may not be known, while such assessments require sequential adaptation, as the status of more affecting events are revealed. For example, in Fig. 1, the evaluation of the probability of event B may be needed at times

when the status of only some of the events A are known, while this probability need to be subsequently adapted when the status of the remaining A events become known. Such sequential adaptations require pertinent sequential computation methodologies for the approximation of conditional probabilities via influence constants and give rise to Time Influence Networks (TINs). We present two different temporal models for the sequential computation of conditional probabilities in a Timed Influence Nets. This enhances the capabilities of the Timed Influence Nets in modeling domains of interest with different time characteristics.

The organization of the paper is as follows: In section 3.2, we present the theoretical formalization and derive initial relationships. In section 3.3, we derive the dynamic programming evolution of the influence constants. In section 3.4, we examine the case where in the generic model, the affecting events are mutually independent, where in section 3.5, the case where the latter events form a Markov chain is examined. In section 3.6, temporal considerations are presented. In section 3.7 we discuss decision model selection and testing. In section 3.8, special forms of the influence constants are discussed. In Section 3.9, we discuss evaluation metrics. In section 3.10, the experimental setup is laid out, while in the final section, 3.11, conclusions are drawn.

3.2 Initial Modeling and Relationships

In this section, we formalize our approach for the development of INs and TINs.

Let us consider an event B being potentially affected by events $\{A_i\}_{1 \leq i \leq n}$. In particular, we are interested in the effect the presence or absence of any of the events in the set $\{A_i\}_{1 \leq i \leq n}$ may have on the occurrence of event B .

Let us first define:

X_1^n : An n -dimensional binary random vector whose j^{th} component is denoted X_j ,
where $X_j = 1$; if the event A_j is present, and $X_j = 0$; if the event A_j is absent.

We will denote by x_1^n realizations or values of the random vector X_1^n . A given realization x_1^n of the binary vector X_1^n describes precisely the status of the set $\{A_i\}_{1 \leq i \leq n}$ of events, regarding which events in the set are present. We name the vector X_1^n , the **status vector** of the affecting events. To quantify the effects of the status vector X_1^n on the event B , we define the **influence constant** $h_n(x_1^n)$ via the following quantitative properties:

$$h_n(x_1^n) = \begin{cases} 1 & \text{; if given } n \text{ affecting events, given the status} \\ & \text{vector } x_1^n, \text{ event } B \text{ occurs surely} \\ -1 & \text{; if given } n \text{ affecting events, given the status} \\ & \text{vector } x_1^n, \text{ the nonoccurrence of event } B \text{ is sure} \\ 0 & \text{; if given } n \text{ affecting events, given the status} \\ & \text{vector } x_1^n, \text{ the occurrence of event } B \text{ is unaffected} \end{cases} \quad (3.1)$$

Let $P(B | x_1^n)$ denote the probability of occurrence of event B, given the status vector x_1^n . Then, the quantitative definition of the influence constant $h_n(x_1^n)$ in (3.1) can be rewritten as follows, where $P(B)$ denotes the unconditional probability of occurrence of the event B.

$$P(B | x_1^n) = \begin{cases} 1 & ; \text{ if } h_n(x_1^n) = 1 \\ P(B) & ; \text{ if } h_n(x_1^n) = 0 \\ 0 & ; \text{ if } h_n(x_1^n) = -1 \end{cases} \quad (3.2)$$

We now extend the definition of all values in $[1, -1]$ of the influence constant, via linear interpolation from (3.2). In particular, we define the influence constant via its use to determine the derivation of the conditional probability $P(B | x_1^n)$ from the unconditional probabilities $P(B)$, where this derivation is derived via linear interpolation from (3.2). We thus obtain.

$$P(B | x_1^n) = \begin{cases} P(B) + h_n(x_1^n)[1 - P(B)] & ; \text{ if } h_n(x_1^n) \in [0, 1] \\ P(B) + h_n(x_1^n)P(B) & ; \text{ if } h_n(x_1^n) \in [-1, 0] \end{cases} \quad (3.3)$$

Defining $\text{sgn } \gamma = \begin{cases} 1 & ; \text{ if } \gamma \geq 0 \\ 0 & ; \text{ if } \gamma < 0 \end{cases}$, we can finally write (3.3) as follows

$$P(B | x_1^n) = P(B) \{1 + h_n(x_1^n)[1 - P(B)]P^{-1}(B)\}^{\text{sgn } h_n(x_1^n)} \times \{1 + h_n(x_1^n)\}^{1 - \text{sgn } h_n(x_1^n)} \quad (3.4)$$

At this point, we present a formal definition of INs and TINs.

Definition 3.1: An Influence Network (IN) is a Bayesian Network mapping conditional probabilities $P(B | x_1^n)$ via the utilization of influence constants as in (3.4). Formally, an Influence Net is a tuple $(\mathbf{V}, \mathbf{E}, \mathbf{C}, \mathbf{A}, \mathbf{B})$, with $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ representing a directed-acyclic graph satisfying the Markov condition (as in BN), where:

- V:** set of nodes representing binary random variables,
- E:** set of edges representing causal influences between nodes,
- C:** set of causal strengths: $E \rightarrow \{[h_1^{(i)}(x_i = 1), h_1^{(i)}(x_i = 0)] \text{ such that } h_i \text{'s} \in [-1, 1]\}$,
- A:** a subset of **V** representing *external* affecting events $\{A_i\}_{1 \leq i \leq n}$ and a status of the corresponding vector X_1^n ,
- B:** Probability distribution of the status vector X_1^n corresponding to the external affecting events $\{A_i\}_{1 \leq i \leq n}$.

A Timed Influence Network (TIN) adds two temporal parameters to the definition of a IN. Formally, a TIN is a tuple $(\mathbf{V}, \mathbf{E}, \mathbf{C}, \mathbf{D}, \mathbf{A}_T, \mathbf{B})$, where **V**, **E**, **C**, and **B** are as defined for INs;

D: set of temporal delays on edges: $\mathbf{E} \rightarrow \mathbf{N}$,

A_T: same as **A** with the addition that the status of each external affecting event is *time tagged* representing the time of realization of its status. In the IN/TIN literature [12, 13, 15, 16, 17, 18, 25], A_T is also referred to as a Course of Action (COA). A COA is, therefore, a time-sequenced collection of external affecting events and their status.

Returning to the influence constant notion, we note that there exist 2^n distinct values of the status vector x_1^n ; thus, there exist 2^n distinct values of the influence constant $h_n(x_1^n)$ as well as of the conditional probabilities in (3.4). In the case where the cardinality of the set $\{A_i\}_{1 \leq i \leq n}$ is one, the influence constant $h_1(x_1)$ equals the constant h in [5]; if $x_1 = 1$ and equals the constant g in [5]; if $x_1 = 0$.

We now proceed with a definition which will lead to a mathematically correct relationship between influence constants and unconditional probabilities.

Definition 3.2: A IN or TIN model is **consistent** if it observes the Bayes Rule.

Let $P(x_1^n)$ denote the probability of the status vector X_1^n at the value x_1^n . We can then express the following simple lemma.

Lemma 3.1

Let the influence constant $h_n(x_1^n)$ be accepted as reflecting accurately the relationship between the affecting events $\{A_i\}_{1 \leq i \leq n}$ and event B. Then the IN or TIN model is consistent iff:

$$\sum_{x_1^n} P(x_1^n) \{1 + h_n(x_1^n)[1 - P(B)]P^{-1}(B)\}^{\text{sgn } h_n(x_1^n)} \{1 + h_n(x_1^n)\}^{1 - \text{sgn } h_n(x_1^n)} = 1 \quad (3.5)$$

or

$$\sum_{x_1^n} P(x_1^n) h_n(x_1^n) \{[1 - P(B)]P^{-1}(B)\}^{\text{sgn } h_n(x_1^n)} = 0$$

Proof: Substituting expression (3.4) in the the Bayes' Rule, $P(B) = \sum_{x_1^n} P(x_1^n)P(B | x_1^n)$, we obtain (3.5).

Expression (3.5) relates the influence constant $h_n(x_1^n)$ to the unconditional probabilities of event B and the status vector X_1^n . This relationship is necessary if the influence constant is accepted as accurately representing the conditional probability $P(B | x_1^n)$ in (3.3). Generally, the influence constant is selected based on a system design assessment provided by experts, while the a priori probabilities $P(x_1^n)$ are accepted to accurately represent the actual model.

Summary

Given the events in Fig. 3.1, given well-established a priori probabilities of the cause events, given the influence constants, the cause-effect conditional probabilities are expressed as follows:

$$P(B | x_1^n) = \begin{cases} a + h_n(x_1^n)[1 - a]; & \text{if } h_n(x_1^n) \in [0, 1] \\ a + h_n(x_1^n)a & ; \text{if } h_n(x_1^n) \in [-1, 0] \end{cases}$$

where

$$a = P(B) = \left[\sum_{x_1^n: \text{sgn } h_n(x_1^n)=1} P(x_1^n) h_n(x_1^n) \right] \left[\sum_{x_1^n} P(x_1^n) |h_n(x_1^n)| \right]^{-1}$$

Influence nets thus utilize expert-provided subjective influence constants, in conjunction with well-established objective a priori probabilities of cause events, to generate conditional probabilities of effect events.

3.3 Evolution of the influence Constant

In section 3.2, we derived the relationship between the conditional probability of event B, and the status x_1^n of its affecting events $\{A_i\}_{1 \leq i \leq n}$, via the influence constant $h_n(x_1^n)$. This relationship is based on the assumption that $\{A_i\}_{1 \leq i \leq n}$ is the maximum set of events affecting event B and that the value x_1^n of the status vector is given. In this section we investigate the case where the status of some of the affecting events may be unknown. Towards this direction, we derive a dynamic programming relationship between the influence constants $h_n(x_1^n)$ and $h_{n-1}(x_1^{n-1})$, where $h_{n-1}(x_1^{n-1})$ is the constant corresponding to the case where the status of the affecting event A_n is unknown. We express a lemma whose proof is in Appendix A of this report. The proof is based on the observation of the Bayes' Rule and the Theorem of Total Probability.

Lemma 3.2

Let the probability $P(B)$ be as in Section II and let $P(x_n | x_1^{n-1})$ denote the probability of the value of the last bit in the status vector X_1^n being x_n , given that the reduced status vector value is x_1^{n-1} . Then, the influence constant $h_{n-1}(x_1^{n-1})$ is given as a function of the influence constant $h_n(x_1^n)$, as shown below.

$$h_{n-1}(x_1^{n-1}) = \begin{cases} Q_n & ; \quad Q_n \in [-1, 0] \\ P(B)[1 - P(B)]^{-1} Q_n & ; \quad Q_n \in [0, P^{-1}(B) - 1] \end{cases} \quad (3.6)$$

where

$$Q_n = \sum_{x_n=0,1}^{\Delta} P(x_n | x_1^{n-1}) \{h_n(x_1^n)\} \{[1 - P(B)]P^{-1}(B)\}^{\text{sgn } h_n(x_1^n)} \quad (3.7)$$

We note that the influence constants are deduced from the same constants of higher dimensionality, as shown in Lemma 3.2. In accordance, conditional probabilities of the event B are produced from the deduced influence constants, via expression (3.4), as:

$$P(B | x_1^{n-1}) = P(B) \{1 + h_{n-1}(x_1^{n-1})[1 - P(B)]P^{-1}(B)\}^{\text{sgn } h_{n-1}(x_1^{n-1})} \times \{1 + h_{n-1}(x_1^{n-1})\}^{1-\text{sgn } h_{n-1}(x_1^{n-1})} \quad (3.8)$$

It is important to note that in the dynamic programming evolution of the influence constants $h_n(x_1^n)$, as well as in the evolution of the conditional probabilities in (3.7), knowledge of the joint probability $P(x_1^n)$ is assumed. This reflects a conjecture by the system designer, based on his /her previous experience regarding the a priori occurrence of the affecting events $\{A_i\}_{1 \leq i \leq n}$. Thus the probability $P(x_1^n)$ used for the construction exhibited by Lemma 3.2 is a design probability and it may not coincide with the actual probabilities of the status vector X_1^n . When full scale dependence of the components of the status vector X_1^n is incorporated within the design probability $P(x_1^n)$, then the relationship between the different dimensionality influence constants is that reflected by Lemma 3.2 and is of dynamic programming nature. In the case where the design probability $P(x_1^n)$ generically reflects either a Markov Chain of events or mutually independent events, then the relationships between the different dimensionality influence constants may be also of recursive nature. The cases of Markovian or independent affecting events, as modeled by the system designer, are examined in sections 3.4 and 3.5.

3.4 The Case of Independent Affecting Events

In this section, we consider the special case where the affecting events $\{A_i\}_{1 \leq i \leq n}$ are assumed to be generically mutually independent. Then, the components of the status vector X_1^n are mutually independent, and:

$$P(x_1^n) = \prod_{i=1}^n P(x_i) \quad ; \quad P(x_1^n | B) = \prod_{i=1}^n P(x_i | B) \quad (3.9)$$

Let us denote by $h_1^{(i)}(x_i)$ the influence constant corresponding to the effect of the event A_i on the occurrence of the event B, when event A_i acts in isolation and when the status value of the event is x_i . Then, from expression (3.4) in section 3.3, we have:

$$P(B | x_i) = P(B) \{1 + h_1^{(i)}(x_i)[1 - P(B)]P^{-1}(B)\}^{\text{sgn } h_1^{(i)}(x_i)} \bullet \{1 + h_1^{(i)}(x_i)\}^{1-\text{sgn } h_1^{(i)}(x_i)} \quad (3.10)$$

We now express a lemma whose proof is in Appendix A.

Lemma 3.3

Let the events $\{A_i\}_{1 \leq i \leq n}$ that affect event B be assumed to be generically mutually independent. Then

$$P(B | x_1^n) = P(B) \prod_{i=1}^n \left\{ 1 + h_1^{(i)}(x_i) [1 - P(B)] P^{-1}(B) \right\}^{\text{sgn } h_1^{(i)}(x_i)} \bullet \left\{ 1 + h_1^{(i)}(x_i) \right\}^{1 - \text{sgn } h_1^{(i)}(x_i)} \quad (3.11)$$

Via the same logic as that in the last part in the proof of Lemma 2, we can show the result expressed in the corollary below.

Corollary 3.1

When the affecting events are assumed to be generically mutually independent then, the influence constant $h_n(x_1^n)$ is given as a function of the single event influence constants $\{h_1^{(i)}(x_i)\}_{1 \leq i \leq n}$, as follows:

$$h_n(x_1^n) = \begin{cases} R_n - 1 & ; \quad \text{if } R_n \in [0, 1] \\ P(B) [1 - P(B)]^{-1} [R_n - 1] & ; \quad \text{if } R_n \in [1, P^{-1}(B)] \end{cases} \quad (3.12)$$

where

$$R_n = \prod_{i=1}^n \left\{ 1 + h_1^{(i)}(x_i) [1 - P(B)] P^{-1}(B) \right\}^{\text{sgn } h_1^{(i)}(x_i)} \times \left\{ 1 + h_1^{(i)}(x_i) \right\}^{1 - \text{sgn } h_1^{(i)}(x_i)} \quad (3.13)$$

The sequence of expressions $\{R_i\}_{1 \leq i \leq n}$ in (3.13) is clearly recursively generated and the conditional probability $P(B | x_1^n)$ is given by $h_n(x_1^n)$ as in (3.4) in section 3.2.

We note that the consistency condition in Lemma 3.1, section 3.2 reduces in a straight forward fashion and by construction to the following condition here:

$$\sum_{x_i=0,1} P(x_i) \left\{ 1 + h_1(x_i) [1 - P(B)] P^{-1}(B) \right\}^{\text{sgn } h_1(x_i)} \left[1 + h_1(x_i) \right]^{1 - \text{sgn } h_1(x_i)} = 1; \forall i$$

or

$$\sum_{x_i=0,1} P(x_i) h_1(x_i) \left\{ [1 - P(B)] P^{-1}(B) \right\}^{\text{sgn } h_1(x_i)} = 0; \forall i$$

3.5 The Case of A Markov Chain of Affecting Events

In this section, we consider the case where the affecting events $\{A_i\}_{1 \leq i \leq n}$ are assumed to form generically a Markov Chain. In particular, we assume that the design probabilities $P(x_1^n | B)$ and $P(x_1^n)$ are such that:

$$\begin{aligned}
P(x_1^n | B) &= \prod_{i=1}^n P(x_i | x_{i-1}, B) \\
P(x_1^n) &= \prod_{i=1}^n P(x_i | x_{i-1})
\end{aligned} \tag{3.14}$$

where

$$P(x_1 | x_0, B) \stackrel{\Delta}{=} P(x_1 | B) \text{ and } P(x_1 | x_0) \stackrel{\Delta}{=} P(x_1)$$

We denote by $h_1^{(1)}(x_1)$ the influence constant corresponding to the effect of the event A_1 on the occurrence of the event B, when the status value of A_1 , is given by x_1 . We denote by $h_2^{(i,i+1)}(x_i, x_{i+1})$ the influence constant corresponding to the effect of the events A_i and A_{i+1} on the occurrence of the event B, when the status values of the (A_i, A_{i+1}) pair are given by (x_i, x_{i+1}) . Then, via (3.4) in section 3.2, we have

$$P(B | x_1) = P(B) \{1 + h_1^{(1)}(x_1)[1 - P(B)]P^{-1}(B)\}^{\text{sg } nh_1^{(1)}(x_1)} \times \{1 + h_1^{(1)}(x_1)\}^{1 - \text{sg } nh_1^{(1)}(x_1)} \tag{3.15}$$

$$\begin{aligned}
P(B | x_i, x_{i+1}) &= P(B) \{1 + h_2^{(i,i+1)}(x_i, x_{i+1})[1 - P(B)]P^{-1}(B)\}^{\text{sg } nh_2^{(i,i+1)}(x_i, x_{i+1})} \times \\
&\quad \{1 + h_2^{(i,i+1)}(x_i, x_{i+1})\}^{1 - \text{sg } nh_2^{(i,i+1)}(x_i, x_{i+1})}; i = 1
\end{aligned} \tag{3.16}$$

We now express a lemma whose proof is in the Appendix.

Lemma 3.4

Let the affecting events $\{A_i\}_{1 \leq i \leq n}$ be assumed to generically form a Markov Chain; thus, $P(x_1^n)$ is assumed to satisfy the equation in (3.14). Then,

$$\begin{aligned}
P(B | x_1^n) &= P(B) \{1 + h_1^{(1)}(x_1)[1 - P(B)]P^{-1}(B)\}^{\text{sg } nh_1^{(1)}(x_1)} \times \{1 + h_1^{(1)}(x_1)\}^{1 - \text{sg } nh_1^{(1)}(x_1)} \times \\
&\quad \prod_{i=2}^n \frac{\{1 + h_2^{(i,i-1)}(x_i, x_{i-1})[1 - P(B)]P^{-1}(B)\}^{\text{sg } nh_2^{(i,i-1)}(x_i, x_{i-1})} \times \{1 + h_2^{(i,i-1)}(x_i, x_{i-1})\}^{1 - \text{sg } nh_2^{(i,i-1)}(x_i, x_{i-1})}}{\{1 + h_1^{(i-1)}(x_{i-1})[1 - P(B)]P^{-1}(B)\}^{\text{sg } nh_1^{(i-1)}(x_{i-1})} \times \{1 + h_1^{(i-1)}(x_{i-1})\}^{1 - \text{sg } nh_1^{(i-1)}(x_{i-1})}} \\
&\stackrel{\Delta}{=} P(B)W_n
\end{aligned} \tag{3.17}$$

where,

$$h_1^{(i)}(x_i) = \begin{cases} Q_{i,i+1} - 1 & ; \text{ if } Q_{i,i+1} \in [0,1] \\ P(B)[1 - P(B)]^{-1}[Q_{i,i+1} - 1] & ; \text{ if } Q_{i,i+1} \in [1, P^{-1}(B)] \end{cases} \tag{3.18}$$

$$Q_{i,i+1}^{\Delta} = \sum_{x_{i+1}=0,1} P(x_{i+1} | x_i) \left\{ 1 + h_2^{(i,i+1)}(x_i, x_{i+1}) [1 - P(B)] P^{-1}(B) \right\}^{\text{sgn} h_2^{(i,i+1)}(x_i, x_{i+1})} \times \left\{ 1 + h_2^{(i,i+1)}(x_i, x_{i+1}) [1 + P(B)] P^{-1}(B) \right\}^{1 - \text{sgn} h_2^{(i,i+1)}(x_i, x_{i+1})} \quad (3.19)$$

As with Corollary 3.1 in section 3.4, we can express the corollary below, in a direct fashion.

Corollary 3.2

When the affecting events $\{A_i\}_{1 \leq i \leq n}$ are assumed to generically form a Markov Chain, depicted by the expression in (3.14), then, the influence constant $h_n(x_1^n)$ is given as a function of the influence constants $\{h_1^{(i)}(x_i)\}$ and $\{h_2^{(i,i-1)}(x_i, x_{i-1})\}$, as below, where W_n is defined in (3.17).

$$h_n(x_1^n) = \begin{cases} W_n - 1 & ; \quad \text{if } W_n \in [0, 1] \\ P(B)[1 - P(B)]^{-1}[W_n - 1]; & \text{if } W_n \in [1, P^{-1}(B)] \end{cases} \quad (3.20)$$

The sequence $\{W_i\}_{1 \leq i \leq n}$ in (17) is clearly recursively expressed; thus, $h_n(x_1^n)$ is recursively evolving. The consistency condition in Lemma 3.1, section 3.2, takes here the following form, by construction.

$$\sum_{x_i=0,1} \sum_{x_{i-1}=0,1} P(x_i | x_{i-1}) \left\{ 1 + h_2^{(i,i-1)}(x_i, x_{i-1}) [1 - P(B)] P^{-1}(B) \right\}^{\text{sgn} h_2^{(i,i-1)}(x_i, x_{i-1})} \times \left\{ 1 + h_2^{(i,i-1)}(x_i, x_{i-1}) \right\}^{1 - \text{sgn} h_2^{(i,i-1)}(x_i, x_{i-1})} = 1; \quad \forall i$$

3.6 Temporal Extension

In sections 3.2 and 3.3, we presented our theoretical foundation for the development of INs and TINs, while in sections 3.4 and 3.5, we focused on the special cases of independent and Markovian affecting events. In this section, we focus on the formalization of the temporal issues involved in the development of TINs. In particular, we are investigating the dynamics of the relationship of the affecting events $\{A_i\}_{1 \leq i \leq n}$ to the affected event B, when the status of the former events are learned asynchronously in time. Without lack in generality – to avoid cumbersome notation – let the affecting events $\{A_i\}_{1 \leq i \leq n}$ be ordered in the order representing the time when their status become known. That is, the status of events A_1 is first known, then that of event A_2 , and so on. In general, the status of event A_k becomes known after the status of the events A_1, \dots, A_{k-1} are known, and this knowledge becomes available one event at the time.

Let us assume that the considered system model implies full dependence of the components of the status vector X_1^n . Then, the influence constants $\{h_i(x_1^n)\}_{1 \leq i \leq n-1}$ are first pre-computed via the dynamic programming expression in Lemma 3.2, section 3.2, utilizing the pre-selected a pri-

ori probabilities $P(x_1^n)$ that are part of the given system parameters. The above influence constants can be recursively computed if the adopted system model implies either generically independent affecting events or affecting events that generically form a Markov Chain, as shown in sections 3.4 and 3.5.

Let T_0 denote the time when the computation of the system dynamics starts. Let T_1 denote the time when the status of event A_1 , becomes known. Let T_k ; $1 \leq k \leq n$ denote the time when the status of event A_k becomes known. Then at time T_k , the conditional probabilities $P(B | x_1^k)$ are computed via expression (3.4), Section II, as,

$$P(B | x_1^k) = P(B) \left\{ 1 + h_k(x_1^k) [1 - P(B)] P^{-1}(B) \right\}^{\text{sgn } h_k(x_1^k)} \times \left\{ 1 + h_k(x_1^k) \right\}^{1 - \text{sgn } h_k(x_1^k)} \quad (3.21)$$

where the probability $P(B)$ is computed via the consistency condition (5).

As the knowledge about the status of the affecting events unravels, the conditional probabilities of event B in (3.21) evolve dynamically in time and finally converge to the probability $P(B | x_1^n)$ at time T_n , when the status of all the affecting events become known.

It is important to point out that the conditional probability in (3.21) is sensitive to the time ordering of the affecting events. That is, for the same value x_1^k of a partial affecting vector, but different time ordering of events, different conditional probabilities values of the affected event B arise. Thus, the order by which the status of the affecting events become known is crucial in the evaluation of the conditional probabilities of event B.

3.7 Selection and Testing of the Decision Model

Model Selection

As we have discussed earlier, the unconditional probabilities $P(x_1^n)$ as well as the influence constant $h_n(x_1^n)$ are design parameters that may not represent the actual parameters correctly. Furthermore, as discussed in section 3.2, the design parameters must be **consistent**, where consistency is represented by the satisfaction of condition (3.5) in Lemma 3.1. Condition (3.5) can be rewritten as follows, in a straightforward fashion.

$$[1 - P(B)] \sum_{x_1^n : \text{sgn } h_n(x_1^n) = 1} P(x_1^n) h_n(x_1^n) = P(B) \sum_{x_1^n : \text{sgn } h_n(x_1^n) = 0} P(x_1^n) h_n(x_1^n) \quad (3.22)$$

which gives:

$$P(B) = \left[\sum_{x_1^n : \text{sgn } h_n(x_1^n) = 1} P(x_1^n) h_n(x_1^n) \right] \left[\sum_{x_1^n} P(x_1^n) h_n(x_1^n) \right]^{-1} ; \text{ when } \sum_{x_1^n} P(x_1^n) h_n(x_1^n) \neq 0 \quad (3.23)$$

Example: Let us consider the case where the only affecting event for B is A_i .

Let
$$P(A_i) \stackrel{\Delta}{=} P(X_i = 1) = p,$$

where then,

$$P(A_i^c) \stackrel{\Delta}{=} P(X_i = 0) = 1 - p.$$

Define h and g as in [5] and let $P(B)$ be what has been called in [5] base probability for the event B. Then, due to (3.22) the above parameters must satisfy the following equation(s):

$$\begin{aligned} \text{either } [1 - P(B)]ph &= P(B)(1 - p)|g| \quad ; \text{ if } h > 0 \text{ and } g < 0 \\ \text{or } [1 - P(B)](1 - p)g &= P(B)p|h| \quad ; \text{ if } h < 0 \text{ and } g > 0 \end{aligned}$$

no other h and g combinations are acceptable. Note that parameters h and g in [5] map to $h_1^{(i)}(x_i = 1)$ and $h_1^{(i)}(x_i = 0)$, respectively, in Definition 3.1, section 3.2.

When new information about the a priori probability $P(x_1^n)$ is obtained, then, $P(B)$ and/or $h_n(x_1^n)$ need to be accordingly adjusted to satisfy the condition in (22). We note that the latter condition involves a number of free parameters; thus even specification of the probabilities $P(B)$ and $P(x_1^n)$ does not specify uniquely the values of the influence constant $h_n(x_1^n)$. Naturally, specification of $P(x_1^n)$ and $h_n(x_1^n)$ uniquely determines the probability $P(B)$, however, as in (3.23).

In the case that the assumed system design model implies generically independent affecting events $\{A_i\}_{1 \leq i \leq n}$, then, for consistency the probability $P(B)$, the probability $P(x_1^n) = \prod_{i=1}^n P(x_i)$ of the status vector and the influence constants $\{h_1^{(i)}(x_i)\}$ are constraint to satisfy the condition:

$$\sum_{x_i=0,1} P(x_i) \{1 + h_1(x_i)[1 - P(B)]P^{-1}(B)\}^{\text{sgn } h_1(x_i)} \{1 + h_1(x_i)\}^{1 - \text{sgn } h_1(x_i)} = 1; \forall i \quad (3.24)$$

Or

$$\sum_{x_i=0,1} P(x_i) h_1(x_i) \{[1 - P(B)]P^{-1}(B)\}^{\text{sgn } h_1(x_i)} = 0; \forall i$$

Model Testing

Since the “consistency” constraints allow for a number of free parameters, we will focus on the influence constant $h_n(x_1^n)$ as the constant to be tested, when information about the probabilities of the events $\{A_i\}_{1 \leq i \leq n}$ and B is obtained. Thus, model testing will involve comparison of the $P(x_1^n)$ and $P(B)$ probabilities assumed in the model with those computed, to test the validity of the assumed influence constant. When the computed $P(x_1^n)$ and $P(B)$ values do not satisfy equa-

tion (23) for the assumed $h_n(x_1^n)$, then a non valid model is declared and a new influence constant $h_n(x_1^n)$ is sought, in satisfaction of the consistency condition in (3.23).

3.8 Some Special Influence Constants

As noted at the end of section 3.7, the influence constant is a important component of the system model: the appropriate choice of this constant needs to be carefully thought out, to accurately reflect the interleaving of partial influences. In this section, we study some specific influence constants, $h_n(x_1^n)$. In particular, we study such constants that are specific analytic functions of the one-dimensional components $h_i(x_i)$; $1 \leq i \leq n$. We note that we are not mapping the $\{h_i(x_i)\}_{1 \leq i \leq n}$ constants onto conditional probabilities $\{P(B | x_i)\}_{1 \leq i \leq n}$. Instead, we are using the constants $\{h_i(x_i)\}_{1 \leq i \leq n}$ to construct a global $h_n(x_1^n)$ influence constant; it is the latter constant which is mapped onto the conditional probability $P(B | x_1^n)$, as in section 3.2.

The $h_n(x_1^n)$ corresponding to the CAST logic

The influence constant presented below is that used by the CAST logic in [4, 5, 9, 10, 11]. In the present case, given the constants $\{h_i^{(i)}(x_i)\}_{1 \leq i \leq n}$ the global influence constant, $h_n(x_1^n)$, is defined as follows

$$h_n(x_1^n) = \left[\prod_{i: h_i(x_i) < 0} (1 - |h_i^{(i)}(x_i)|) - \prod_{i: h_i(x_i) > 0} (1 - |h_i^{(i)}(x_i)|) \right] \times \left[\max \left(\prod_{i: h_i(x_i) < 0} (1 - |h_i^{(i)}(x_i)|), \prod_{i: h_i(x_i) > 0} (1 - |h_i^{(i)}(x_i)|) \right) \right]^{-1} \quad (3.25)$$

In agreement with the results in section 3.2, and via (5) in Lemma 1, the global constants $h_n(x_1^n)$ and the probabilities $P(x_1^n)$ and $P(B)$ must satisfy the consistency condition

$$\sum_{x_1^n} P(x_1^n) \{1 + h_n(x_1^n)[1 - P(B)]P^{-1}(B)\}^{\text{sgn } h_n(x_1^n)} \{1 + h_n(x_1^n)\}^{1 - \text{sgn } h_n(x_1^n)} = 1 \quad (3.26)$$

Via (4), the conditional probabilities $P(B | x_1^n)$ are then given, by the following expression:

$$P(B | x_1^n) = P(B) \{1 + h_n(x_1^n)[1 - P(B)]P^{-1}(B)\}^{\text{sgn } h_n(x_1^n)} \times \{1 + h_n(x_1^n)\}^{1 - \text{sgn } h_n(x_1^n)} \quad (3.27)$$

For maintaining the consistency condition in (3.26), the conditional probability $P(B | x_1^{n-1})$ is defined via the influence constant $h_{n-1}(x_1^{n-1})$ as in Lemma 3.2, Section 3.2, where,

$$P(B | x_1^{n-1}) = P(B) \{1 + h_{n-1}(x_1^{n-1})[1 - P(B)]P^{-1}(B)\}^{\text{sgn } h_{n-1}(x_1^{n-1})} \times \{1 + h_{n-1}(x_1^{n-1})\}^{1 - \text{sgn } h_{n-1}(x_1^{n-1})}$$

and

$$h_{n-1}(x_1^{n-1}) = \begin{cases} Q_n - 1 & ; \quad Q_n \in [0, 1] \\ P(B)[1 - P(B)]^{-1}[Q_n - 1] & ; \quad Q_n \in [1, P^{-1}(B)] \end{cases}$$

$$Q_n \stackrel{\Delta}{=} \sum_{x_n=0,1} P(x_n | x_1^{n-1}) \{1 + h_n(x_1^n)[1 - P(B)]P^{-1}(B)\}^{\text{sgn } h_n(x_1^n)} \times [1 + h_n(x_1^n)]^{1 - \text{sgn } h_n(x_1^n)}$$

A $h_n(x_1^n)$ Constant Representing Extreme Partial Values

In this part, we first define the effect of the constants $\{h_1^{(i)}(x_i)\}_{1 \leq i \leq n}$ on the event B as follows:

If at least one of the constants $\{h_1^{(i)}(x_i)\}_{1 \leq i \leq n}$ equals the value 1, then event B occurs surely, if in addition $\sum_{i=1}^n h_1^{(i)}(x_i) > 0$

If at least one of the constants $\{h_1^{(i)}(x_i)\}_{1 \leq i \leq n}$ equals the value -1, then the nonoccurrence of event B is sure, if in addition $\sum_{i=1}^n h_1^{(i)}(x_i) < 0$

The events $\{A_i\}_{1 \leq i \leq n}$ do not affect the event B if $\sum_{i=1}^n h_1^{(i)}(x_i) = 0$

The above conditions translate to the following initial expressions for the conditional probability $P(B | x_1^n)$, where x_1^n is the value of the status vector of the affecting events $\{A_i\}_{1 \leq i \leq n}$:

$$P(B | x_1^n) = \begin{cases} 1 & ; \text{ if } \max_{1 \leq i \leq n} h_1^{(i)}(x_i) = 1 \text{ and } \sum_{i=1}^n h_1^{(i)}(x_i) > 0 \\ P(B) & ; \text{ if } \sum_{i=1}^n h_1^{(i)}(x_i) = 0 \\ 0 & ; \text{ if } \min_{1 \leq i \leq n} h_1^{(i)}(x_i) = -1 \text{ and } \sum_{i=1}^n h_1^{(i)}(x_i) < 0 \end{cases} \quad (3.28)$$

Via linear interpolation from the above expression we obtain the general expression of the conditional probability $P(B | x_1^n)$, as a function of the influence constants $\{h_1^{(i)}(x_i)\}_{1 \leq i \leq n}$, as follows:

$$P(B | x_1^n) = \begin{cases} P(B) + \max_{1 \leq i \leq n} (h_1^{(i)}(x_i)) [1 - P(B)] & ; \quad \sum_{i=1}^n h_1^{(i)}(x_i) > 0 \\ P(B) & ; \quad \sum_{i=1}^n h_1^{(i)}(x_i) = 0 \\ P(B) + \min_{1 \leq i \leq n} (h_1^{(i)}(x_i)) P(B) & ; \quad \sum_{i=1}^n h_1^{(i)}(x_i) < 0 \end{cases} \quad (3.29)$$

Defining the operators $O(x) = \begin{cases} 1 & ; \quad x > 0 \\ 0 & ; \quad x < 0 \end{cases}$ and $U(x) = \begin{cases} 1 & ; \quad x \geq 0 \\ 0 & ; \quad x < 0 \end{cases}$, we can rewrite equation (29) in a compressed form as follows.

$$P(B | x_1^n) = P(B) \left\{ 1 + P^{-1}(B) [1 - P(B)] \max_{1 \leq i \leq n} h_1^{(i)}(x_i) \right\}^{O(\sum_{i=1}^n h_1^{(i)}(x_i))} \left\{ 1 + \min_{1 \leq i \leq n} h_1^{(i)}(x_i) \right\}^{-U(\sum_{i=1}^n h_1^{(i)}(x_i))} \quad (3.30)$$

Next, we express a lemma regarding the consistency condition for our present model, evolving from the application of the Bayes' Rule and the Theorem of Total Probability on (3.30). The lemma is the parallel to Lemma 3.1 in section 3.2, for the model in the present case.

Lemma 3.5

For the influence model expressed in (3.30), the probabilities $P(B)$, $P(x_1^n)$ and the influence constants $\{h_1^{(i)}(x_i)\}_{1 \leq i \leq n}$ must satisfy the following condition:

$$[1 - P(B)] \sum_{x_1^n : \sum_{i=1}^n h_1^{(i)}(x_i) > 0} P(x_1^n) \max_{1 \leq i \leq n} h_1^{(i)}(x_i) + P(B) \sum_{x_1^n : \sum_{i=1}^n h_1^{(i)}(x_i) < 0} P(x_1^n) \min_{1 \leq i \leq n} h_1^{(i)}(x_i) = 0 \quad (3.31)$$

From the consistency condition in (3.31), we notice that when examining all the values of the status vector X_1^n , it is necessary that some x_1^n vector values exist such that $\max_{1 \leq i \leq n} h_1^{(i)}(x_i)$ is positive and that some x_1^n vector values exists such that $\min_{1 \leq i \leq n} h_1^{(i)}(x_i)$ is negative.

Temporal Issues

Here, we will assume that the very existence of the affecting events is revealed sequentially. Let then the existence and the status of the events $\{A_i\}_{1 \leq i \leq n}$ be revealed sequentially in time, from A_1 to A_n , where the status of events A_1 to A_k is known at time T_k . At time T_k , the partial status vector x_1^k is expressed and for each one of its values, the probability $P(x_1^k)$ and the quantities, $S_k(x_1^k) = \sum_{i=1}^k h_1^{(i)}(x_i)$, $F_k(x_1^k) = \max_{1 \leq i \leq k} h_1^{(i)}(x_i)$ and $G_k(x_1^k) = \min_{1 \leq i \leq k} h_1^{(i)}(x_i)$ are computed. Next, the probability $P(B)$ is computed from (31) as follows:

$$P(B) = P_k(B) = \left[\sum_{x_1^k : S_k(x_1^k) > 0} P(x_1^k) F_k(x_1^k) - \sum_{x_1^k : S_k(x_1^k) < 0} P(x_1^k) G_k(x_1^k) \right]^{-1} \sum_{x_1^k : S_k(x_1^k) > 0} P(x_1^k) F_k(x_1^k) \quad (3.32)$$

Given each x_1^k value, the probability $P(B)$ in (3.32) is then used to compute the conditional probability $P(B | x_1^k)$, as,

$$P(B | x_1^k) = P_k(B) \left\{ 1 + P_k^{-1}(B) [1 - P_k(B)] F_k(x_1^k) \right\}^{O(S_k(x_1^k))} \left\{ 1 + G_k(x_1^k) \right\}^{-U(S_k(x_1^k))} \quad (3.33)$$

At time T_{k+1} , upon the revelation of the existence and the status of the affecting event A_{k+1} , for each status vector x_1^{k+1} , the quantities, $S_{k+1}(x_1^{k+1}) = S_{k+1}(x_1^k) + x_{k+1}$, $F_{k+1}(x_1^{k+1}) = \max(F_k(x_1^k), h(x_{k+1}))$, $G_{k+1}(x_1^{k+1}) = \min(G_k(x_1^k), h_1(x_{k+1}))$ are first recursively computed. Then, the probability $P(B)$ is recomputed as

$$P(B) \stackrel{\Delta}{=} P_{k+1}(B) = \left[\sum_{x_1^{k+1}: S_{k+1}(x_1^{k+1}) > 0} P(x_1^{k+1}) F_{k+1}(x_1^{k+1}) - \sum_{x_1^{k+1}: S_{k+1}(x_1^{k+1}) < 0} P(x_1^{k+1}) G_{k+1}(x_1^{k+1}) \right]^{-1} \times \sum_{x_1^{k+1}: S_{k+1}(x_1^{k+1}) > 0} P(x_1^{k+1}) F_{k+1}(x_1^{k+1}) \quad (3.34)$$

The probability in (3.31) is used to compute the conditional probability below.

$$P(B | x_1^{k+1}) = P_{k+1}(B) \left\{ 1 + P_{k+1}^{-1}(B) [1 - P_{k+1}(B)] F_{k+1}(x_1^{k+1}) \right\}^{O(S_{k+1}(x_1^{k+1}))} \left\{ 1 + G_{k+1}(x_1^{k+1}) \right\}^{1 - U(S_{k+1}(x_1^{k+1}))} \quad (3.35)$$

We note that the time evolution of the conditional probabilities $P(B | x_1^k)$ is different for different time orderings of the affecting events $\{A_i\}_{1 \leq i \leq n}$.

A linear $h_n(x_1^n)$ Constant

Here, we assume that the effects of events $\{A_i\}_{1 \leq i \leq n}$ on event B are weighted by a known set

$\{w_i\}_{1 \leq i \leq n}$ of weights, such that $w_i \geq 0; \forall i$ and $\sum_{i=1}^n w_i = 1$. Then, given the constants $\{h_1^{(i)}(x_i)\}_{1 \leq i \leq n}$,

we define $h_n(x_1^n)$ as follows, for some given value $\alpha : 0 \leq \alpha < 1$:

$$h_n(x_1^n) = \begin{cases} (1 - \alpha)^{-1} \sum_{i=1}^n w_i h_1^{(i)}(x_i) & ; \left| \sum_{i=1}^n w_i h_1^{(i)}(x_i) \right| \leq 1 - \alpha \\ 1 & ; \sum_{i=1}^n w_i h_1^{(i)}(x_i) \geq 1 - \alpha \\ -1 & ; \sum_{i=1}^n w_i h_1^{(i)}(x_i) \leq -(1 - \alpha) \end{cases}$$

A nonzero α value translates to the probability of event B being equal to one, not only when all the $\{h_1^{(i)}(x_i)\}_{1 \leq i \leq n}$ values equal one, but also when a predefined weighted majority exceeds a total weighted sum of $1 - \alpha$. Similarly then, the event B occurs with zero probability when the weighted sum of the $\{h_1^{(i)}(x_i)\}_{1 \leq i \leq n}$ values is less than $-(1 - \alpha)$, rather than only when it equals -1.

The relationships between the $h_n(x_1^n)$ and $h_{n-1}(x_1^{n-1})$ influence constants and the probabilities $P(B)$, $P(x_1^n)$ and $P(B | x_1^n)$ are as in IX.A.

A $h_n(x_1^n)$ constant representing Noisy OR Format

Given the constants $\{h_1^{(i)}(x_i)\}_{1 \leq i \leq n}$, we define here $h_n(x_1^n)$ as follows; where α is such that $0 \leq \alpha \leq 1$:

$$h_n(x_1^n) = \left\{ 1 - (1 - \alpha)^{-1} \prod_{i=1}^n (1 - |h_1^{(i)}(x_i)|) \right\}^{\text{sgn}(h_n(x_1^n))} \left\{ -1 + \alpha^{-1} - \alpha^{-1} \prod_{i=1}^n (1 - |h_1^{(i)}(x_i)|) \right\}^{1 - \text{sgn}(h_n(x_1^n))} \quad (3.36)$$

Then, via (3.3) and (3.5) in section 3.2, we obtain:

$$P(B) = \alpha \quad (3.37)$$

$$1 - P(B | x_1^n) = \prod_{i=1}^n (1 - |h_1^{(i)}(x_i)|) \quad (3.38)$$

The expression in (3.38) represents the Noisy-OR format [1, 4], where the probabilities in the latter are here substituted by the absolute values of the one-dimensional influence constants $\{h_1^{(i)}(x_i)\}_{1 \leq i \leq n}$.

Influence Constant Comparison

Figure 3.2 shows an example IN with a binary event B known to be affected by the events $\{A_i\}_{1 \leq i \leq 3}$. The edges connecting the external affecting events $\{A_i\}_{1 \leq i \leq 3}$ to the event B are shown annotated with the constants $[h_1^{(i)}(x_i = 1), h_1^{(i)}(x_i = 0)]$ for each i, where $x_i = 0, 1$ represents one of the two states of an affecting event A_i . A global influence constant $h_3(x_1^3)$ is then designed using all four (i.e., A-D) special influence functions presented in this section. Table 3.1 shows the computed values of $h_3(x_1^3)$ and corresponding $P(B | x_1^3); \forall x_1^3$ for each of the four cases. For illustration purposes, we also assume that the joint probability $P(x_1^3); \forall x_1^3$ values are computed by assigning $P(x_3 = 1) = 0.8, P(x_3 = 0) = 0.2$ and $P(x_i = 1) = P(x_i = 0) = 0.5$; for $i = 1, 2$ and by assuming $\{A_i\}_{1 \leq i \leq 3}$ to be mutually independent.




Fig. 3.2 Example TIN

TABLE 3.1 Comparison of Influence Constants

x_1	x_2	x_3	$h_3(x_1^3)$				$P(B x_1^3)$			
			A. CAST Log- ic Based	B. Extreme Partials	C. Linear Constant	D. Noisy- OR	A. CAST Logic Based	B. Extreme Partials	C. Linear Constant	D. Noisy- OR
0	0	0	-0.33	-0.33	-0.073	-0.095	0.335	0.335	0.452	0.463
0	0	1	-0.699	-0.33	-0.101	-0.095	0.150	0.335	0.452	0.449
0	1	0	0.66	0.66	0.498	0.326	0.83	0.83	0.663	0.749
0	1	1	0.242	0.0	0.0	0.326	0.621	0.5	0.663	0.5
1	0	0	0.33	0.33	0.176	-0.095	0.665	0.665	0.452	0.588
1	0	1	-0.33	-0.33	-0.044	-0.095	0.335	0.335	0.452	0.478
1	1	0	0.847	0.66	1.00	0.326	0.923	0.83	0.663	1.00
1	1	1	0.66	0.66	0.196	0.326	0.83	0.83	0.663	0.598

From the values included in Table 3.1, we notice the sensitivity of the computed probability of event B on the selected structure of the aggregate influence constant. Different such structures reflect different environments and their choice is at the discretion of an expert.

3.9 Evaluation Metrics

As already repeatedly stated, the INs and TINs studied in this paper are basically BNs whose conditional probabilities are approximated by expert provided influence constants. Thus, the architectural and computational complexities involved are similar to those in BNs [8, 31, 32, 33, 34], while the complexities involved in the computation of influence constants depend on the specific structure of the latter (see Section VIII). The evolution of lower dimensionality conditional probabilities from high dimensionalities ones, as in Lemma 3.2, section 3.2, is of dynamic programming nature inducing polynomial complexity. As stated in section 3.7, the accuracy of a IN or TIN model is determined by the accuracy of the selected influence constants. The accuracy of the latter may be tested and they may be subsequently adjusted appropriately.

3.10 Experimental Setup

In this section, we lay out the steps involved in an experimental setup. Given an event B, determine **all** the events $\{A_i\}_{1 \leq i \leq n}$ **known** to be affecting its occurrence. Given B, all the known affecting events $\{A_i\}_{1 \leq i \leq n}$, and the causal strengths $[h_1^{(i)}(x_i = 1), h_1^{(i)}(x_i = 0)]$ between each A_i and B, design an influence constant $h_n(x_1^n)$, where x_1^n signifies the value of the status vector of the events $\{A_i\}_{1 \leq i \leq n}$, and where $-1 \leq h_n(x_1^n) \leq 1; \forall x_1^n$ values. The $h_n(x_1^n)$ constant may have one of the forms presented in section 3.8. If **all** in (b) is given, then upon a given probability of the status vector X_1^n , say $P(x_1^n); \forall x_1^n$ values, the probability of event B is given by the following equation, named the consistency equation.

$$\sum_{x_1^n} P(x_1^n) \{1 + h_n(x_1^n)[1 - P(B)]P^{-1}(B)\}^{\text{sgn } h_n(x_1^n)} \{1 + h_n(x_1^n)\}^{1 - \text{sgn } h_n(x_1^n)} = 1$$

whose equivalent form is:

$$P(B) = \left[\sum_{x_1^n : \text{sgn } h_n(x_1^n)=1} P(x_1^n) h_n(x_1^n) \right] \left[\sum_{x_1^n} P(x_1^n) h_n(x_1^n) \right]^{-1}, \text{ if the denominator is non zero}$$

When **all** the affecting events $\{A_i\}_{1 \leq i \leq n}$ are known, but the status of some of them are unknown, then, the probability $P(B)$, as computed in step (c) is used to compute the conditional probability $P(B | x_1^k)$, when the status vector of only k affecting events is known as:

$$P(B | x_1^k) = P(B) \{1 + h_k(x_1^k)[1 - P(B)]P^{-1}(B)\}^{\text{sgn } h_k(x_1^k)} \bullet \{1 + h_k(x_1^k)\}^{1 - \text{sgn } h_k(x_1^k)}$$

where $h_k(x_1^k)$ is computed in a dynamic programming fashion from the influence constant $h_n(x_1^n)$ in (b); as follows:

$$h_{n-1}(x_1^{n-1}) = \begin{cases} Q_n - 1 & ; \quad Q_n \in [0,1] \\ P(B)[1 - P(B)]^{-1}[Q_n - 1] & ; \quad Q_n \in [1, P^{-1}(B)] \end{cases}$$

$$\text{for } Q_n = \sum_{x_n=0,1}^{\Delta} P(x_n | x_1^{n-1}) [1 + h_n(x_1^n)[1 - P(B)]P^{-1}(B)]^{\text{sgn } h_n(x_1^n)} \bullet [1 + h_n(x_1^n)]^{1 - \text{sgn } h_n(x_1^n)}$$

We note that in the above expression, the affecting events $\{A_i\}_{1 \leq i \leq n}$ are assumed ordered as of the revealing of their status in time. Different such ordering results in different evolutions of the conditional probabilities $P(B | x_1^k)$.

When the existence as well as the status of the affecting events are sequentially revealed, then at time k, $P_k(B)$ and $P_k(B | x_1^k)$ are computed as in (c) and (d) where n is substituted by k in the latter.

Example 1: The following example illustrates the steps (a) to (e) with the help of an example TIN. Figure 3.3 shows a IN with a binary event B known to be affected by the events $\{A_i\}_{1 \leq i \leq 4}$.

The edges connecting the external affecting events $\{A_i\}_{1 \leq i \leq 4}$ to the event B are shown in Fig. 3, annotated with the constants $[h_1^{(i)}(x_i = 1), h_1^{(i)}(x_i = 0)]$ for each i, where $x_i = 0, 1$ represents one of the two states of an affecting event A_i . A global influence constant $h_4(x_1^4)$ is then designed using the CAST logic expression (25) in section 3.8. Table 3.2 shows the computed values for $h_4(x_1^4); \forall x_1^n$. The joint probability $P(x_1^4); \forall x_1^4$ values are computed by assigning $P(x_i = 1) = P(x_i = 0) = 0.5; \forall i$ and by assuming $\{A_i\}_{1 \leq i \leq 4}$ to be mutually independent (Lemma 3). The probability of occurrence of event B, i.e., $z = 1$, is now calculated with the consistency equation, and is given as $P(z = 1) = 0.5$. Assuming the status of all the affecting events to be known, the conditional probabilities $P(B | x_1^4); \forall x_1^4$ are calculated via expression (26), and are shown in Table 3.2.




Fig. 3.3 Example TIN

The assumption in step (iii), regarding the knowledge of the status of all the affecting events, may not be valid at times. Such is the case of a TIN with delays on edges (see Definition 3.1), reflecting variations in the times when the status of the affecting events become known. To illustrate this notion, we add temporal information to the IN in Fig. 3.3. The added temporal information together with the underlying graph is shown in Fig. 3.4. The time assigned to an affecting event A_i is the instance at when it assumes a state, i.e., $x_i = 0$ or 1 . Prior to that time, the state of the event is assumed unknown. As stated in Definition 3.1, this combination of the external affecting events' status and their timing is also termed a Course of Action (COA), in the TIN literature.

TABLE 3.2 Conditional Probabilities

x_1	x_2	x_3	x_4	$h_4(x_1^+)$	$P(z=1 x_1^+)$
0	0	0	0	-0.990000	0.005000
0	0	0	1	-0.999900	0.000050
0	0	1	0	-0.900000	0.050000
0	0	1	1	-0.999000	0.000500
0	1	0	0	-0.900000	0.050000
0	1	0	1	-0.999000	0.000500
0	1	1	0	-0.000001	0.499999
0	1	1	1	-0.990000	0.005000
1	0	0	0	0.990000	0.995000
1	0	0	1	0.000001	0.500001
1	0	1	0	0.999000	0.999500
1	0	1	1	0.900000	0.950000
1	1	0	0	0.999000	0.999500
1	1	0	1	0.900000	0.950000
1	1	1	0	0.999900	0.999950
1	1	1	1	0.990000	0.995000

The temporal information in the TIN, Fig. 3.4, determines the dynamics of the relationship between the affecting events and the affected event B; specifically, the times when the status of the affecting events are revealed to B. Figure 3.5 shows a IN equivalent, obtained by mapping the status of the affecting events and their effects on the event B, on a timeline. This mapping determines the number of affecting events ‘k’ at different time points (or time slices). For the temporal case presented in section VI, the existence of all the affecting events is known to the event B a priori; their status, however, remain unknown until revealed, as determined by the COA and the delays on the edges. The probability $P(B)$, as calculated in step (c), is used to compute the conditional probabilities $P(B | x_1^k); k = 1, 2, 4$, i.e., $P(B | x_1^1)$, $P(B | x_1^2)$, and $P(B | x_1^4)$, as illustrated in the figure. Table 3.3 shows the values for $P(B | x_1^1)$ and $P(B | x_1^2)$, as computed by the corresponding $h_1(x_1^1)$ and $h_2(x_1^2)$. The posterior probability of B captures the impact of an affecting event on B and can be plotted as a function of time for a corresponding COA. This plot is called a Probability Profile [12, 27]. Fig. 3.6 shows the resulting probability profile for the illustrative example. The plotted values in the profile are shown with bold letters in Tables 3.3-3.4. The overall complexity is polynomial.

For the temporal case presented in section IX, the existence as well as the status of the affecting events are not known a priori but are determined by the given COA and the delays on the edges. At time k, $P_k(B)$ and $P_k(B | x_1^k)$ are computed as in (c) and (d) where n is substituted by k in the latter. Table 3.4 shows the computed values of $P_k(B)$ and $P_k(B | x_1^k); k = 1, 2, 4$ and Fig. 3.6(b) shows the resulting probability profile.




Fig. 3.4. Example TIN with COA and Edge Delays





Fig. 3.5. Temporal Model for the Example TIN

TABLE 3.3 Posterior Probabilities of B


x_1	$P(z=1 x_1^1)$	x_1	x_2	$P(z=1 x_1^2)$
0	0.076381	0	0	0.013887
1	0.923619	0	1	0.138875
		1	0	0.861125
		1	1	0.986113

TABLE 3.4 Probability Profile values

x_1	$P_1(B)$	$P(z=1 x_1^1)$	x_1	x_2	$P_2(B)$	$P(z=1 x_1^2)$	x_1	x_1	x_4	x_3	$P_4(B)$	$P(z=1 x_1^4)$		
0	0.5	0.005	0	0	0.5	0.005	0	0	0	0	0.5	0.005		
1		0.995	0	1		0.005	0	0	0	1		0.005		
			1	0		0.995	0	0	1	0		0.005		
			1	1		0.995	0	0	1	1		0.005		
							0	1	0	0		0.005		
							0	1	0	1		0.005		
							0	1	1	0		0.95		
							0	1	1	1		0.005		
							1	0	0	0		0.005		
							1	0	0	1		0.995		
							1	0	1	0		0.05		
							1	0	1	1		0.995		
							1	1	0	0		0.995		
							1	1	0	1		0.995		
							1	1	1	0		0.995		
							1	1	1	1		0.995		



(a). For Temporal Case I



(b). For Temporal Case II

Fig. 3.6. Probability Profile for the Example COA

Example 2: In multi-node connected network structures, given a set of external unaffected affecting events $\{A_i\}_i$, given influence constants $\{h_n(x_1^k)\}_k$, pertinent conditional probabilities are constructed hierarchically, as the structure of the network dictates. Consider, for example, the network in Fig. 3.7, below. In this network, the affecting events $A_i; i = 1, 2, 3, 4$ are external and unaffected by other events, while events B and C are affected, B being affecting as well. Let us denote the status of event $A_i; i = 1, 2, 3, 4$ by x_i , the status of event B by y and the status of event C by z , where y, z and $\{x_i\}_{1 \leq i \leq 4}$ are 0-1 binary numbers. Let the influence constants $h(x_1, x_2), h(x_3, x_4)$ and $h(y, x_3, x_4)$ be given. Let also the joint probability $P(x_1, x_2, x_3, x_4)$ be given.




Fig. 3.7. A Multi-node Network

We then compute all the pertinent probabilities in the above network following the steps stated below:

1. Compute the probability $P(y)$ from the consistency condition:

$$\sum_{x_1, x_2} P(x_1, x_2) \{1 + h(x_1, x_2)[1 - P(y)]P^{-1}(y)\}^{\text{sgn } h(x_1, x_2)} \{1 + h(x_1, x_2)\}^{1 - \text{sgn } h(x_1, x_2)} = 1$$

where,

$$P(x_1, x_2) = \sum_{x_3, x_4} P(x_1, x_2, x_3, x_4)$$

2. Compute $P(y | x_1, x_2)$ from,

$$P(y | x_1, x_2) = P(y) \{1 + h(x_1, x_2)[1 - P(y)]P^{-1}(y)\}^{\text{sgn } h(x_1, x_2)} \{1 + h(x_1, x_2)\}^{1 - \text{sgn } h(x_1, x_2)}$$

3. Compute $P(y, x_3, x_4)$ as:

$$P(y, x_3, x_4) = \sum_{x_1, x_2} P(y | x_1, x_2) P(x_1, x_2, x_3, x_4)$$

4. Compute $P(z)$ from the consistency condition

$$\sum_{y, x_3, x_4} P(y, x_3, x_4) \{1 + h(y, x_3, x_4)[1 - P(z)]P^{-1}(z)\}^{\text{sgn } h(y, x_3, x_4)} \{1 + h(y, x_3, x_4)\}^{1 - \text{sgn } h(y, x_3, x_4)} = 1$$

5. Compute $P(z | y, x_3, x_4)$ from,

$$P(z | y, x_3, x_4) = P(z) \{1 + h(y, x_3, x_4)[1 - P(z)]P^{-1}(z)\}^{\text{sgn } h(y, x_3, x_4)} \{1 + h(y, x_3, x_4)\}^{1 - \text{sgn } h(y, x_3, x_4)}$$

6. Compute $P(z | x_1, x_2, x_3, x_4)$ from,

$$P(z | x_1, x_2, x_3, x_4) = \sum_y P(z, y | x_1, x_2, x_3, x_4) = \sum_y P(z | y, x_3, x_4) P(y | x_1, x_2)$$

3.11 Conclusion

In this chapter, we presented a comprehensive approach to Influence Nets including conditions for model consistency and dynamic programming evolution of the influence constants, as well as temporal issues and model testing methodologies. We revisited the earlier CAST logic [4, 5] based approach to Timed Influence Network (TIN) modeling [13, 15, 17], by redefining the design parameters for a TIN model, reevaluating the cases of independence and (partial) dependence among external affecting events, introducing new methods for aggregating joint influences from design parameters, and by offering new insights into the temporal aspects of causal influences modeled inside a TIN. The presented approach successfully overcomes the deficiencies in the CAST logic based TIN modeling and the inconsistencies therein. It also does not require any additional design information than that already available in a TIN constructed via CAST logic parameters; the entire repository of situational models developed earlier [15, 17, 18] may be simply reanalyzed (without any modifications) using the new set of computational tools introduced in this paper. We analyzed and evaluated our approach and tested it for a specific TIN. This illustrative application is presented in Chapter 5. The approach produces consistent and stable in time results.

Chapter 4

Meta-Model Driven Construction of Timed Influence Nets

Faisal Mansoor, Abbas K. Zaidi, Alexander H. Levis

4.1 Introduction

The analysis and decision problems often require modeling of subjective, informal, and uncertain concepts in a domain in order for an analyst to capture the required behavior of the domain. Influence Net (IN) [36], a variant of Bayesian Networks (BN), is an approach for modeling cause-and-effect relationships among variables of a domain. The construction of an IN requires a subject matter expert (SME) to model the parameters of the domain – random variables – as nodes in a network. The arcs (or directed edges) in the network represent the cause-and-effect relationships between the random variables. The nodes in an IN and their interdependencies may represent the inter effects between political, military, economic, social, infrastructure, and information (PMESII) factors present in an area of interest. The strengths of these dependencies are captured in terms of a small (i.e., linear) number of influence constants (as opposed to an exponential number of conditional probabilities in a BN). The IN approach was developed in recognition of the fact that most domain experts and situation analysts do not think in terms of conditional probabilities (as required for a BN) while relating affecting and effect variables in a domain. The INs provide an intuitive elicitation, update, and merge interface to an underlying BN that only requires specification of qualitatively described individual influences between each pair of an affecting and an affected variables. The approach then combines these individual influences to calculate the aggregate effect of multiple affecting variables on an effect variable in terms of conditional probability values of a resulting BN.

Wagenhals and Levis [13] have added a special set of temporal constructs to the basic formalism of Influence Nets. The Influence Nets with these additional temporal constructs are called Timed Influence Nets (TINs). A fully specified TIN model is characterized by the causal relationships between propositions and the values of the parameters, i.e., strength of influences [22], and temporal delays associated with these relationships. TINs have been experimentally used in the area of Effects Based Operations (EBOs) and Adversarial Modeling for evaluating alternative courses of actions and their effectiveness to mission objectives in a variety of domains, e.g., war games [15], and coalition peace operations. A number of analysis tools have been developed over the years for TIN models to help an analyst in solving problems of interest [22 - 24]. In this sequel, the term Influence Net (IN) will be used generically to refer to both INs and TINs.

The lack of familiarity with, or enthusiasm for, these analytical representations (i.e., BNs and/or TINs) prevents most domain experts and analysts from developing such models on their own and using them for the analysis tasks assigned. The tools implementing some of these formalisms [37 - 38] require prior knowledge and experience in modeling and, therefore, do not provide any assistance to such users. There is, however, a growing community of analysts who makes use of these analytical and quantitative formalisms resulting in a small, but expanding, repository of models addressing different PMESII aspects of a domain. There is, therefore, a

need not only to facilitate the model building task, but also to utilize the existing models for building quick, larger and better domain models without requiring experienced domain experts, at least, in the early stages of a domain modeling exercise.

This chapter introduces a meta-modeling approach that facilitates generalizing an entire class of problem-specific TINs in the form of a meta-model, called Template TIN. For example, a causal relation in an existing TIN might model, “If the Kurd population in the Northern provinces of Iraq finds its rights respected in the new administration, then it will be more cooperative in the new development plans.” A simple generalization of this could be, “If an *<ethnic minority>* in a *<geographic administrative unit>* finds its rights respected, then it will *<participate>* in the *<development activity>*.” A Template TIN captures such generalized relation using abstract entities characterizing a problem domain. It can be constructed by generalizing several TINs, or can be directly constructed by an expert using the template specification language. A set of stored templates can then be instantiated for a particular situation by substituting abstract entities with concrete instances characterizing a situation.

A Template TIN provides a means for leveraging past, tested TIN models that may have been constructed by other experts or team of experts in addressing a problem similar to the one under consideration. It simplifies the Influence Net construction process by providing an analyst with a repository of templates capturing different fragments of a generalized understanding of the problem domain in terms of possible causal relationships among domain variables. These templates are not intended to prescribe a solution or a model but are a means to enhance an analyst’s search for better understanding of the domain and to facilitate the process of building a more pertinent model of the domain. A Template IN identifies a set of concepts that are considered *relevant* by some analyst or a team of analysts for a problem-specific domain. These concepts are described at an abstract level and are required to be instantiated when a model is being constructed for a new domain. Exploring available knowledge bases for information required for instantiating a Template TIN is also a complex and time-consuming task, especially if that information is not implicitly available in the form of an expert of the new domain. With increasing popularity and use of structured knowledge representation and reasoning tools, it is now possible to automate the data exploration and Template TIN instantiation process. In the presented framework, we use an OWL [39] ontology not only as the knowledge representation for domain data, but also as a mechanism to reason about this data while constructing a situational assessment model as a Timed Influence Net. For a fully automated instantiation of a Template TIN with the data in an ontology, the approach proposes a mapping scheme that provides a definition of abstract concepts present in the Template TIN in terms of concepts and properties available in the ontology. These definitions are constructed as a set of mapping rules. The mapping rules are SPARQL (Protocol and RDF Query Language) [40] queries that use the OWL reasoning engine Pellet [41] to identify relevant data in the ontology to be used for TIN instantiation.

The rest of the chapter is organized as follows. In section 4.2, we present the architecture of the developed meta-model driven ontology based TIN construction approach. Section 4.3 presents how this approach can aid in developing situation assessment model for some class of problems. Section 4.4 contains an example while 4.5 concludes the paper with a discussion on future research directions.

4.2 The Methodology

Figure 4.1 shows an overview of the meta-model driven ontology based TIN construction approach. The following subsections describe each component of the architecture in Fig. 4.1.

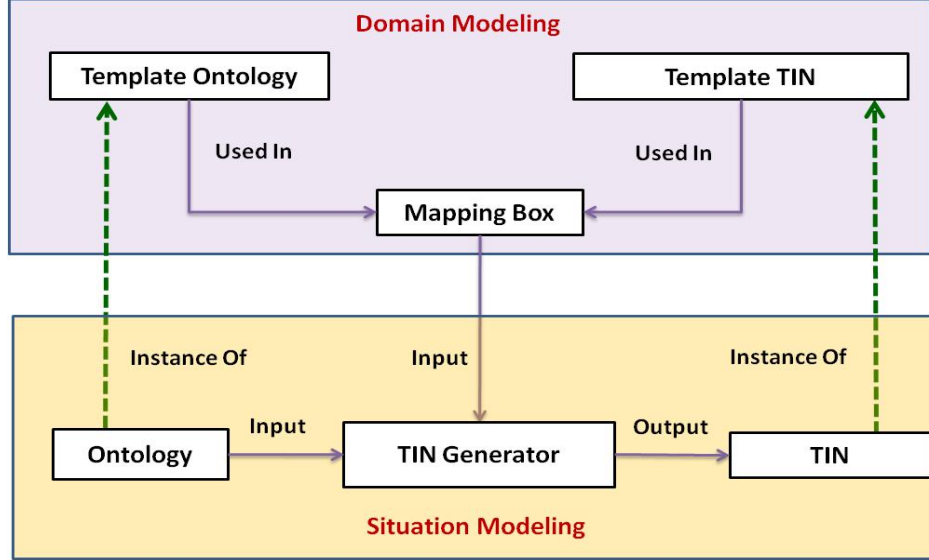


Fig. 4.1 Architecture of the Approach

Ontology

Ontology is an explicit conceptualization of some domain of discourse. We can define an ontology as a knowledge base composed of Terminology Box (TBox) and Assertion Box (ABox); $K = (TBox, ABox)$, where:

- TBox is a finite set of concepts and a finite set of relations between the concepts.
- ABox is a finite set of instances, relations between instances and relations between instances and concepts in TBox.

In Fig. 4.1, the terms Template Ontology and Ontology refer to the TBox and the ABox of an ontology, respectively.

Template Timed Influence Net (TIN)

An Influence Net for a problem instance involves a pre-specified set of random variables with fixed cause-and-effect relationships. The goal of Template Timed Influence Net is to capture the abstract relationships between classes of causes and effects characterized by a problem domain. Template TINs extend TINs just as first-order logic extends propositional logic.

Template TINs are Influence Nets except that the nodes in them contain labels formed by variables instead of terms representing domain instances. Formally, a Template Influence Net is described as follows:

Definition 4.1 - A Template TIN \mathcal{T}_{IN} is a tuple (V, E, C, B) where $G(V, E)$ is a Directed Acyclic Graph (DAG), and


- V is a set of random variables that makes up the nodes of an Timed Influence Net. Where
 - All the variables in the Timed Influence Net have binary states.
 - For each random variable $v \in V$ there exists a set of slot variables $svar(v)$ such that each slot variable in $svar(v)$ represents an abstract domain entity.
- E is a set of directed links that connect pairs of nodes in V .
- C is a set of causal strength parameter and the time delay triples. Each triple in C is associated with an edge in E (the causal strength parameters are usually denoted as h and g values).
- B represents a set of baseline and prior probabilities associated with non-root and root nodes respectively.

Definition 4.2 (Template Influence):- A Template Influence \mathcal{T}_I in a Template TIN is a tuple $\{sr, tr, e, \mathcal{M}(\mathcal{T}_I)\}$, where $sr, tr \in V, e \in E$, and sr is the source node of e and tr is the target node of e . In addition, $\mathcal{M}(\mathcal{T}_I)$ is a function providing a mapping from $svar(sr) \cup svar(tr)$ to situation specific entities.

A Template TIN is a collection of several Template Influences. An example Template Influence is shown in Fig. 2a. The terms $[?leader]$ and $[?funding]$ represent the slot variables $svar(sr)$ and $svar(tr)$, respectively.

Mapping Box

Mapping Box defines influences present in a Template TIN in terms of concepts and relation available in a Template Ontology. Specifically, Mapping Box is a set of mappings where each mapping is defined as a pair consisting of a Template Influence and an ontology query (Fig. 2b), the query establishes the link between Template Influence and ontology.



(a) Template Influence

```
Select ?leader, ?funding
Where
  (?leader instanceOf KnownTerrorist)
  (?leader leadershipRank high)
  (?observationL instanceOf C2Observation)
  (?observationL hasParticipant ?leader)
  (?funding instanceOf FinancialSource)
  (?observationF instanceOf FinancialAcquisitionObservation)
  (?observationF hasFinancialSource ?funding)
```

(b) Ontology Query

Fig. 4.2 An Example Mapping

Timed Influence Net Generator

Given an ontology (i.e., TBox and ABox both) describing a particular situation, TIN Generator uses the abstract definitions available in the Mapping Box to produce a TIN specialized for

the situation described by the input ontology. For the example in Fig. 4.2 this would amount to running the query in Fig. 4.2b for identifying the instances in the ontology that match with the conditions defining the slot variables [?leader] and [?funding] in the Template Influence (Fig. 4.2a). The results of the query are then substituted for the slot variables. The following is a formal definition of this substitution process.

Definition 4.3 (Substitution):- Let t_i be a term denoting a situation specific entity belonging to a domain of interest. Then a substitution $\theta = \{v_1/t_1, \dots, v_n/t_n\}$ is an assignment of term t_i to variable v_i . Applying a substitution θ to Template Influence \mathcal{T}_I yields the instantiated Influence $\mathcal{T}_I\theta$ where all occurrences of the variable v_i are simultaneously replaced by the term t_i .

As shown in Fig. 4.1, the presented TIN construction approach is a two-phase process consisting of a Domain-Modeling phase and a Situation-Modeling phase. In the Domain-Modeling phase, ontology and TIN templates are used to develop a generalized mapping that can be applied to any ontology compatible with the Template Ontology. Domain-Modeling is a process done only once. When a Mapping Box is created, instantiating a TIN from a given instance ontology describing a particular situation becomes a completely automated process.

4.3 Castalia

The described meta-model driven ontology based TIN construction process has been implemented as part of the Pythia [38] suite of applications. The implemented software package, called Castalia, takes as input (a) an OWL [39] ontology expressed in Protégé [41], (b) mapping rules expressed in SPARQL [40], and (c) Template TIN developed using Pythia application [38] for instantiating TINs. Pellet [41] is used as the ontology reasoning and query engine by Castalia. The output of Castalia can be imported in Pythia as a Timed Influence Net for subsequent analysis.

An implementation view of the architecture in Fig. 4.1 is shown in Fig. 4.3.

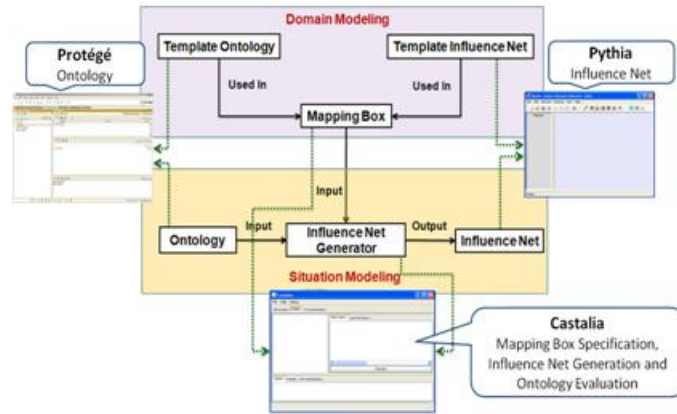


Fig. 4.3 Architecture with Respective Applications

From a procedures point of view, Figs. 4.3 and 4.4 illustrate how a team of analysts and knowledge engineers can use Protégé, Pythia, and Castalia to automate the Influence Net construction process.

The process in Fig. 4.4 comprises of two phases: Domain Modeling and Situation Modeling. In the Domain Modeling phase, Protégé is used to develop the Template Ontology and Pythia is used to develop the Template TIN. The Template Ontology can be developed by a knowledge engineer with technical knowledge of OWL and Protégé, and some understanding of the problem domain. A Template TIN can be developed by a domain expert with no or little help from a knowledge engineer. A Template TIN can also be derived by generalizing already developed TIN models. The derivation of Template TIN in the latter case can be done by a knowledge engineer with no or little help from an analyst. The two meta models are then used to construct the MBox using Castalia, which contains a graphical user interface module for developing the MBox. The construction of MBox requires both the knowledge engineer for SPARQL syntax and the analyst to describe the mapping. Once the MBox is available, it can be used to develop the situation specific Influence Nets during the Situation Modeling phase. It should be noted that the Domain Modeling phase is a one-time effort. Given a new situation described using an ontology which can be automatically constructed using the available text-extraction ontology building tools, the Influence Net Generator module in Castalia automatically generates an Influence Net specialized for the situation. The generated Influence Net is compatible with Pythia and can be opened in Pythia to perform different kind of analyses. The Situation Modeling phase is where the effort put in the Domain Modeling phase pays off: given a repository of such domain models in the form of templates, situation models for new problems can be easily instantiated by an analyst by merely selecting a Template TIN and clicking a button in Castalia to instantiate it with information in an OWL ontology. An analyst using Castalia does not need to know anything about SPARQL, OWL, and/or Protégé. Moreover, Castalia also contains a module for ontology assessment, which computes the TBox and ABox fitness measures during the Influence Net generation process.




Fig. 4.4 Construction Process

4.4 Application

To illustrate the approach described in this paper, a detailed investigative report [42] on the 1998 bombings of the US embassies in Kenya and Tanzania was used to develop and populate an OWL ontology. The class hierarchy of the ontology is shown in Fig. 4.5. The Template Ontology was developed using the concepts derived from an understanding of the general nature of such incidents.

A Timed Influence Net model was constructed for the Kenya incident, using the information in the report, to capture the events leading up to the bombing. The Kenya TIN model was then transformed into a Template TIN to represent a generalized model for a terrorist attack on a US interest abroad. The Template TIN derived from the Kenya based TIN is shown in Fig. 4.6. The Template TIN is a collection of several Template Influences. The nodes in this template represent abstract concepts derived by replacing instances from the Kenya TIN with the slot variables.

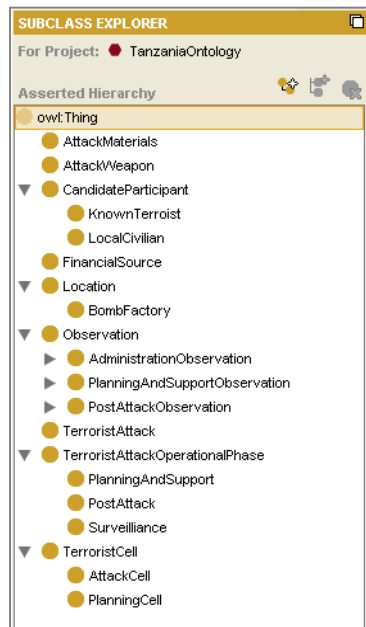


Fig. 4.5 Class Hierarchy of the Kenya and Tanzania Bombing Ontology

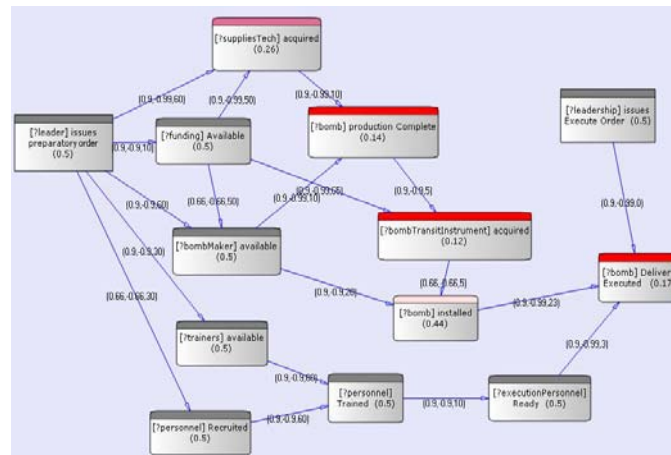


Fig. 4.6 Template TIN used in the Application

A Mapping Box (i.e., MBox) was then constructed with the help of concepts in the Template Ontology (i.e., OWL ontology’s TBox) and in the Template TIN. The Following is an example of a mapping rule, expressed in SPARQL, which provides a definition of the relation between a leader and available funding as captured in a Template Influence. The rule states that a leader is a

known terrorist and has a high leadership rank. It also states that funding is a financial source and a leader has funding.

```
SELECT ?leader ?funding
{
  ?leader rdf:type this:KnownTerrorist.
  ?leader this:leadershipRank "high"^^xsd:string.
  ?obs1 rdf:type this:C2Observation.
  ?obs1 this:hasParticipant ?leader.
  ?funding rdf:type this:FinancialSource.
  ?obs2 rdf:type this:FinancialAcquisitionObservation.
  ?obs2 this:hasFinancialSource ?funding.
}
```

The Template Ontology was populated with information in the OWL ontology's ABox. The ABox used for this illustration contained data from the Tanzania bombing incident only. The Tanzania instance ontology was provided to Castalia that used it and the MBox to generate a TIN specialized for the Tanzania incident. The construction of a new instance TIN was automatically done by Castalia which replaced the variables in each of the Template Influences by the values available in the instance ontology with the help of mapping rules in the MBox. Figure 4.7 shows the generated Timed Influence Net. As can be seen in the generated TIN, not all variables were instantiated with values from the Tanzania ontology. For example, Castalia reasoning engine did not find an instance for the slot variable [?trainer] in the Tanzania data. In other words, the data available for the Tanzania incident had no person identified as the potential trainer of the bomber. The generated TIN, however, succeeded in capturing a number of key elements, i.e., people involved and equipment, used in the bombing.

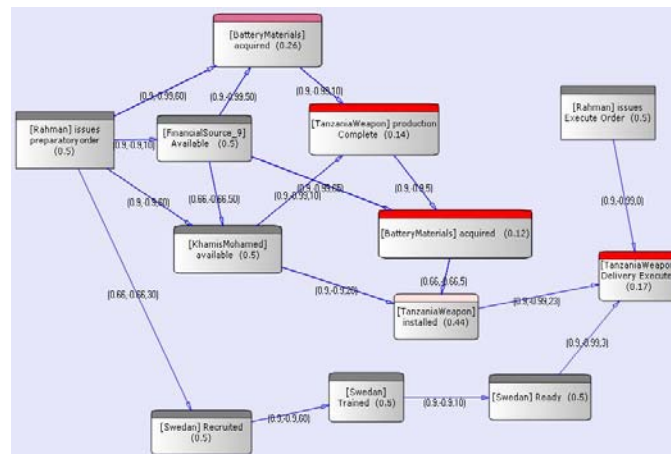


Fig. 4.7 Instantiated Timed Influence Net for Tanzania Bombing

This generated TIN could be used by a SME or analyst as a fragment in a larger domain model that might require it as part of a more complex situation involving vehicle borne explosives or IED attacks. The instantiated TIN can also be used for course of action analysis as well as a host of other analyses provided by the TIN suite of tool. As mentioned earlier, the templates are not prescriptions to be used for building future models, but useful references that a SME/analyst might like to consult either during the process of building a new model or as a start-

ing point for it. The generated TIN can also be used to study temporal aspects of the influences in the TIN and to do a course of action analysis indicating how long it takes for input events to cause some desired or undesired effects on output nodes.

4.5 Conclusion

Template TIN models a problem at a generic level using abstract entities characterizing the problem domain, allowing an analyst to model an entire class of Timed Influence Nets using a compact representation. However they also lack specialized domain modeling constructs like objects, properties, inheritance etc. that makes it difficult to construct and maintain probabilistic models for complex domains. This limitation of Influence Nets was overcome by using ontologies, which provide a highly expressive language for representing complex domains. The presented approach uses ontologies along with Template TINs to automate the Influence Net construction process. We believe that, given the time and expertise required for Influence Net construction by hand, an automated approach for Influence Net construction would prove vital for Influence Net's widespread adaption and use.

The mapping box used in the approach acts as a bridge between Template TIN and ontologies, and are expressed using SPARQL. In the presented approach, an analyst will have to manually specify the MBox. One possible way to further facilitate an analyst would be to automate the MBox specification using automated inductive learning techniques.

It is assumed that Castalia will benefit a growing community of TIN users that include both government (e.g., NASIC, NPS, JIEDDO, AU) and private (e.g., Raytheon, ANSER, and other corporations supporting DOD and DHS) organizations for rapid construction and deployment of situational influence models for intelligence assessment, course of action planning and assessment in EBO, and adversarial modeling problems. The repository of Template Timed Influence Nets can also be used for training future analysts in different problem domains. The update and the re-use of the templates will also facilitate automated generation of situational models for assessment and planning purposes in a new theatre of operations.

Chapter 5

Adversary Modeling Applications

5.1 Modeling Uncertainty in Adversary Behavior: Attacks in Diyala Province, Iraq, 2003–2006

Claudio Cioffi-Revilla and Pedro Romero

5.1.1 Introduction

Uncertainty is a universal characteristic of conflict behavior and low-intensity warfare, guerrilla, insurgency, and other forms of violence that accompany civil war and transnational conflict seem not to be an exception. How can the uncertainty of adversary behavior—its seemingly haphazard nature—be understood or grasped in order to better prepare or mitigate its effects? Which theoretical principles and modeling tools might be tested with available data? How can empirical findings be used to improve simulations, particularly in areas such as validation, verification, and calibration?

Modeling-based analyses can offer new insights for analysts and policymakers, and this study applies well-established concepts, principles, and models from the theory of political uncertainty and from complexity theory—the two core methodological approaches used in this study—to the analysis of conflict events during the first three years of the second Iraq War, 2003–2006, in the province of Diyala. Preliminary findings show that neither the time between attacks T or the severity of attacks S (fatalities) have a “normal” (i.e., bell shaped or Gaussian) or log-normal distribution that is characteristic of equilibrium systems. Instead, both variables showed “heavy tails” in the upper extreme range, symptomatic of non-equilibrium dynamics; in some cases approximating a power law with critical or near critical exponent value of 2. The empirical hazard force analysis in both cases showed that intensity was high for the first epoch in both variables, namely between March 2003 and June 2004, but even higher in the following period ending in March 2006. Moreover, the average empirical hazard rate clearly increased throughout the three epochs, supporting the authors’ hypothesis. Although these findings are limited to Diyala province in Iraq, and do not necessarily apply anywhere else in the country or region, Diyala province is linked to several other provinces and neighboring Iran—via the ancient strategic passage linking Khânaqin (Iraq) and Qasr-e Shirin (Iran) across the Zagros mountains.

Analysts and policymakers are always interested in understanding uncertainty, and the uncertainty of warfare continues to dominate much of the scientific modeling literature, consistent with the fundamental nature of this complex phenomenon. This common interest should be developed. In terms of political uncertainty theory applied to the analysis of war [113], [114] - that is, the first methodological approach employed in this study - Fearon [115] and others have applied similar estimation techniques to model the duration of civil wars, classified in five types, arriving at two main results. First, the “sons of soil” and contraband-financed civil war types last longer than other types (coups/popular revolutions, anti-colonial wars, and wars in eastern Europe or former Soviet Union countries). Second, the standard predictors for duration of civil wars

(e.g., ethnic diversity, GDP per capita, level of democracy, ideological effects) have a negligible effect on war duration.

In the same tradition, Bennett and Stam [116]¹ apply a parametric Weibull regression to predict the duration of the ongoing second U.S.–Iraq war. Based on a set of predictor variables—such as the strategies used and the quality of the terrain, while other factors (e.g., population, military surprise) are held fixed—the Bennett-Stam model predicts a likely duration of 83 months, or almost 7 years since the fall of the Saddam Hussein regime. However, it must be noted that such predictions, based on uncertainty-theoretic models, are probabilistic expectations, not deterministic forecasts.

In terms of complexity theory and power law analysis applied to conflict analysis—the second methodological approach used in this study—the seminal work is by Richardson [118], [119], based on his data set of “deadly quarrels,” which included international conflicts and civil wars between 1820 and 1945. An early revision of his work and a discussion of the different theories behind his empirical work is found in Rapoport [120]. Wilkinson [121] and Cioffi-Revilla and Midlarsky [122] present replications of Richardson’s results with larger and more diverse data sets.

This chapter proceeds as follows. Section 5.1.2 presents the methods used for data analysis and model testing, based on the theory of political uncertainty and social complexity theory. The essence of these methods is to use events data as signals for understanding latent, underlying dynamics that are causally responsible for observed conflict. Although the methods are statistical, mathematical, and computational, they are essentially information extraction procedures for understanding adversary conflict dynamics. The next section presents the results in technical and non-technical language. The fourth section presents a discussion of the main results and some general conclusions, including discussion of policy significance. The discussion of policy implications is innovative for the integrated multidisciplinary methods used in the analysis, which combined political uncertainty theory and complexity or complex systems theory.

5.1.2 Method

Let X denote a conflict-related random variable, such as the time-interval between attacks T (measured in days), the severity S of each attack (measured by fatalities or deaths), distance D from the previous attack, or other variables associated with an attack event. Formally, a conflict process $P(\mathbf{X}_{<\tau>})$ is modeled as an n -tuple of random variables with realizations ordered in historical time τ (so-called “epochal time” [123]), where each r.v. is defined by its set of associated probability functions $p(x)$ and $\Phi(x)$, or p.d.f. and c.d.f., respectively.

Figure 5.1 illustrates the overall methodological process used in this study, as detailed in the following sections. Empirically, our analysis is based on 2002–2006 high frequency (daily) conflict events data collected independently at the Lawrence Livermore National Laboratory by E. O’Grady [124], [125] as detailed below. We conducted synchronic analyses based on the entire population of data, as well as diachronic analyses based on epochs. In particular, we examined results based on the three data “epochs” proposed by the International Crisis Group (ICG)², based on organizational hypotheses.

¹ Their model is explained in detail in the earlier paper by Bennett and Stam, [117]

² International Crisis Group (ICG). 15 February 2006, available at <http://www.crisisgroup.org/home/>

The events data were analyzed with two distinct but interrelated types of quantitative/ computational methods: (i) hazard force analysis, founded on the theory of political uncertainty and (ii) power law analysis from complexity theory [126]. Although traditionally autonomous from each another, in this study we exploited the synergy of these two analytical methods to obtain new inferences that advance our understanding of adversary conflict behavior.

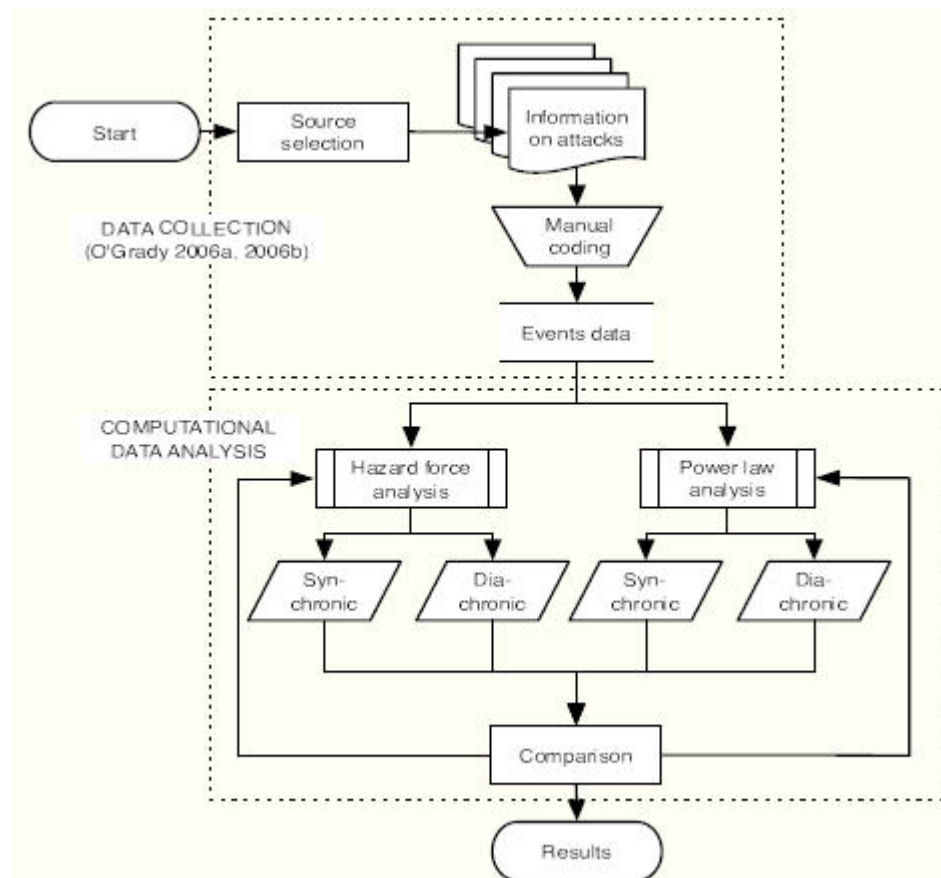


Fig. 5.1 Overall events data analysis process conducted in this study, starting with O’Grady’s [124], [125] data on attacks. Hazard force analysis and power law analysis are parallel computational data analysis processes

5.1.2.1 Data

This study used the dataset Iraq Event Database: Diyala Province on adversary conflict events recently compiled by O’Grady (2006), which is based on unclassified sources. The data set contains $N = 335$ attack events that took place in the Province of Diyala, Iraq, between March 2003 and March 2006. The two coded variables used in this study were “date of event”, used for computing time between attacks (T) and “total fatalities” (used as proxy for severity S), defined as “the total figure of people killed (includes terrorist/ insurgent and non-terrorist/insurgent).” Moreover: “The fatalities count reports the number killed in situ (in Diyala), not those that were fatally injured and subsequently passed away, e.g., in a hospital in Europe or CONUS [continental United States].” These comprise military and civilian non-terrorist plus terrorist fatalities.

O’Grady’s new events dataset is uncommon and scientifically valuable, because— inter alia—it provides a count of high-frequency daily events with fatalities (and other variables not used in this study). Most other conflict data sets that record fatalities contain only low-frequency events (e.g., wars). By contrast, most high-frequency datasets (e.g., COPDAB, WEIS, or KEDS) do not report fatalities. [127]

5.1.2.2 Analyses

As shown in Fig. 5.1 above, we analyzed the conflict events data using two distinct analytical methods, hazard force analysis and power law analysis, as described in the next subsections. Both methods were applied to the same data ($N = 335$ events) for time-between-attacks T , and severity S . In turn, each analysis was conducted synchronically and diachronically, as explained below.

Temporal Analyses

For each variable (time T and severity S) and type of analysis (hazard force and power law) we first conducted an overall synchronic analysis, based on all the data for the entire period (March 2003 to May 2006), followed by a more historically de- tailed diachronic analysis. The latter was based on three epochs of the Iraq conflict hypothesized by the International Crisis Group (ICG):

- “Phase 1”: March 2003 to June 2004. According to the ICG, this initial epoch was characterized by “competition” among insurgent groups that had only erratic coordination and little or no organizational capacity. During period I, Iraqi rebel groups were small, not very mobile and many of the first attacks signaled a lack of expertise in handling mortars or other explosive devices. Moreover, they used their small world networks (family, neighbors, mosques) to propagate by means of leaflets their message of resistance and recruit members or, in any case, to foster similar initiatives by other people. Websites—such as iraqresistance.net—were used as a channel to communicate the message to people outside their locality but also outside Iraq. This later strategy was aimed at Muslims willing to fight against the coalition forces.
- “Phase 2”: July 2004 to June 2005. During this epoch “consolidation” would have taken place within groups of attackers. During this period, small successful groups merged with others and started to apply more often a strategy of hit- and-run ‘a la guerrilla in order to avoid frontal combat. Also, an improvement regarding how to handle explosive devices and the like allowed to them to focus their attacks on specific targets.
- “Phase 3”: July 2005 to May 2006. This third epoch would have been characterized by the ICG as having increasing “confidence”—even insurgent optimism— indicative of increased organizational capacity on the part of attackers. This third period would also have been oriented towards justifying religiously their kidnaps and killings of members of the U.S. coalition, foreign civilians, and even Iraqis (mostly Shi’ites) working with the coalition.

The significance of these three “phases” (*epochs*, in quantitative conflict analysis terminology) stems from their application to the overall conflict in Iraq, applying to the whole country; they are not specific to Diyala province. The authors are not aware of any periodization specific to Diyala. The main theoretical motivation for these epochs—and additional reason why epochs matter—is that conflict dynamics, in terms of forces of onset \mathbf{F}_T and forces of severity \mathbf{F}_S , which drive the onset and severity of attacks, undergo fundamental changes across epochs due to the

increasing organizational capacity of the attackers. Hazard force analysis and power law analysis aim at detecting such latent forces, as described in the next subsections. The ICG epochs should therefore mark significant transitions within an overall politico-military process affected by these forces.

Hazard Force Analysis

The hazard force or intensity function producing the observed realizations of a conflict process $P(\mathbf{X}_{<t})$ is defined as follows.³

Definition 5.1 (Intensity function) The intensity function $H(x)$ of a c.r.v. X is defined by the ratio of the value of the p.d.f. to the value of the complementary c.d.f.

of X . Formally, $H(x)$ is defined by the equation

$$H(x) = \frac{p(x)}{[1 - \Phi(x)]}, \quad (5.1)$$

where $p(x)$ and $\Phi(x)$ are the p.d.f. and c.d.f. of X , respectively.

Note that, although the intensity or hazard force $H(x)$ is a latent or non-observable variable, equation 5.1 renders $H(x)$ measurable, because both $p(x)$ and $\Phi(x)$ can be computed from a sufficiently large set of observed realizations $\hat{x}_i \in X$.

Accordingly, by (5.1), the specific qualitative form of $H(x)$ (constant, increasing, decreasing, non-monotonic) depends directly on the form of the associated probability functions (c.d.f. or p.d.f.). Specifically, four cases are fundamentally important for analyzing attacks. To illustrate, let $X = T$, the time interval between attacks, measured—for instance—in days.

Case 1. Constant intensity: $H(t) = k$. In this special or *equilibrium* case the propensity for the next attack to occur—i.e., the hazard rate or event intensity—does not change between realizations, consistent with the notion that escalating and mitigating forces of conflict are in balance. This also corresponds to the Poisson case and simple negative exponential density, with $p(t) = ke^{-kt}$ and $\bar{t} = 1/k = \sigma^2(t)$. This case is known to have the strongest empirical support for many types of conflict, both internal and international, following Richardson's [118] pioneering work on wars of all magnitudes. In terms of the ICG epochs mentioned earlier, we expected to detect a constant or slightly decreasing intensity during Period 1 (March 2003 to June 2004), because the attackers were supposed to have been in competition among themselves and attacks were erratic.

Case 2. Increasing intensity: $dH/dt > 0$. In this case the hazard force or event intensity would increase between attacks, symptomatic of a fundamentally unstable situation where attacks occur under rising pressure or increasing propensity. This situation is akin to a driven threshold system that triggers attack event as forces build up. In terms of the ICG epochs, we expected to observe

³ The original interpretation of (5.1) as an intensity or force is probably due to D. R. Cox [128], based on Bartholomew [129]. For a more detailed description of hazard force analysis, including examples from conflict processes and computational issues, see [114] chs. 2–4, containing numerous references. Unfortunately, most of the standard social statistical and econometric literature (e.g., Greene [130]) treats the estimation of $\hat{H}(x)$ as just another case of regression, ignoring the much deeper dynamical implications used in this study.

increasing force intensity during Period 2 and (even more so) in Period 3, given the rising organizational capacity of attackers.

Case 3. Decreasing intensity: $dH/dt < 0$. In this case the hazard force or event intensity would decrease between attacks, symptomatic of a stable situation where attacks occur under diminishing pressure or decreasing propensity. This situation is akin to a leaky threshold system that dissipates forces as they build up. For example, conflict resolution mechanisms (nonviolent processes) may be responsible for dissipation and decreasing propensity for attacks. In terms of the ICG epochs, we expected to see this force pattern only in Period 1, if at all.

The above three cases are covered by the two-parameter Weibull model:

$$H(x) = kt^{\beta-1} \quad (5.2)$$

where k and β are the scale and shape parameters, respectively. Thus, the estimated exponent $\hat{\beta}$ computed directly from the data supports the follow inferences concerning the causal conflict dynamics driving the incidence of attacks:

$$\hat{\beta} < 1 : \text{decreasing conflict force }) \text{ stable situation} \quad (5.3)$$

$$\hat{\beta} = 1 : \text{constant conflict force }) \text{ borderline situation} \quad (5.4)$$

$$\hat{\beta} > 1 : \text{increasing conflict force }) \text{ unstable situation} \quad (5.5)$$

Clearly, these three conflict situations are qualitatively distinct, and from a policy perspective they obviously correspond to desirable, indifferent, and undesirable conditions, respectively. Interestingly, the mean or first moment of T is given by

$$\bar{t} = k\Gamma(1 + 1/\beta) \quad (5.6)$$

where Γ is the gamma function. Therefore, commonly used heuristic estimates based of mean values (e.g., “the average time lapsed between attacks”) are not generally valid and instead must be computed exactly because \bar{t} is notoriously sensitive to $\hat{\beta}$

Finally, a fourth qualitative case in the qualitative form of the conflict force is also interesting:

Case 4. Non-monotonic intensity. After an attack occurs, the conflict force may rise (as in Case 2), but then subside, as in a lognormal function. Alternatively, the conflict force may subside following an attack and then begin to rise again sometime after, as in a so-called “bathtub” function. These non-monotonic situations were also considered in our analysis, given their plausibility. In terms of the ICG epochs, their logic seemed mostly linear, ruling out non-monotonic forces.

Summarizing our hazard force analysis, conflict events data on time intervals between attacks (T) and fatalities produced by each attack (or severity S) were used to compute the corresponding empirical hazard functions, $H(t)$ and $H(s)$, respectively. These empirical functions were then closely examined to determine their qualitative shape and draw inferences concerning conflict conditions. This procedure was repeated for the entire population of data, as well as for each of the three ICG epochs. The initial expectation was that these estimates would yield mostly Case 1 (constant force), consistent with many earlier studies, with rising value of k as the epochs progressed (as argued by the ICG).

Power Law Analysis

Power law analysis is a complexity-theoretic method for drawing inferences from a set of conflict data. Here we used the so-called type IV power law, which is defined as follows.⁴

Definition 5.2 (Power law) *A power law of a conflict process $P(X_{<\tau>})$ is a parametric distribution model where increasing values $x_i \in X$ of the conflict variable X occur with decreasing frequency, or $f(x)/x^{-b}$, with $b > 0$. Formally, $f(x)$ in this case is a p.d.f. given by (5.7) where a and b are scale and shape parameters, respectively.*

$$p(x) = \frac{a(b-1)}{x^b} \quad (5.7)$$

From this 2-parameter hyperbolic equation for the p.d.f. it can be easily shown that the complementary cumulative density function (c.c.d.f.), defined as $1 - \Phi(x) \equiv \Pr(X > x)$ (a.k.a. survival function when $X = T$, or $S(x)$), has the following form in log-log space:

$$\log[1 - \Phi(x)] = a' - (b-1) \log x \quad (5.8)$$

which, finally, yields

$$\Phi(x) = 1 - \frac{a}{x^{(b-1)}} = 1 - ax^{1-b} \quad (5.9)$$

The penultimate expression is commonly used for empirical analysis, because it can be obtained directly from the set of observed values \hat{x}_i .

The empirical estimate \hat{b} is of interest because the first moment of a power law is given by

$$\begin{aligned} E(x) &= \int_{\min\{x\}}^{\infty} xp(x)dx = a(b-1) \int_{\min\{x\}}^{\infty} x^{1-b} dx \\ &= \frac{a(b-1)}{2-b} x^{2-b} \Big|_{\min\{x\}}^{\infty} \end{aligned} \quad (5.10)$$

which goes to infinity as b goes to 2. In other words, there is no mean size (no expected value $E(x)$ exists) for the conflict variable X (such as onset times T or severity S) when X is governed by a power law with exponent b approaching the critical value of 2, or $(b-1) < 1$ (below unit elasticity). This is an insightful theoretical result for numerous social variables, such as organizational sizes, fatalities in warfare [118], [122] and terrorist attacks. The critical threshold $b = 2$ marks the dynamical boundary between conflict regimes that have a finite average and computable size ($b > 2$) and a highly volatile regime that lacks an expected value or mean size ($b \leq 2$). This is a theoretical insight directly derived from the empirically estimated value of the power law exponent b .

⁴ Other types of power laws include the rank-size law or Zipfian, various algebraic forms, and others [131]. In this study we applied the type IV power law because in the case of conflict data (attacks) it seems to provide the most powerful complexity-theoretic inferences.

Based on previous studies, we expected that (a) T should obey the simple (one parameter) negative exponential p.d.f. of a Poisson process,

$$p(t) = \lambda e^{-\lambda t} \quad (5.11)$$

where $\lambda = 1/\bar{t}$; and (b) S should obey a power law. Moreover, with respect to the diachronic epochs (ICG periods) discussed earlier, we expected t to increase across periods (epochal time) and λ to approach criticality as the attackers gained strength.

Summarizing our power law analysis, conflict events data on time interval between attacks (T) and the severity of attacks (S) were used to compute the corresponding empirical power law functions $\log[1 - \Phi(t)]$ and $\log[1 - \Phi(s)]$, for onsets and severity (fatalities), respectively. These empirical functions were then closely examined to determine their qualitative shape and draw inferences concerning conflict conditions. We also examined the p.d.f.s directly using kernel estimation. This procedure was repeated for the entire population of data (synchronic analysis), as well as for each ICG epoch (diachronic). The initial expectation was that these estimates would yield mostly a poor fit of the power law for the overall synchronic analysis, but increasingly good fit and decreasing exponent (towards criticality) as the epochs progressed and the attackers became more organized.

5.1.3 Findings

First are presented the temporal findings for the analysis of time between attacks T (the next subsection), including both synchronic and diachronic patterns, followed by a parallel presentation of findings for the severity of attacks S (the subsection after). Table 5.1 summarizes the overall descriptive statistics for both processes, T and S .

Time Between Attacks

Overall (Synchronic) Patterns. There were 107 occurrences where $T = 0$, meaning more than one attack took place in a given day. The large and positive skewness implies that the right tail of the distribution is more pronounced. Kurtosis is substantially larger than zero, implying a leptokurtic feature. These moments suggest a distribution for T with non-normal characteristics.

Another insightful indicator is the ratio of the mean to the standard deviation, which in this case is 0.37 and closer to 0 than to 1. This could imply a hyper-exponential process or a high degree of political uncertainty, because the mean of T is significantly smaller than the variance (by a factor of 24).

Figures 5.2 and 5.3 plot the empirical c.d.f. and p.d.f, respectively, consistent with the non-normal results reported in Table 5.1.

TABLE 5.1 Onset of attacks T (days between events)

Descriptive Statistics: Whole Sample				
Mean	Variance	Skewness	Kurtosis	Mode
3.25	78.15	8.13	80.65	0
Std. Dev	Median	Max	Min	N
8.84	1	104	0	334




Fig. 5.2 Cumulative probability density for time between attacks T, Diyala Province, Iraq. March, 2003 - March, 2006




Fig. 5.3 Probability density for time between attacks T, Diyala Province, Iraq. March, 2003 - March, 2006

Both graphs suggest a distribution with a pronounced exponential pattern. Note also that by the 10th day the c.d.f. amounts to 95%, implying that short intervals between attacks are by far the most frequent. In other words, following an attack the probability of another attack is very high within intervals no greater than 10 days. The surprisability of the process, or difference between the mean and the median, yielded 2.3 days, which is another indication of the volatility of attacks.

Normality Tests. The results in Table 5.1, showing that the mean of 3.25 days is clearly lower than the standard deviation of 8.84 days, as well as the empirical distributions in Figures 2 and 3, consistently imply that the data might not be normally distributed. A formal test of normality was applied to corroborate these preliminary results. The Shapiro-Wilk test was implemented to test the null hypothesis that the data are normally distributed and the results are reported in Table 5.2.

We also tested the hypothesis of a lognormal distribution, by computing the log-transformation of T . Because the p-value is less than 5% the null hypothesis is rejected in both cases.

In addition, the variable T does not correspond to a lognormal distribution either. In both cases the probability or p-value is very small or close to zero.

TABLE 5.2 Shapiro-Wilk Test

Time Between Events					
Variable	Obs	W	V	z	p-value>z
T	334	0.3571	150.67	11.83	0.0000
Ln(T)	227	0.9425	9.56	5.23	0.0000

Hazard Forces. We applied the Kaplan-Meier method for estimating the empirical survival function $\hat{S}(t)$, and results are shown in Figure 5.4. Recall that the K-M method is non-parametric, so it does not impose a specific structure on the data. In this case the estimated value of the probability of no attack within time t should be interpreted as the product of the probabilities of not attack occurring at t and the preceding periods. In our particular case the K-M estimator tell us that the survival function for attacks drops off sharply in the days following an attack and slowly settling to zero after about 10 days.




Fig. 5.4 Empirical survival function $\hat{S}(t)$, for time between attacks T , Kaplan-Meier estimate, Diyala Province, Iraq. March, 2003 - March, 2006

Figure 5.5 shows the K-M estimate of the hazard force function. In terms of the Weibull hazard model given by (5.2), Figure 5.5 implies that $\hat{\beta} < 1$. Specifically, this empirical hazard force for the complete period is decreasing until approximately the 10th day, after which it shows some volatility around a value of 0.1. The average empirical hazard force is 0.0877, a value that will be more meaningful when we analyze the data within shorter periods (epochal, diachronic analysis).




Fig. 5.5 Diyala Province, Iraq. March, 2003 - March, 2006

Power Laws and Criticality. Our alternative hypothesis, given the rejection of the normality (or log-normal characteristics) of the data, is that the data follow a power law as defined earlier in the methodological section. A univariate regression model was carried out, in order to test the linearized power law using an ordinary least squares (OLS) procedure. The logarithmic transformation was applied to the complementary c.d.f. (or survival function) and T , as described earlier in the Methods section.

Figure 5.6 shows the scatter plot for both variables, the linear regression fit, and a 95% confident interval. Both point estimates, also inserted in the figure, are statistically significant at the 1% level. The slope estimate is the most relevant in this type of analysis, which is minus 1.03. The R^2 is 0.95 and the standard errors for the constant term and the slope are 0.0085 and 0.0153, respectively. Although the overall fit is close, there is clearly some systematic departure in the pattern for the upper range.




Fig. 5.6 The empirical complementary c.d.f. for time between attacks T in log-log space, Diyala Province, Iraq. March, 2003 - March, 2006

Epochal (Diachronic) Patterns. The null hypothesis tested was that the estimates for the hazard force would yield a roughly constant intensity within each epoch with increasing average mean values for the hazard force across epochs. Figures 7 through 9 plot the empirical hazard force functions for the three epochs.

The pattern in Figure 5.7 is not clear-cut or smooth because there are fluctuations starting at zero up to 0.2 for the first forty days, then the hazard rate drops to zero until day 95th where it spikes up to 0.45. The average empirical hazard rate for Period 1 is 0.094. Also, after the twentieth day the pattern is not very different from Figure 5.5, including even the sudden spike around day 95th.




Fig. 5.7 Diyala Province, Iraq. Period 1, March, 2003 - June, 2004. Source: Prepared by the authors based on O'Grady's (2006a, 2006b)

TABLE 5.3 Severity of attacks S (fatalities data were either normally distributed or belonged to a lognormal distribution)

Descriptive statistics whole sample				
Mean	Variance	Skewness	Kurtosis	Mode
4.17	74.0	5.19	32.66	1
Std. Dev	Median	Max	Min	N
8.6	2	71	0	335

In Figure 5.8 for Period 2 we observe a different pattern from the one in Period 1.




Fig. 5.8 Diyala Province, Iraq. Period 2, July, 2004 - June, 2005. Source: Prepared by the authors based on O'Grady's (2006a, 2006b)

However, before the twentieth day the empirical hazard rate for Period 2 is not very different from the pattern for the empirical hazard rate within the whole period (see synchronic results earlier). That is to say, in Period 2 the hazard rate decreases during the first twenty days and then spikes up beyond 0.5. This last fluctuation of the data, however, might be an artifact of the computation of the data rather than a reflection of the actual intensity of the events. The average empirical hazard rate for Period 2 is 0.212, not counting the last point beyond 0.7. This is twice as high as the average hazard rate for Period 1. Therefore, during this period from June, 2004 until June, 2005, there was a doubling in the hazard force driving attacks in Diyala province. There is no such a drastic difference between the plot in Figure 5.8 and the one for Period 3 in Figure 5.9, which spans the period from summer 2005 until March, 2006. The average empirical hazard rate in Period 3 was 0.217, not including the last point above 0.8. Thus, during this epoch there was not a substantial increase in relation to what happened in Period 2.




Fig. 5.9 Diyala Province, Iraq. Period 3, July, 2005 - March, 2006

All in all, the analysis of Diyala attacks by epochs is consistent with the increased organizational capacity hypothesized by the ICG in terms of an increase in the average empirical hazard forces.

Power Laws. Figures 5.10 through 5.12 report the results of the diachronic power law analysis for individual epochs.

Figure 5.10 for the first epoch shows a linear regression OLS slope estimate of -0.62, which is statistically significant at 1% level of confidence. The standard errors reported for the constant and the slope terms are low: 0.0374 and 0.0399, respectively, with $R^2 = 0.9$. Figure 5.11 shows results for the second epoch, with a steeper slope estimate of -1.15, also statistically significant at 1% level of confidence. The respective standard errors reported for both regression terms are: 0.0179 and 0.0366, $R^2 = 0.9$. Figure 5.12 shows the slope estimate to be -1.25, again statistically significant at 1% level of confidence. The standard errors are 0.0174 and 0.0406, respectively, with $R^2 = 0.92$. Note that the slope becomes increasingly steeper.




Fig. 5.10 Diyala Province, Iraq. Period 1, March, 2003 - June, 2004

In general, the power law for the whole period is the closest to a linear relationship between the complementary c.d.f. and T. However, all periods show some upper range bending, even if slight in some cases. We cannot make a formal test to determine if the slope coefficients for every linear regression are equal because of the difference in the number of observations.




Fig. 5.11 Diyala Province, Iraq. Period 2, July, 2004 - June, 2005




Fig. 5.12 Diyala Province, Iraq. Period 3, July, 2005 - March, 2006

Severity of Attacks

Overall (Synchronic) Patterns. Table 5.3 shows a summary of the descriptive statistics for the severity S of attacks. The statistical properties for S are not very different from those discussed earlier for the time between events (T). We observe positive skewness and kurtosis, again meaning that we should find a pronounced right tail and leptokurtic distribution. The ratio between the mean and the standard deviation is 0.485 which also suggests a non-normal pattern. The mode is 1 and the median is 2. That is to say, the most typical number of fatalities produced by an attack was one death.

The empirical c.d.f. and p.d.f. are plotted in Figs. 5.13 and 5.14, respectively. We observe the similarity with the respective figures for T . The pronounced right tail with a few values at the end of the distribution, or “dragon tail” indicating the presence of extreme events with unduly high frequency/probability.




Fig. 5.13 Diyala Province, Iraq. March, 2003 - March, 2006




Fig. 5.14 Diyala Province, Iraq. March, 2003 - March, 2006

Normality Tests. Table 5.4 shows the results of our Shapiro-Wilk test for the normality of attack severity S (fatalities). As before, the null hypotheses were that the data were either normally distributed or belonged to a lognormal distribution.

Because both p-values are less than 5%, or even 1%, the null hypothesis can be rejected in both tests. These results provide more confident about our previous claim that the data for severity S is not normally distributed or even belong to a lognormal distribution. This provides additional justification for the power law analysis.

TABLE 5.4 Shapiro-Wilk Test

Severity of attacks S (fatalities)					
Variable	Obs	W	V	z	p-value>z
Fatalities	311	0.4846	113.37	11.12	0.0000
Ln(Fatalities)	241	0.9571	7.54	4.69	0.0000

Hazard Forces. Figures 5.15 and 5.16 show the K-M estimate for the c.c.d.f. and the hazard force, respectively. The K-M estimate for the c.c.d.f. for severity S reflects the cumulative probability of an additional fatality beyond a given level s . For the whole period, that probability is less than 0.25 beyond the first five fatalities and it decreases faster than the Kaplan-Meier curve for the variable time between events T . Beyond $S = 20$ deaths the probability is very close to zero, although the hazard force highlights the probability of extreme events.




Fig. 5.15 Empirical complementary cumulative probability function for severity of attacks S (fatalities), Kaplan-Meier estimate, Diyala Province, Iraq. March, 2003 - March, 2006




Fig. 5.16 Diyala Province, Iraq. March, 2003 - March, 2006

The empirical hazard rate for severity S , shown in Figure 5.16, starts off at a value close to 0.4 and decreases steadily by the twentieth day, after which it fluctuates around 0.1. The average empirical hazard rate for the whole period is 0.1472. These results have added uncertainty, because in this series there are 24 events with missing data. In general, the intensity in fatalities decreases in Diyala up to a value of around 20, after which it fluctuates with spikes around 50 and 65—not exactly well-behaved. Power Laws and Criticality Figure 5.17 shows the plot of the empirical c.c.d.f. of S in log-log space, including the observed data points, the best-fitting OLS line, and 95% confidence intervals.




Fig. 5.17 Empirical c.c.d.f. of severity S (fatalities) in log-log space, Diyala Province, Iraq. March, 2003 - March, 2006

The linear regression fitted curve through the OLS approach yields a slope estimate of -0.94 , which is statistically significant at the 1% level of confidence. Unfortunately, once again, the upper range falls off exponentially. Nonetheless, the standard errors reported are: 0.0098 and 0.0158 for the constant and slope coefficients, respectively; and the R^2 is 0.94. Overall, while the data are not normally-distributed, they also fall short of a perfect fit to a power law, indicating perhaps another fat-tailed distribution.

Epochal (Diachronic) Patterns. In Figures 5.18, 5.19, and 5.20 we report results the three epochs, in a similar way as for time between attacks (T). We observe a generally decreasing and fluctuating pattern from approximately 0.4 to 0 in Period 1 (without taking into account the last computed value that climbs to almost 1 due to rounding errors). The average in this first epoch was 0.3737 (omitting the last point). This pattern is quite different from the empirical hazard rate for the complete period in Diyala, but it could be due to fewer observations. On the other hand, a similar pattern in Period 2 to the empirical hazard rate of the whole series is observed in Figure 5.16. Its average hazard rate was 0.1719, again higher than the value for the whole series. And lastly, in Period 3 the average hazard rate is 0.1852, slightly higher than in the second epoch. Therefore, the average hazard force for severity S (fatalities) dropped substantially from Period 1 to Period 2, but then increased slightly in Period 3—a pattern not consistent with the organizational dynamics hypothesized by the ICG. This pattern in overall force mitigation may have been due to increased effectiveness of the coalition forces in Periods 2 and 3 relative to Period 1.




Fig. 5.18 Diyala Province, Iraq. Period 1, March, 2003 - June, 2004




Fig. 5.19 Diyala Province, Iraq. Period 2, July, 2004 - June, 2005




Fig. 5.20 Diyala Province, Iraq. Period 3, July, 2005 - March, 2006

Power Laws and Criticality. Figure 5.21 shows the results for the power law analysis in Period 1: the slope estimate is -1.15 , which is statistically significant at the 1% level of confidence. Standard errors reported for the constant and slope coefficients are: 0.0427 and 0.0895, respectively; and the R^2 is 0.88. In Fig. 5.22 for Period 2 the slope estimate is -0.95 , also statistically significant at 1% level of confidence. Standard errors reported for the constant and slope coefficients are: 0.0108 and 0.0185, respectively; and the R^2 is 0.96. In Figure 5.23 the slope came to -0.85 ,

again statistically significant at 1% level of confidence, standard errors for the constant and slope were 0.0198 and 0.0299, respectively, and the R^2 is 0.9. In general, Periods 1 and 2 clearly show the best fits to a power law, although here again the very highest values tend to deviate. Thus, these two epochs might be reflecting a similar evolution to the one observed for the complete period previously. It is not feasible to make a formal test to compare the slope coefficients across epochs and the whole period due to the difference in the number of observations, however, they indicate a general movement toward a flatter and hence more lethal extreme range.




Fig. 5.21 Diyala Province, Iraq. March, 2003 - June, 2004




Fig. 5.22 Diyala Province, Iraq. July, 2004 - June, 2005




Fig. 5.23 Diyala Province, Iraq. July, 2005 - March, 2006

5.1.4 Discussion

The findings reported in this article suggest new insights and implications for research and policy. The utility of these findings is to illuminate the political and military context of conflict, and to address the questions raised in the introduction. The following discussion focuses on the main findings and selected policy implications.

Main Empirical Findings

The results for the onset of attacks T (time between events) in the analysis of overall synchronic patterns showed a non-normal distribution with a heavy right tail. The formal normality tests (Shapiro-Wilk) also rejected the null hypothesis for the presence of a lognormal distribution in the data. The empirical c.d.f. and p.d.f. both allow one to visualize this non-normal pattern in the distribution of T . These statistical properties suggest a high degree of political uncertainty, far from the equilibrium conditions of normality with marked central tendency. The Kaplan-Meier estimate of the survival function $\hat{S}(t)$ demonstrated that T has a higher probability of realizing very short time spans between attacks, with rapidly increasing cumulative probability (much faster than Poisson). In addition, the empirical hazard force function showed that the intensity of the force for attacks to take place decreased up to approximately the tenth day, after which this intensity fluctuated below 0.1. The average hazard rate for the complete period was 0.088 attacks/day, but varied across epochs.

The power law analysis of onset times T yielded a point estimate for the slope of the inverse relationship between the c.c.d.f. and T of 1.03, with a statistically significant 1% chance of being wrong. This is basically a perfect inverse relation between these two logarithmic variables. More importantly, the exponent is therefore 2.03, which is critical given the usual level of imprecision in these data. The complexity-theoretic implication of this finding is that, for the overall period,

extreme time spans are far more likely than would be normally assumed, and (by Equation 5.12) the first moment is practically nonexistent from a theoretical perspective.

For the *epochal diachronic patterns* the authors found that for the first epoch of T the empirical hazard rate did not evolve in a clear-cut fashion, fluctuating basically around 0.1 until the fortieth day and going to zero thereafter. During Period 1 the average empirical hazard rate was 0.094, which is close to the level for the overall period. For the second and third epochs, however, this average increased to around 0.21, which is at least twice as the early period. In both epochs the hazard rates were decreasing until the twentieth day. In terms of the initial hypotheses, these results are generally supportive.

All the slope estimates of the power law analyses of T for the three epochs were statistically significant at the 1% level of confidence, with high values of R^2 and—more importantly—very small standard errors. For the first epoch the slope estimate was -0.62 , which is critical, -1.15 for the second epoch, and -1.25 , or away from criticality for the third and last epoch. However, the authors also note the occurrence of systematic deviations of the highest values, down from the theoretically expected fit of the power law. Most likely, the upper tail for the distribution of T was exponential, consistent with earlier literature, not power law.

Results for the severity of attacks S (fatalities) also resemble in some ways the nonnormality characteristics of T for the analysis of *overall synchronic patterns*. Severity showed a pronounced right tail according to its skewness (5.19) and the empirical p.d.f. plot confirmed such a pattern. The mode of severity was one casualty per attack in Diyala. Furthermore, according to the Shapiro-Wilk test the statistical distribution of S did not belong either to a normal or log-normal distribution, which is also consistent with the fat tail. A Kaplan-Meier estimate constructed for S also showed an overall similar pattern as that found for T . However, in this case the c.c.d.f. decreased even faster and was less than 25% after the first five fatalities. The corresponding empirical hazard force of S for the whole period started at 0.4 and decreased steadily until about the twentieth day, after which it fluctuated below 0.1 with an average value of 0.1472.

The power law analysis of S for the whole period yielded a slope estimate of -0.94 , which was also statistically significant at 1% level of confidence with a high R^2 value of 0.94. This demonstrated an almost perfect inverse relationship between the c.c.d.f. and S . The *epochal diachronic patterns* for the three epochs of S showed a decreasing and fluctuating pattern from 0.4 to zero in Period 1. This pattern was different from the empirical hazard rate for the complete period in Diyala. However, a similar development in Period 2 to the empirical hazard force of the whole series was observed in the middle panel of Table 4. And, lastly, Period 3 also showed a somewhat similar process to Period 2 but with higher values at points close to 40 fatalities. All in all, the respective averages in these epochs for S were higher than for the overall period.

Finally, the slope estimates for each of the three epochs in the power law analysis of the severity of attacks hovered around the critical value of 2.0 (2.15, 1.95, and 1.85, in chronological order). All of them are statistically significant at 1% level of confidence and high fit. Compared to the whole period, epochs 1 and 2 seemed to be closer to the process of the overall period.

None of these findings are available through plain observation or even field visits to Iraq. Although more traditional methods provide significant information of a different nature, these analytical results provide reliable insights concerning conflict dynamics. Such insights shed new

light on insurgent activity and underlying processes. As such, these insights can help inform policymakers on the effectiveness of policies implemented or under consideration.

Policy Implications

The following discussion of policy implications moves from some basic aspects of theoretical science in applied domains to institutional issues. Throughout, the science–policy nexus dominates the discussion, but several important themes are only summarized due to space limitations.

To begin, the scientific principle according to which “*there is nothing more practical than a good theory*” (Lewin, [132]) is or should be as valid for conflict analysis as it has been for social psychology—a science that evolved from humanistic origins dating back to Aristotle. In fact, as Vansteenkiste and Sheldon [133] have noted, Lewin intended to convey a two-way relationship between scientists and practitioners, such that the two would gain from each others’ insights and specialized familiarity with information, issues, and methods—as well as toolkits. Whereas computational conflict scientists could and should develop research that yields more actionable results, practitioners could and should make greater use of available scientific progress, including viable areas of social science. The difficulties for each are many but the potential payoff is significant.

Kline’s thesis is as true for conflict scientists as it is for physicists—some of whom, such as L. F. Richardson (founder of scientific conflict analysis) have made contributions to the science of conflict. Another way to appreciate the power of scientific approaches to conflict analysis is by recalling a thesis formulated by the late mathematician Morris Kline [134] that scientists do not learn mathematics for its own sake, but because mathematics provides a unique and powerful method for discovering fundamental features of the real empirical world that are not accessible through other methods—including direct observation, measurement, or experience. Gravity, pressure, and radiation are among the many natural phenomena that are understood through the exclusive medium of mathematics, even when one can observe their effects. Much the same is true of the conflict features revealed by the medium of theories such as those applied in this study. Conflict hazard rates (the latent intensity for attacks), half-life (the greater-than-even-odds tipping point for attacks to occur), and criticality (the phase transition to an extreme threat environment) are specific features of adversarial attacks that are known exclusively through the medium of mathematics, not through direct experience or plain observation.

Within a politico-military context, the *situational awareness dashboard* of conflict analysts and policymakers could be significantly enriched by adding newpanels for viewing computational indicators, such as those applied in this analysis or others with comparable theoretical foundation. For example, application of these methods soon after the ICG Phase I (i.e., after March 2003) would have revealed the gathering momentum of the insurgency (at least in Diyala), perhaps in time to have avoided the entrenchment and maturation of effective insurgent networks by reformulating an appropriate policy.⁵ To use an analogy, such latent indicators—based on political uncertainty theory, social complexity theory, and other mathematical or computational social science theories—are akin to measuring pressure changes before the onset of a storm, or radiation prior to blast pressure. Further testing of such indicators is necessary, now that theoretical and methodological foundations exist. A better dashboard—or “computational radar screen”—could help policy analysts and practitioners navigate with reduced risk through complex threat environments where traditional assessments have proven to be insufficient.

Although this study was conducted post-hoc, by necessity, *real-time or near real-time analysis* of uncertainty and complexity models is becoming increasingly feasible. This is also significant within a politico-military context. Already the increased interest in open source data and analysis on the part of the intelligence community is stimulating a new generation of information processing tools that will one day provide real-time capabilities in events analysis and related methodologies [135]. In addition, the merging of real-time facilities with advanced data visualization and cartographic tools (e.g., social GIS, spatial social science models)—combined with Moore’s Law—will soon render feasible information awareness environments that would have been close to unthinkable just a few years ago. Real-time events data analysis will provide significant support not just for intelligence analysts but also for planners, decision makers, and others that can benefit from feedback.

Besides these improvements, *sequential event process modeling of attacks*—such as for suicide bombings or Improvised Explosive Device (IED) attacks—could prove helpful for practitioners, as well as challenging from a scientific perspective. For instance, a detailed empirically based event process model (sometimes known as a “business model” in organizational theory) of IED attacks could shed significant light on the attackers’ vulnerabilities, by revealing actionable information that a defender could exploit to prevent attacks or mitigate their effects. Models like this already exist for weapons of mass destruction ([136], chap. 15); they should be developed for a broad variety of insurgency and irregular warfare attacks. More specifically, event process models should focus on phases in the overall life cycle of an attack:

1. Decision making: Attackers deciding to act, including cognitive processes and alternative choice mechanisms;
2. Planning: Attackers organizing the schedule for implementing the attack, including operational security;
3. Preparation: Attackers coordinating the tasks necessary to execute the attack;
4. Execution: Attackers carrying out the attack that causes undesirable effects for the defender;
5. Effects: Consequences affecting the defender;
6. Recovery: Defenders restoring partially or fully restoring their condition, including socio-psychological aspects;
7. Investigation: Defenders engaging in a fact-finding campaign to apprehend attackers and their confederates; and, last but not least;
8. Prosecution: Defenders apprehending and processing attackers through the criminal justice system.

The simple fact that the operational causal structure of an attack’s processes is serialized—not parallelized—holds fundamental and inescapable policy and practical implications: all serialized behavior is vulnerable to disruption by elimination of one or more necessary conjunctions. Effective defenders must therefore learn how to exploit the inescapable serialization of an attacker’s process—by making the difficult life of insurgents almost impossible or as difficult as possible.

Of course, when it comes to the complex conflict dynamics of insurgency and asymmetric warfare, another important consideration within a politico-military context is that *not all the ne-*

cessary conflict science is known—not even for selected regions of the world or for subsets of actors—and much will remain unknown for a long time, even as better data and better theories are developed and become available to the policy community. But this situation in the politico-military domain of national security is not different from what occurs in medicine, engineering, or economics; and yet, public policy in these areas does attempt to draw on the best existing scientific understanding. Understanding what one does not know is as important as mastering what one does know.

It is important to increase the *availability and desirability of scientific knowledge on conflict*. The main findings from this study—summarized in the previous section—offer some new insights that are worth considering in the domain of policy analysis and planning. This study—and others like it that apply computational social science approaches to the analysis of real-world conflict events [137], [138], [139], [140], [141]—begin to indicate that some new systematic approaches could eventually become available to policy analysts and practitioners. Much remains to be demonstrated, but some evidence of increasing relevance is already available.

The *specific policy relevance of findings* such as those reported in this study of uncertainty and complexity patterns in adversary behavior must be judged directly in terms of new and testable insights and understanding. These may eventually permit different courses of action, or validation of policies that have been enacted on the basis of different criteria. For example, the hazard force analysis is capable of illuminating the conflict process by revealing phases of stability and instability that are otherwise not directly observable, even through the direct measurement of trends in attack frequencies or fatalities. Likewise, power law analysis can extract signals—such as the trajectory of the exponent in Equation (5.2)—capable of detecting the transformation of a threat environment or the increased likelihood of extreme attacks. Again, the application of these methods on a real-time or near real-time basis soon after March 2003 would have revealed the same gathering momentum as this study—conducted several years later. *The deteriorating conditions detected by power law exponents on the right-hand panels of Tables 5.5 and 5.6 provide unambiguous signals of an increasingly dangerous threat environment, indicating the increasing need for a counterinsurgency campaign that should have begun back in early 2004 at the latest—as opposed to three years later.* Moreover, such policy-relevant indicators could have been scrutinized by the scientific community, just like scientists discuss indicators and other metrics in numerous fields of public policy ranging from environment to health.

Besides anticipating the rise of the insurgency in Iraq, deteriorating hazard forces and increasing criticality could have anticipated the process of *ethno-sectarian segregation and humanitarian crisis with refugee flows* within Iraq as well as to neighboring countries. This is because, based on well-established concepts and principles of social science, social segregation—not just in situations like those in Iraq, but also in many urban areas—is an emergent collective phenomenon that is driven by many individual localized decisions that depend on tolerance for ethnic or sectarian diversity. In turn, such tolerance depends on trust and social bonds of reciprocity, collaboration, and expectations in terms of time horizon. When violence increases—as it did with incipient insurgency—fear in the populace also increased, leading to mistrust (ethnically diverse but formerly trusted neighbors can no longer be trusted), which leads to movement to regain security, which results in a collective pattern of segregation. Although the long chain of events may give the appearance of a Rube Goldberg process, the social scientific understanding of segregation processes has solid foundations in the pioneering work of Thomas Schelling [142] and others. Today, agent-based models of segregation offer unique and powerful computational

tools for understanding ethno-sectarian segregation in irregular conflicts and—with added necessary refinements—for exploring and designing better preventive and mitigating policies. Some [135] have recently argued that one desirable course of action would be a comprehensive thrust to increase the policy relevance of scientific conflict analysis to increase national capability in this area—and in a timely fashion consistent with due scientific processes concerning testing, replication, peer review, publication, and other quality control mechanisms. This too, like Lewin’s adage, is a *two way interaction between science and policy*: The computational social science of conflict can benefit from greater exposure to policy concerns (not limited to national security), and policy analysts and practitioners can benefit from new insights and understanding derived from science. The science of conflict (and peace) will always benefit from direct challenges originating from the policy community, and—vice versa—the national security policy community will benefit from advances in the relevant areas of social science that investigate conflict.

Admittedly, *practical policy solutions unfounded in science can sometimes suffice*, assuming some luck. Indeed, the Romans were able to build bridges that were sufficiently reliable to advance their military and strategic purposes—and indeed many Roman bridges are still intact and fully operable today—without any scientific understanding of the true laws of mechanics. Although this is certainly true—one does not need a complete science of conflict to improve current performance against adversaries—there is no denying that modern bridges built by modern science and engineering have vastly superior performance characteristics than their earlier Roman counterparts. The same is true for designing more effective counterinsurgency policies: much can be gained in terms of experience and other practical data, but a great deal more can be attained by exploiting scientific knowledge based on testable ideas and valid theories.

Ultimately, *scientific analysis of adversary threat environments can provide alternative views and insights that add value*, based on replicable methods and inter-subjective standards that are less personal or affected by biases. As well, the growing body of scientific knowledge about conflicts of many kinds—not just the insurgency and irregular warfare type of attacks examined in this study—might yet find its way into the policy process, much in the same way as knowledge from the economic sciences and the biological sciences has contributed to better economic policies and public health policies, respectively. Such a prospect leads to a final point concerning policy dimensions of scientific approaches to conflict analysis.

From an *institutional perspective*, the national security policy of the American polity— comprised of foreign and defense policies—is distributed across a number of departments and agencies; components of the national system of government. However, the distribution of science and engineering expertise or receptivity across these components, or even within them, is far from even. Some government institutions are more appreciative of science than others. The result of this uneven landscape is not only a differential appreciation for science across departments and agencies, but cultural and attitudinal differences that render the adoption of scientific methods and greater systematic rationality problematic in some quarters—especially those affected by ideology. C. P. Snow’s “two cultures” coexist, often under considerable stress, throughout many areas of the national security establishment— including the legislative branch. Advancing the role of science in the area of national security policy is a complex organizational process that involves not only scientists and practitioners, but the institutions and norms within which they operate. The same is true in allied countries that share similar concerns to America’s.

5.1.5 Summary

Neither time between attacks T or severity of attacks S (fatalities) have a normal or log-normal distribution. Instead, both variables showed heavy tails, symptomatic of non-equilibrium dynamics, in some cases coming close to approximating a power law with critical or near critical exponent value of 2. The empirical hazard force analysis in both cases showed that the intensity was high for the first occurrences in both variables, namely between March 2003 and June 2004. Moreover, the average empirical hazard rate clearly increased throughout the three epochs, supporting the article's main hypotheses. These findings—and the underlying theoretical approach and methodology— demonstrate the potential value of adversarial models for conflict analysis. From an applied policy perspective, the article highlighted the additional knowledge contributed by these kinds of analysis, including the fact that real-time or near real-time implementation of these methods could have revealed the surge of insurgents in Diyala, Iraq, relatively soon after March 2003. These and related methods from the computational social science of conflict should be viewed within the broader context of science and policy.

5.2 Timed Influence Nets Applied to the Suppression of IEDs in Diyala, Iraq

Lee W. Wagenhals and Alexander H. Levis

5.2.1 Introduction

A case study was developed to demonstrate the capability of Timed InfluenceN to develop and analyze courses of action. The specific issue that the case study addressed was stated as follows: given a military objective and a set of desired effects derived from statements of commander's intent, develop and analyze alternative courses of actions (COAs) that will cause those desired effects to occur and thus achieve the military objective. Specifically, the case study demonstrated the use of a TIN tool called Pythia that has been developed at George Mason University. This demonstrated the use of the tool to create knowledge about an adversary and the population that potentially supports or resists that adversary and the use of the TIN to analyze various COAs.

A scenario was chosen based on the problem of suppressing the use of Improvised Explosive Devices (IEDs) in a specific province of Iraq, denoted as province D in the year 2005. Specifically, it is assumed that IED incidents have increased along two main east-west routes between the capital town C of the province and a neighboring country M. Both roads are historically significant smuggling routes.

There were hundreds of documents about Iraq in general and D province in particular that were reviewed to get a better understanding of the situation. The province includes substantial fractions of Kurdish, Shia, and Sunni populations as well as other minorities. It was noted that the northern route was in the predominantly Kurdish region and the southern route was in a predominantly Shia region. A dynamic tension existed between these regions particularly with regard to the flow of commerce (overt and covert) because of the revenue the flow generated. It was noted that some revenue was legitimate, but a significant amount was not and was considered covert. Increased IEDs in one region tended to suppress the trade flow in that region and caused the flow to shift to the other. Consequently, each region would have preferred to have the IEDs suppressed in its region, but not necessarily in the neighboring region. The IED perpetrators needed support from the local and regional populations as well as outside help to carry out their attacks. The support was needed for recruiting various individuals to help manufacture the IEDs and to carry out the operations necessary to plant them and set them off. It was postulated that improving the local economy and the quality of the infrastructure services would reduce the local and regional support to the insurgents. Of course, this required effective governance and willingness on the part of the workers to repair and maintain the infrastructure that in turn required protection by the Iraqi security and coalition forces.

5.2.2 Model Development

With this basic understanding, the following steps were taken to create the TIN. First the overall key effects were determined to be:

- 1) IED attacks are suppressed on routes A and B (note these were modeled as separate effects because it may be possible that only one of the routes may have the IED attacks suppressed),
- 2) Covert economic activity improves along each of the two routes.
- 3) Overall overt economic activity increases in the region.

- 4) Insurgent fires are suppressed,
- 5) Local support for the insurgents exist and
- 6) Regional support for the insurgents exists.

Nodes for each of these effects were created in the Pythia TIN modeling tool. It was noted that suppression of IED attacks on one route could have an inverse effect on the covert economic activity on the other, but each could improve the overall overt economic activity. The suppression of the insurgent fires positively affected both covert and overt economic activity.

The next step was to identify the key coalition force (Blue) actions that would be evaluated as part of the potential overall COA. To be consistent with the level of model abstraction the follow high level actions were considered: 1) Blue coalition forces (CF) exercise their standard Tactics, Techniques, and Procedures (TPPs) (including patrols, searches, presence operations, and the like). 2) Blue Coalition Forces actively conduct surveillance operations. 3) Blue CF actively conduct Information Operations. 4) Blue CF continue to train the local Iraqi security forces and police. 5. Blue CF broker meetings and discussions between various Iraqi factions (Green).

Of course, it is not possible to just connect these actions to the key effects and, therefore, several other sub-models were constructed and then linked together to produce the final model. These models include a model of the process the insurgents must use to conduct IED operations, a sub-model for the infrastructure and economic activity, and a sub model of the political and ethno-religious activities. In addition, it was recognized that the region was being influenced by outside sources, so these also were added to the model.

The sub model of the insurgent IED activities was based on the concept of how the insurgents develop an IED capability. They must have the IEDs, the personnel to carry out the IED operation, the communication systems to coordinate the operation and the surveillance capability to determine where to place the IED and when to set it off. Each of these in turn requires additional activities. For example, the personnel must be recruited and trained. The IEDs must be manufactured, and this requires material and expertise. Furthermore, the insurgents must be motivated to use their capability. Much of this capability relies on support by the local and regional population and funding and material from outside sources. The nodes and the directed links between them were added to the TIN model to reflect the Insurgents' Activities.

The economic and infrastructure sub-model included nodes for each of the main essential services: water, electricity, sewage, health, and education. It also included financial institutions (banks, etc.) and economic activities such as commerce and retail sales of goods. The nodes for the economic and infrastructure aspect of the situation were linked to the local and regional support as well as to the overall effect on the overt economic activity.

Of course, the economic and infrastructure services will not function properly without the support of the Political and Ethno-Religious entities in the region. Thus a sub-model for these factors was also included. To do this, three facets of the region were considered: the religious activities including Shia, Sunni, and Kurdish (who are either Shia or Sunni) groups, political party activities (Shia, Sunni, and Kurdish), and the Shia, Sunni, and Kurdish activities within the government structure including the civil service and the police and law enforcement institutions. The nodes for all of these activities were created and appropriate links were created between

them. Links were also created to other nodes in the model such as local and regional support of the insurgents, economic activity and infrastructure development.

Finally, the outside influences were added to the model. These include external support for the insurgents, anti-coalition influences from neighboring countries, and external financial support for the local government and the commercial enterprises of the region. All of these nodes were modeled as actions nodes with no input links. With this model design, analysts could experiment with the effects of different levels of external support, both positive and negative, on the overall outcomes and effects.

The complete model is shown in Fig. 5.24. The model has 62 nodes, including 16 nodes with no parents, and 155 links.

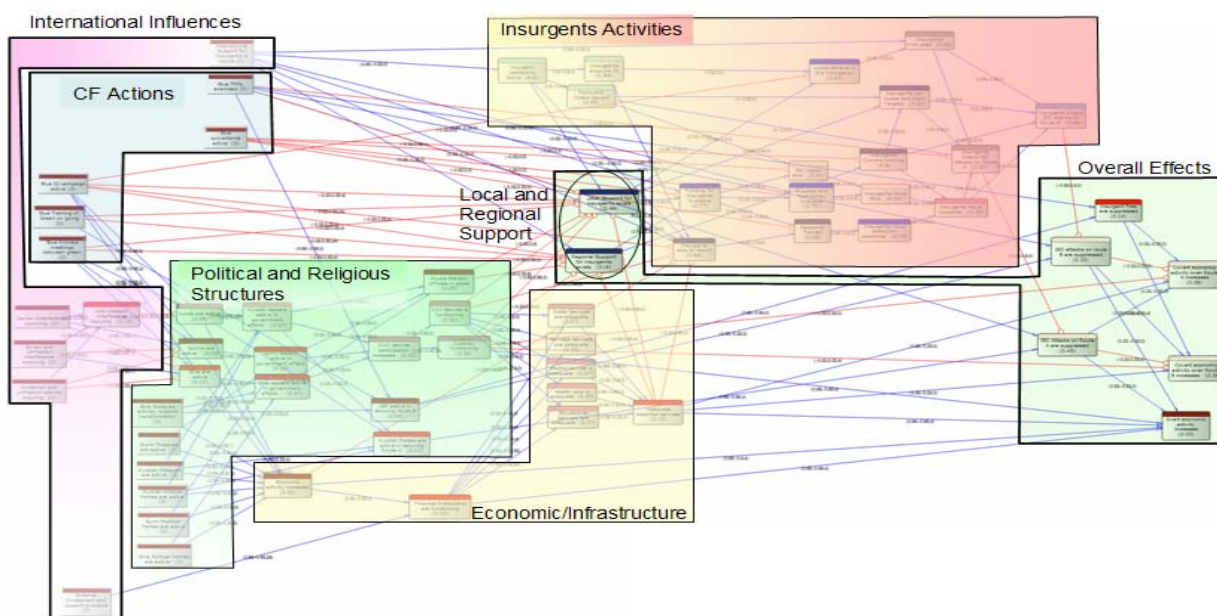


Fig. 5.24 Complete model of the case study TIN

Once the structure of the models was completed, the next step was to assign the values to the parameters in the model. This was done in two steps. First, the strengths of the influences (the g and h parameters on each link) and the baseline probability of each node were selected. This may seem like a daunting task given the subjective nature of the problem and the number of links and nodes. However, TINs and the Pythia tool limit the choices that can be made for these parameters. For each link, the model determines the impact of a parent node on a child node first if the parent is true and then if the parent is false. The choices range from very strongly promoting (meaning nearly 100%), strong (quite likely, but not 100%), moderate (50% or greater, but less than strong), slight (greater than 0% but not likely), or no effect. The modeler can also select a similar set of inhibiting strengths ranging from very strongly inhibiting to no effect. The second set of parameters is the baseline probabilities of the node. These are set to a default value of 0.5 meaning that the probability of the node being true is 0.5 given no other influences or causes (we don't know). In many cases, the default value was selected.

At this point it is possible, if not prudent, to perform some analysis on the model to observe its behavior. We will describe this in detail shortly. The final step in creating the TIN model

was to assign the temporal parameter values to the nodes and the links. The default value for these is 0. With all values set to 0 the model is identical to an ordinary Influence Net. The process for assigning the time delay values is similar to that for assigning the strengths of the influences and the baseline probabilities. For each link, the modeler determines how long it will take for the child node to respond to a change in the probability of the parent node. In some cases the change is instantaneous, so the default value of 0 is appropriate. In others, a time delay may be expected. Part of this process requires that the modeler establish the time scale that will be used in the model and thus what actual time length of one unit of delay is. Any unit of measure can be selected from seconds to days, weeks, months or even years. In this particular model each time delay unit was set to be one week. In setting the time delay of the arcs, it may also be useful to set the time delay of the nodes. Again the default value for this delay is 0. This delay represents processing delay. It reflects the concept that if there is a change in one or more of the parent nodes, once the child node realizes that the change has occurred, there may be some time delay before it processes this new input and changes its probability value.

5.2.3 Model Validation

Once the complete TIN was created, a validation of the model was undertaken. This was done by consulting with several subject matter experts who had been in the region and were familiar with the situation. Each node and link was checked to see if the node and the relationships to and from that node made sense. In short, we were confirming that the overall structure of the model made sense. Several suggestions were made and the changes were incorporated. Once the structure had been vetted, then the parameters were checked. This was done link by link and node by node. First the strengths of the influences were checked, then the baseline probabilities, and finally the time delays.

5.2.4 Analysis

Once the TIN model was finished and validated, two levels of analysis were accomplished to demonstrate the utility of the approach. The first level is the logical level. This can be done without using the parameters because it only requires the structure of the model. At this level of analysis the model shows the complex causal and influencing interrelationships between Blue CF, the external influence, the religious and political factions, the adversary (Red), and the local and regional population (Green). This particular model shows that while Blue CF has some leverage, there are many other outside influences that also can affect the outcome of any actions that Blue may take. The model identifies these influences and how they may help inhibit the progress that is made as a result of Blue CF actions. Furthermore, the model shows relationships between the actions and activities of major religious and ethnic groups and effects on government activities (police, judiciary, public works and service, etc.). It shows the impact of the adequacy of government and public services on support of the insurgency. It captures the IED development, planning, and employment processes and the impact of the other activities, the status of public services, and coalition interventions on those processes. Finally the model captures interaction of IED attack suppression on two major trade routes (suppressing one route increases attacks on the other). In short, the model has captured Blue's understanding of a very complex situation and can help articulate concepts and concerns involved in COA analysis and selection.

The second level of analysis involves the behavior of the model. It is divided into a static quantitative and a dynamic temporal analysis. The static quantitative analysis requires the structure of the model and the non temporal parameters to be set. The temporal, time delay param-

ters should be set to the default value of 0. This analysis enables one to compare COAs based on the end result of taking the actions in the COA. In the Province D model, four major COAs were assessed as shown in Fig. 5.25. This table has four parts, an Action stub in the upper left corner, the Action or COA matrix to the right of the Action stub, an Effects stub below the Action stub, and the Effects matrix adjacent to the Effects stub. In the COA matrix, the set of COAs that have been evaluated are listed with an X showing the actions that comprise the COA. The Effects matrix shows the corresponding effects as the probability of each effect.

Actions	Situation (COA) 1	Situation (COA) 2	Situation (COA) 3	Situation (COA) 4
International Interference	X	X	X	X
External Financial Support		X	X	X
CF TTPs and Surveillance		X	X	X
CF IO, training, brokering			X	X
Iraqi political and religious group participation				X
EFFECTS				
Local and Region Support for Insurgents Exists	0.97	0.92	0.26/0.36	0.22/0.14
IED Attacks Suppressed on Route A / B	0.17/0.15	0.31/0.34	0.67/0.68	0.85/0.74
Insurgent's fires suppressed	0.14	0.65	0.9	0.93
Public services adequate	0.12	0.39	0.39	0.55
Overt Economic Activity Increasing	0.02	0.08	0.31	0.89
Covert Economic Activity Increasing along routes A and B	0.37	0.50	0.56	0.57

Fig. 5.25 Static Quantitative COA Comparison

COA 1 was a baseline case in which only international interference and support to the insurgency occurs. There is no action from the Blue CF, no external financial support to the infrastructure and the economy, and the religious and political factions are not participating in the governance of the area. The overall effects are shown in the lower part of the matrix. The results for this COA are very poor. There is support for the insurgency and it is very unlikely that the IED attacks will be suppressed on either route. With an ineffective local government, the basic services are inadequate which encourages the support to the insurgency and there is little chance for economic increase.

COA 2 represents the case where external financial support is provided and the coalition forces are active both in presence operations and in conducting surveillance. However, Information Operations, training of Iraqi forces and workers, and brokering of meetings and agreement between Iraqi factions are not occurring. In addition, the political and religious groups are not participating in positive governance and support to civil service. In this case, there is some improvement compared to COA 1, but still there are many problems. Local support for the insurgents is still very strong, although there is some suppression of the IED attacks and insurgent fires due to the activities of the coalition forces. As a result there is some improvement in public services and an increase in covert and overt economic activity, due in part to the reduction in IED attacks and insurgent fires.

The third COA contains all of the actions of COA 2 plus the addition of coalition force information operations, training of Iraqi security and police forces as well as civilian infrastructure

operations and significant brokering of meetings and agreements between the various Iraqi agencies and factions. The result is a significant improvement in the suppression of the IED attacks and insurgent fires due to the improved capabilities of the Iraqi security and police forces and the significant drop in the local and regional support of the insurgents. There is also a significant improvement in the covert and overt economic activity. However, there is little change in the adequacy of the public services, due primarily to the lack of effective participation of the Iraqi governance function.

The last COA has all actions occurring. In addition to the activities of the previous three COAs, COA 4 includes the active participation of the Iraqi religious and political groups in the governance activities. It results in the highest probabilities of achieving the desired effects. While there is still some likelihood of local and regional support for the insurgents (0.22 and 0.14, respectively), many of the IED attacks are suppressed as are the insurgent fires. The result is significant increases in overt economic activity and moderate increase in the covert economic activity. Public services are still only moderately adequate, with room for improvement.

While the static quantitative analysis provides a lot of insight into the potential results of various COAs, it does not address the questions of how long it will take for the results to unfold or what should the timing of the actions be. The dynamic temporal analysis can provide answers to these types of questions.

Having created the TIN model with the time delay information, it is possible to experiment with various COAs and input scenarios. Figure 5.26 shows an example of COA and input scenarios that illustrate such an experiment. The second column of the Table in Fig. 6 shows a summary of the input nodes that were used in the experiment. They are divided into two types, those listed as Scenario and those listed as COA Actions. The scenario portion contains actions that may take place over which limited control is available. These set the context for the experiment. The second group contains the actions over which control exists, that is the selection of the actions and when to take them is a choice that can be made. The last column shows the scenario/action combinations that comprise the COA/Scenario to be examined. The column provides a list of ordered pairs for each Scenario Action or COA Action. Each pair provides a probability (of the action) and a time when that action starts. For example, the listing for the second scenario actions is [0.5, 0] [1.0, 1] which means that the probability of Country M and Country L interfering is 0.5 at the start of the scenario and changes to 1.0 at time = 1. In this analysis, time is measured in weeks.

The entries under the column labeled “COA 4a” mean that the scenario/under which the COA being tested is one in which there is immediate and full support for the insurgency (financial, material, and personnel) from international sources, and it is expected to exist throughout the scenario. The same is true for support from Country S. Countries M and L are modeled with the probability of providing support at 0.5 initially, but it immediately increases to 1.0 at week 1. All of the COA actions are assumed to not have occurred at the start of the scenario, thus the first entry of each is [0, 0]. The coalition force (Blue) actions start at week 1 with a probability of 1.0, meaning that all of the elements of Blue actions start at the beginning. With regard to religious activities, the Kurds begin at week 1 with probability 1.0. The Shia and Sunni have a probability of 0.5 starting at week 10 and then increase to 1.0, becoming fully engaged at week 20. In terms of political activity, the Kurds and Shia become fully active at week 1. The Shia become more likely to be active at week 10, fully active at week 20, then become less likely to

be active at week 30 (probability 0.5) and then become fully active again at week 40. Finally, the External Financial support begins at week 26.

	Action	COA 4a: List [p, t]
Scenario	Int'l Support to Insurgents	[1.0, 0]
Actions	Interference by countries M and L	[0.5, 0], [1.0, 1]
	Interference by country S	[1.0, 0]
COA	Blue TTPs activated	[0, 0], [1.0, 1]
Actions	Blue Surveillance, IO, Training, Brokering	[0, 0], [1.0, 1]
	Shia and Sunni Religious Activity	[0, 0], [0.5, 10], [1.0, 20]
	Kurd Religious Activity	[0, 0], [1.0, 1]
	Kurd and Shia Political Activity	[0, 0], [1.0, 1]
	Sunni Political Activity	[0, 0], [1.0, 20], [0.5, 30], [1.0, 40]
	International Investment	[0, 0], [1.0, 26]

Fig. 5.26 Dynamic Temporal Analysis Input

To see what the effect of this input scenario on several key effects, the model is executed and the probabilities of the key effects as a function of time are plotted as shown in Fig. 5.27. In the figure, the probability profiles of four effects are shown: IEDs are suppressed on Routes A and B and Local and Regional support for the Insurgents exists.

Figure 5.27 shows that the probability of suppression of the IED attacks on the two routes increases significantly under this scenario. This means that the number of IED attacks should decrease, more on Route A than on Route B. The improvement can be expected to occur more rapidly along Route A than along Route B by about 35 weeks or 8 months. Route A is the northern route that is controlled by the Kurds and Route B is the southern route controlled by the Shia and Sunni. This can be attributed to the rapid and steadfast political and religious activities of the Kurds as opposed to the more erratic activities of the others as modeled in the input scenario (Fig. 5.26). Also note that it is expected to take 80 to 100 weeks (nearly 2 years) for the full effect to occur. Fig. 5.27 also shows a significant decline in support for the insurgents both by the local and the regional populace with the local support decreasing more as the situation with respect to governance and services improves.

Of course it is possible to examine the behavior of any of the nodes in the model, by plotting their probability profiles. This can increase the understanding of the complex interactions and dependencies that in the situation that have been expressed in the TIN model. The TIN model provides a mechanism to experiment with many different scenarios and COAs. Questions like what will happen if some of the Blue CF actions are delayed or what will happen if the Shia or Sunni decide not to participate after some period of time can be explored. By creating plots of the probability profile of key effects under different scenarios, it is possible to explore the differences in expected outcomes under different scenarios. This can be illustrated by changing the input scenario. Suppose that it is believed to be possible to get other countries or external organizations to reduce the support to the insurgents by some means, for example diplomatic or military action. It is postulated that we could reduce the likelihood of such support to about 50% but it will take 6 months to do this. The results can be modeled by changing the input scenario of Fig. 5.26. In this case the first line of Fig. 5.26 is changed from [1.0, 0] to [1.0, 0] [0.5, 26]. All of the other inputs remain the same. Figure 5.28 shows a comparison of effect of this change on

the suppression on IED attacks along Route B. The reduction in international support for the insurgents at week 26 can cause a significant improvement in the suppression of the IED attacks along Route B (and a corresponding improvement along Route A, not shown). The improvement begins about 6 months after the reduction in international support or about 1 year into the scenario. Thus, decision makers may wish to pursue this option.




Fig. 5.27 Probability Profiles of Scenario (COA) of Fig. 5.26

5.2.5 Observations and Conclusions

Creating TIN models of situations provides a representation of knowledge about a situation that is derived from an understanding of the capabilities of an adversary and the interactions and dependencies of that adversary with the local and regional social, religious, and economic condition. Once created, the TIN model can be used to conduct computational experiments with different scenarios and COAs. In a sense, it provides a mechanism to assess various COAs based upon comparisons of the change in the probability of key effects over time

It is important to emphasize that the purpose of these models is to assist analysts in understanding the potential interactions that can take place in a region based on actions taken by one or perhaps many parties. It is not appropriate to say that these models are predictive. They are more like weather forecasts, which help us to make decisions, but are rarely 100% accurate and are sometimes wrong. To help deal with this uncertainty, weather forecasts are continually updated and changed as new data become available from the many sensors that make a variety of observations in many locations. Since these models cannot be validated formally, the appropriate concept is that of credibility. Credibility is a measure of trust in the model that is developed over

time through successive use and comparison of the insights developed through the model and the occurrence of actual events and resulting effects.




Figure 5.28: Comparison of the Effect of Different Scenarios

The techniques described in this paper can make an important contribution to a variety of communities that need to evaluate complex situations to help make decisions about actions they may take to achieve effects and avoid undesired consequences. The approach offers at least three levels of analysis, a qualitative evaluation of the situation based on the graph that shows the cause and effect relationships that may exist in the environment, and two levels of quantitative evaluation. The first level of quantitative analysis is static, and shows, in a coarse way, what the likelihood of different effects occurring is given different sets of actions. The second quantitative level is dynamic and shows how the scenario may play out over time. The relevant aspect is that the approach allows the inclusion of diplomatic, information, military, and economic (DIME) instruments and highlights their cumulative effects.

The models can be used to illustrate areas of risk including undesired effects, and risks associated with the amount of time it will take to achieve desired effects. It should also be noted that these models are not likely to be created on a one time basis. It can be expected that the understanding of the situation will continue to evolve requiring updates or even new models to be created. Perhaps the best contribution is that the technique offers a standardized way to analyze and describe very complex situations.

During the ten years that such models have been applied to different domains and problems, a number of lessons have been noted.

The first lesson is that these models are best suited to addressing issues at the operational/strategic level and are unsuitable for the tactical level. At the tactical level, we need to expand

the range of attrition-type combat models to include the influences of the whole spectrum of instruments of national power. A very difficult issue is the determination of the interactions among the various instruments. For example, what is the effect of a diplomatic initiative when coupled with information operations and should the latter precede, be concurrent or follow the former?

The second lesson is that consideration of temporal issues is critical to the understanding of effects based operations applied to transnational terrorist networks. While the results of conventional military operations focused on attrition may be well understood, it is very difficult (not enough data yet) to estimate how some of the non military actions will affect the future recruitment by the terrorist organization. Even issues such as persistence are not well understood and, certainly, not quantified yet.

The third lesson is a critical one. It is much too early to establish general purpose TIN models that can be applied to different circumstances by changing the contained data. It is not even clear that this is a desirable approach or one that is technically sound for this class of problems. Rather, the way the technology and the tools are developing is to provide the analysts the capability to put together models (in a given domain about which the analyst is knowledgeable and for which SMEs are available) to address specific issues in the order of several hours. This approach has been tried successfully at the Global War games at the Naval War College in 2000 and 2001. At this time, the state of the art has taken two directions: (a) the development of template TINs for routine analyses and (b) the extraction directly from unstructured data using ontologies draft TINs that the analyst or modeler can then improve.

5.3 Enhanced Influence Nets Case Study

Abbas K. Zaidi, Faisal Mansoor, Titsa P. Papantoni-Kazakos, Alexander H. Levis

5.3.1 Introduction

In this section, we apply the algorithms developed in Chapter 3 to an illustrative TIN. We also provide a comparison of the latter results with those previously obtained via the use of the CAST logic. The model used in this section was presented by Wagenhals et al. in 2001 [18] to address the following scenario: As described in [18], internal political instabilities in Indonesia have deteriorated and ethnic tensions between the multiple groups that comprise Indonesia have increased. Religion has been a major factor in these conflicts. Members of one of the minority (2%) religious groups have banded together to combat disenfranchisement. These members have formed a rebel militia group. Armed conflicts recently occurred between those rebels and the Indonesian military. The rebels fled to eastern Java where they have secured an enclave of land. This has resulted in a large number of Indonesian citizens being within the rebel-secured territory. Many of these people are unsympathetic to the rebels and are considered to be at risk. It is feared that they may be used as hostages if ongoing negotiations break down with the Indonesian government. The food and water supply and sanitation facilities are very limited within the rebel-secured territory.

Several humanitarian assistance (HA) organizations are on the island, having been involved with food distribution and the delivery of public health services to the urban poor for several years. So far, the rebels have not prevented HA personnel from entering the territory to take supplies to the citizens. The U.S. and Australian embassies in Jakarta are closely monitoring the situation for any indications of increasing rebel activity. In addition, Thailand, which has sent several hundred citizens to staff numerous capital investment projects on Java, is known to be closely monitoring the situation.

5.3.2 Modeling

To reflect the situation stated above, a TIN was first created in [18] and is shown in Fig. 5.29. This TIN models the causal and influencing relationships between (external) affecting events (on the left side and along the top of the model in Fig. 5.29) and the overall effect of concern which is the single node with no parents on the right-hand side of the model. In this case, the effect is “Rebels decide to avoid violence”. The actionable (external) events in this model include a combination of potential coalition, UN, and rebel actions. The coalition actions include actions by the US government, its military instrument of national power, actions by the Government of Indonesia, and actions by Thailand.

For purposes of illustration and comparison of results, we have selected a part of this network, as shown in Fig. 5.30.

The (external) affecting events in the TIN of Fig. 5.30 are drawn as root nodes (nodes without incoming edges). The text in each node, e.g., “1—Coalition Deploys Forces to Indonesia,” represents a node ID and a statement describing the binary proposition. In Fig. 5.30, $\{A_i\}_{0 \leq i \leq 4}$ represents the set of the external affecting events, where the index ‘i’ depicts the node ID. The marginal probabilities for the external affecting events are also shown inside each node. In this illustration, we assume all external affecting events to be mutually independent (Section 3.4.)

the influence of a parent node's state (as being either true or false) on its child node. The third element in the inscription depicts the time it takes for a parent node to influence a child node. For instance, in Fig. 5.30, event “1—Coalition Deploys Forces to Indonesia” influences the occurrence of event “7—Coalition Secures APOD and SPOD” after 3 time units.

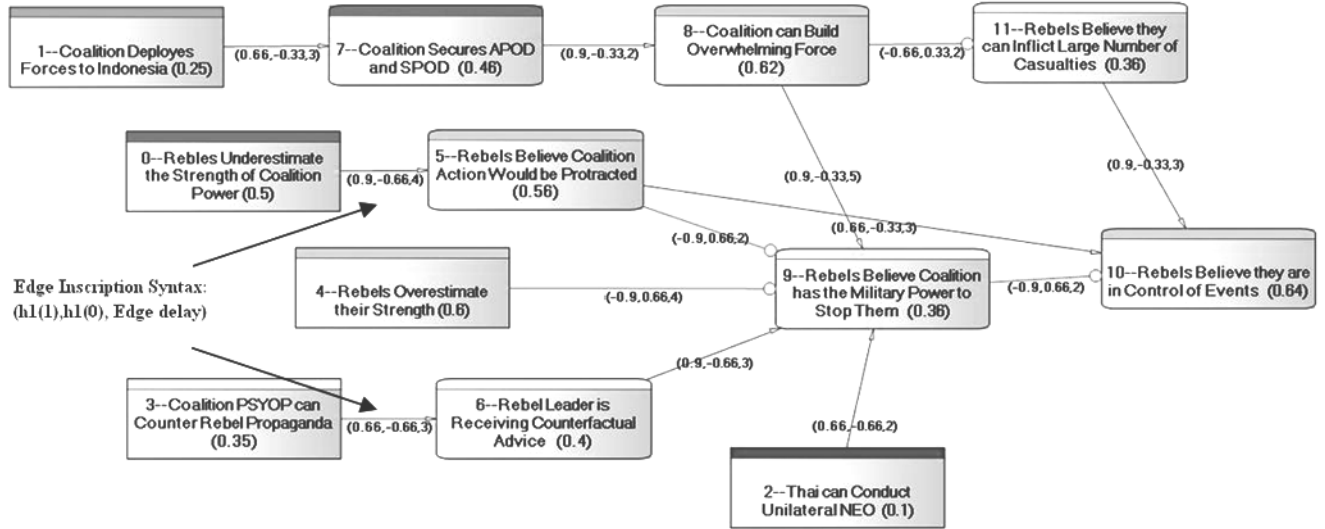


Fig. 5.30 Sample TIN for Analysis

The purpose of building a TIN is to evaluate and compare the performances of alternative courses of actions described by the set A_T in the definition of TINs. The impact of a selected course of action on the desired effect is analyzed with the help of a *probability profile*. The following is an illustration of such an analysis with the help of two COAs, given below:

COA1: All external affecting events are taken simultaneously at time 1 and are mutually independent.

COA2: Events {0, 2, 4} are taken at time 1, simultaneously, and events {1, 3} are taken at time 2, simultaneously.

The two COAs can also be described as in Table 5.5.

TABLE 5.5 The two Courses of Action

Event	COA1		COA2	
	Time	Status	Time	Status
0 -- Rebels Underestimate the Strength of Coalition Power	1	1 (= True)	1	1
1 -- Coalition Deploys Forces to Indonesia	1	1	2	1
2 -- Thai can Conduct Unilateral NEO	1	1	1	1
3 -- Coalition PSYOP can Counter Rebel Propaganda	1	1	2	1
4 -- Rebels Overestimate their Strength	1	1	1	1

Note that the simultaneous occurrence of external affecting events does not necessarily imply simultaneous revealing of their status on an affected node; the time sequence of revealed affecting events is determined by both the time stamp on each affecting event and the delays on edges. Because of the propagation delay associated with each edge, influences of actions impact the affected event progressively in time. As a result, the probability of the affected event changes as time evolves. A probability profile draws these probabilities against the corresponding time line. In Fig. 5.31, probability profiles generated for nodes “9—Rebels Believe Coalition has the Military Power to Stop Them” and “10—Rebels Believe they are in Control of Events,” using the CAST logic based approach in [4, 5, 13, 15, 17] are shown.




Fig. 5.31 Probability Profiles Generated by the CAST Logic Approach

For the same TIN model as in Fig. 5.30 and the corresponding course of actions, we used the approach presented in this paper and produced pertinent results for the following two cases:

Case I

For this illustration, we utilize the influence constant model presented in section 3.8. A and the temporal case presented in section 3.6. The influence constants $\{h_i(x_1^n)\}_{1 \leq i \leq n-1}$ are first pre-computed via the dynamic programming expression in Lemma 3.2, section 3.3. The resulting probability profiles for the two affected events/propositions in the TIN are shown in Fig. 5.32.




Fig. 5.32 Probability Profiles for Case I

Case II

For this illustration, we utilize the influence constant model presented in section 3.8. A and the temporal case where the existence of an affecting event is assumed unknown to an affected event unless it reveals itself and makes its status known to the affected event. The conditional probabilities, in this case, are computed real-time by eq. (3.33). The resulting probability profiles for the two affected events/propositions in the TIN are shown in Fig. 5.33.

Comparing Figs. 5.31 and 5.32, we note that when the existence of all the external affecting events are initially known, then the approach in this paper produces results that are more accurate and consistent than those produced by the CAST logic based approach. This was expected, since the present approach has eliminated the inconsistencies that the CAST logic based approach suffers from. Unlike the CAST logic based approach, the probability profiles generated by the new approach only record the posterior probabilities resulting from the impacts of the external affecting events and do not assume any default initial values; in profiles of Figs. 5.31 and 5.32 the first impact is recorded at time '3'. Comparing Figs. 5.32 and 5.33, we note that, as expected, when the existence of the external affecting events are revealed sequentially in time then, there is a relatively high level of instability in time evolution, as compared to the case where the existence of all the external affecting events is initially known. The selection of a influence constant and of temporal models for a TIN under construction/analysis is a design issue and is reflected by the differences in the resulting probability profiles.

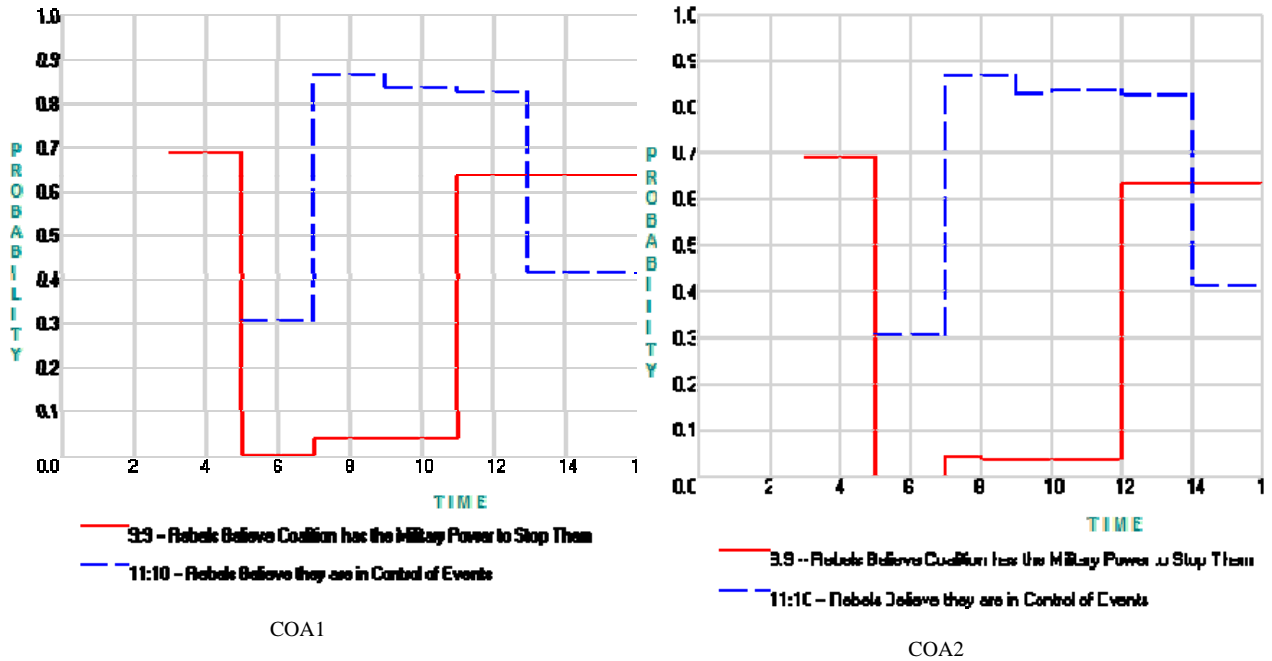


Fig. 5.33. Probability Profiles for Case II

5.4 Conclusion

In this chapter, several applications were presented. The first one dealt with the analysis of the data regarding the placement and effect of IEDs in the Diyala province of Iraq. The second case, illustrated the application of Timed Influence nets to the development and evaluation of potential Courses of Action for suppressing IEDs in Diyala. The third case illustrated how the expanded theory of Influence Nets relaxes the assumptions regarding causality while producing the same results with the original model when the restrictive assumptions apply.

PART III: MODELS OF ORGANIZATIONS

Chapter 6: Computationally Derived Models of Adversary Organizations

Chapter 7: Extracting Adversarial Relationships from Texts

Chapter 8: Inferring and Assessing Informal Organizational Structures from an Observed Dynamic Network of an Organization

Chapter 9: Simulating the Adversary: Agent-based Dynamic Network Modeling

Chapter 10: Adversary Modeling – Applications of Dynamic Network Analysis

Chapter 6

Computationally Derived Models of Adversary Organizations

Ashraf M. AbuSharekh, Smriti K. Kansal, A. Erkin Olmez, and Alexander H. Levis

6.1 Introduction

The effort to model organizational behavior with mathematical models has a long history. The groundbreaking work of Marshak & Radner [43] looked at the communications between organization members; today we would call this connectivity and associated information flows. Drenick [44] proposed a mathematical theory of organization in which a number of fundamental system theoretic ideas were exploited to draw insights for the design of organizations consisting of members who process tasks under time constraints – a form of Simon’s [45] bounded rationality. Levis [46] and his students developed a discrete event dynamical model and a set of rules that governed the allowed interactions – whether they represented forms of information sharing or of commands. This model, expressed mathematically in the language of Colored Petri Nets [47], allowed the design of organizational architectures that could meet accuracy and timeliness constraints while not exceeding the workload limitations of the decision makers. Essentially, the organization members conducted information processing and decision making tasks, often supported by decision support systems in order to reduce workload, while increasing accuracy and timeliness of the organizational response [48].

The basic model of the single decision maker evolved over time in order to accommodate more complex interactions and allow for different types of internal processing by the organization members [49]. The early focus was on small teams in which several members needed to be organized to perform a demanding, time-sensitive task. The objective was to achieve organizational performance without causing excessive workload that would lead to performance degradation.

A key objective, relating structure to behavior, meant that the structure and attributes of the simulation models must be traceable, in a formal way, to the architecture design. Hence the use of the term “executable” model which denotes that there is a formal mathematical model used for simulation with characteristics that are traceable to the static designs. The mathematical model can also be used for analysis, i.e., properties of the model and performance characteristics can be determined from the mathematical description. A wealth of theoretical results on discrete event dynamical systems, in general, and Colored Petri nets, in particular, can be applied to the executable model.

More recently, the problem of modeling adversary organizations about which we may have limited information has received renewed attention. Adversaries may have differences in equipment or materiel, differences in command structures, differences in constraints under which they can operate, and, last but not least, differences in culture. The differences in equipment and in operational constraints can be handled easily in the existing modeling framework. Differences in command structures require some additional work to express these differences in structural and quantitative ways. The real challenge is how to express cultural differences in these, primarily mechanistic, models of organizations.

Other considerations that drive the design problem are the tempo of operations and whether the adversary has an explicit organization, as a military force would have, or an implicit one, as a loosely coupled terrorist organization may have. This work focuses on the ability to introduce attributes that characterize cultural differences into the mechanistic model for organization design and use simulation to see whether these parameters result in significant changes in structure. The objective, therefore, is to relate performance to structural features but add attributes that characterize cultural differences. Specifically, the attributes or dimensions defined by Hofstede [50] are introduced in the design process in the form of constraints on the allowable interactions within the organization.

In sections 6.2 and 6.3, the modeling approach is described briefly since it has been documented extensively in the literature. In sections 6.4 and 6.5, the Hofstede dimensions are introduced and then applied to the organization design algorithm. In sections 6.6 and 6.7, two illustrative examples are presented – one focuses on the design of adversary organizations and one on coalition organizations. In the final section, 6.8, advantages and shortcomings of this approach are discussed.

6.2 The Decision Maker Model And Organizational Design

The five-stage interacting decision maker model [49] had its roots in the investigation of tactical decision making in a distributed environment with efforts to understand cognitive workload, task allocation, and decision making. The five-stage model allows the algorithm in each stage to be defined and makes explicit the input and output interactions of the decision maker with other organization members or the external environment. It also has a well-defined algorithm for characterizing workload. This model has been used for fixed as well as variable structure organizations [51].

The five-stage decision maker (DM) model is shown in Fig. 6.1. The DM receives signals from the external environment or from another decision maker. The Situation Assessment (SA) stage represents the processing of the incoming signal to obtain the assessed situation that may be shared with other DMs. The decision maker can also receive situation assessment signals from other decision makers within the organization; these signals are then fused together in the Information Fusion (IF) stage to produce the fused situation assessment. The fused information is then processed at the Task Processing (TP) stage to produce a signal that contains the task information necessary to select a response. Command input from superiors is also received. The Command Interpretation (CI) stage then combines internal and external guidance to produce the input to the Response Selection (RS) stage. The RS stage then produces the output to the environment or to other organization members.




Fig. 6.1 Model of the Five-Stage Decision Maker

The key feature of the model is the explicit depiction of the interactions with other organization members and the environment. These interactions follow a set of rules designed to avoid

deadlock in the information flow. A decision maker can receive inputs from the external environment only at the SA stage. However, this input can also be another decision maker's output. A decision maker can share his assessed input with another organization member; this is depicted as an input to the IF stage when the decision maker is receiving a second input. This input must be generated from another decision maker and can be the output of the SA or RS stage. In the CI stage, the decision maker can receive commands. This is also internally generated and must originate from another decision maker's RS stage. Thus the interactions between two decision makers are limited by the constraints enumerated above: the output from the SA stage, can only be an internal input to another decision maker's IF stage, and an internal output from the RS stage can only be input to another decision maker's SA stage, IF stage, or CI stage.

The mathematical representation of the interactions between DMs is based on the connector labels e_i , s_i , F_{ij} , G_{ij} , H_{ij} and C_{ij} of Fig. 6.2; they are integer variables taking values in $\{0, 1\}$ where 1 indicates that the corresponding directed link is actually present in the organization, while 0 reflects the absence of the link. These variables can be aggregated into two vectors e and s , and four matrices F , G , H and C . The interaction structure of an n -decision-maker organization may be represented by the following six arrays: two $n \times 1$ vectors e and s , representing the interactions between the external environment and the organization:

$$e = [e_i], \quad s = [s_i] \quad \text{for } i = 1, 2, \dots, n$$

and four $n \times n$ matrices F , G , H and C representing the interactions between decision makers inside the organization. Since there are four possible links between any two different DMs, the maximum number of interconnecting links that an n decision-maker organization can have is

$$k_{\max} = 4n^2 - 2n$$

Consequently, if no other considerations were taken into account, there could be $2^{k_{\max}}$ alternative organizational forms. This is a very large number: 290 for a five-person organization.




Fig. 6.2 One-sided Interactions Between Decision Maker i and Decision Maker j

In the Petri net representation of the DM model, the transitions stand for the algorithms, the connectors for the precedence relations between these algorithms, and tokens for the messages that flow between the DMs. If the tokens need to be distinct, i.e., carry information, then a Colored Petri net representation is used. Other organization components can be modeled using the same basic five-stage model, but eliminating one or more of the stages. For example, a processor that receives sensor data and converts it to an estimate of a vector variable can be modeled by a single SA transition, while a data fusion algorithm can be modeled by an IF transition. With this

model of the organization member and its variants used to model other components, it is now possible to formulate the problem of designing decision-making organizations.

6.3. The Lattice Algorithm

The analytical description of the possible interactions between organization members forms the basis for an algorithm that generates all the architectures that meet some structural constraints as well as application-specific constraints that may be present. The set of structural constraints rules out a large number of architectures. The most important constraint addresses the connectivity of the organization - it eliminates information structures that do not represent a single integrated organization. Remy and Levis [52] developed an algorithm, named the Lattice algorithm, that determines the maximal and minimal elements of the set of designs that satisfy all the constraints; the entire set can then be generated from its boundaries. The algorithm is based on the notion of a simple path - a directed path without loops from the source to the sink. Feasible architectures are obtained as unions of simple paths. Consequently, they constitute a partially ordered set. The algorithm receives as input the matrix tuple $\{\mathbf{e}, \mathbf{s}, \mathbf{F}, \mathbf{G}, \mathbf{H}, \mathbf{C}\}$ of dimension n , where n is the number of organization members.

There are some structures corresponding to combinations of interactions between components that do not have a physical interpretation; e.g., DMs can exchange information - \mathbf{F}_{ij} and \mathbf{F}_{ji} can coexist - but commands are unilateral- either \mathbf{C}_{ij} or \mathbf{C}_{ji} or none, but not both. Those structures should be eliminated, if realistic organizational forms are to be generated. The structural constraints define what kinds of combinations of interactions need to be ruled out. A set of four different structural constraints is formulated that applies to all organizational structures being considered.

- R1 A directed path should exist from the source to every node of the structure and from every node to the sink.
- R2 The structure should have no loops; i.e., the organizational structures should be acyclical.
- R3 There can be at most one link from the RS stage of a DM to each one of the other DMs; i.e., for each i and j , only one element of the triplet $\{\mathbf{G}_{ij}, \mathbf{H}_{ij}, \mathbf{C}_{ij}\}$ can be nonzero.
- R4 Information fusion can take place only at the IF and CI stages. Consequently, the SA and RS stages of each DM can have only one input.

Constraint R1 eliminates structures that do not represent a single integrated organization and ensures that the flow of information is continuous within an organization. Constraint R2, allows acyclical organizations only¹. Constraint R3 states that the output of the RS stage of one DM or component can be transmitted to another DM or component only once: it does not make much sense to send the same information to the same decision maker at several different stages. Constraint R4 prevents a decision maker from receiving more than one input at the SA stage. The rationale behind this limitation is that information cannot be merged at the SA stage; the IF stage has been specifically introduced to perform such a fusion.

¹ This restriction is made to avoid deadlock and circulation of messages within the organization.

Any realistic design procedure should allow the designer to introduce specific structural characteristics appropriate to the particular design problem. To introduce user-defined constraints that will reflect the specific application the organization designer is considering, appropriate 0s and 1s can be placed in the arrays $\{e, s, F, G, H, C\}$. The other elements will remain unspecified and will constitute the degrees of freedom of the design. The complete set of constraints is denoted by \mathbf{R} .

A feasible structure is one that satisfies both the structural and the user-defined constraints. The design problem is to determine the set of all feasible structures corresponding to a specific set of constraints. Note that this approach is not, by design, concerned with the optimal organizational structure, but with the design of a whole family of feasible structures. At this stage, we are only concerned with the structure and information flows, i.e., the development of the set of feasible organizational forms. This set will become the admissible set in the problem of incorporating cultural constraints.

The notion of subnet defines an order (denoted $<$) on the set of all well defined nets of dimension n . The concepts of maximal and minimal elements can therefore be defined. A maximal element of the set of all feasible structures is called a maximally connected organization (MAXO). Similarly, a minimal element is called a minimally connected organization (MINO). Maximally and minimally connected organizations can be interpreted as follows. A MAXO is a well defined net such that it is not possible to add a single link without violating the set of constraints \mathbf{R} . Similarly, a MINO is a well defined net such that it is not possible to remove a single link without violating the set of constraints \mathbf{R} . The following proposition is a direct consequence of the definition of maximal and minimal elements: For any given feasible structure P , there is at least one MINO P_{\min} and one MAXO P_{\max} such that $P_{\min} < P < P_{\max}$. Note that the net P need not be a feasible. There is indeed no guarantee that a well-defined net located between a MAXO and a MINO will fulfill the constraints \mathbf{R} , since such a net need not be connected. To address this problem, the concept of a simple path is used.

The following proposition characterizes the set of all feasible organizational structures: P is a feasible structure if and only if P is a union of simple paths, i.e., P is bounded by at least one MINO and one MAXO. Note that in this approach the incremental unit leading from a feasible structure to its immediate super-ordinate is a simple path and not an individual link. In generating organizational structures with simple paths, the connectivity constraint $R1$ is automatically satisfied.

The Lattice algorithm generates, once the set of constraints \mathbf{R} is specified, the MINOs and the MAXOs that characterize the set of all organizational structures that satisfy the designer's requirements. The next step of the analysis consists of putting the MINOs and the MAXOs in their actual context to give them a physical instantiation. If the organization designer is interested in a particular (MINO, MAXO) pair because it contains interactions that are deemed desirable for the specific application, he can further investigate the intermediate nets by considering the chain of nets that is obtained by adding simple paths to the MINO until the MAXO is reached.

This methodology provides the designer of organizational structures with a rational way to handle a problem whose combinatorial complexity is very large. Having developed a set of organizational structures that meets the set of logical constraints and is, by construction, free of structural problems, we can now address the problem of incorporating attributes that characterize cultures.

6.4. Modeling Cultural Attributes

Hofstede [50] distinguishes dimensions of culture that can be used as an instrument to make comparisons between cultures and to cluster cultures according to behavioral characteristics. Culture is not a characteristic of individuals; it encompasses a number of people who have been conditioned by the same education and life experience. Culture, whether it is based on nationality or group membership such as the military, is what the individual members of a group have in common [53].

To compare cultures, Hofstede originally differentiated them according to four dimensions: *uncertainty avoidance (UAI)*, *power distance (PDI)*, *masculinity-femininity (MAS)*, and *individualism-collectivism (IND)*. The dimensions were measured on an index scale from 0 to 100, although some countries may have a score below 0 or above 100 because they were measured after the original scale was defined in the 70's. The original data were from an extensive IBM database for which 116,000 questionnaires were used in 72 countries and in 20 languages over a six-year period. The hypothesis here is that these dimensions may affect the interconnections between decision makers working together in an organization.

The power distance dimension can be defined as "the extent to which less powerful members of a society accept and expect that power is distributed unequally" [50]. An organization with a high power distance value will likely have many levels in its hierarchy and convey decisions from the top of the command structure to personnel lower in the command structure; centralized decision making. Organizations with low power distance values are likely to have decentralized decision making characterized by a flatter organizational structure; personnel at all levels can make decisions when unexpected events occur with no time for additional input from above.

Uncertainty avoidance can be defined as "the extent to which people feel threatened by uncertainty and ambiguity and try to avoid these situations"[50]. An organization which scores high on uncertainty avoidance will have standardized and formal procedures; clearly defined rules are preferred to unstructured situations. In organizations with low scores on uncertainty avoidance, procedures will be less formal and plans will be continually reassessed for needed modifications. Klein et al. [54] hypothesized that during complex operations, it may not be possible to specify all possible contingencies in advance and to take into account all complicating factors.

The trade-off between time and accuracy can be used to study the affect of both power distance and uncertainty avoidance in the model [55]. Messages exchanged between decision makers can be classified according to three different message types: information, control, and command ones [56]. Information messages include inputs, outputs, and data; control messages are the enabling signals for the initiation of a subtask; and command messages affect the choice of subtask or of response. The messages exchanged between decision makers can be classified according to these different types and each message type can be associated with a subjective parameter. For example, uncertainty avoidance can be associated with control signals that are used to initiate subtasks according to a standard operating procedure. A decision maker with high uncertainty avoidance is likely to follow the procedure regardless of circumstances, while a decision maker with low uncertainty avoidance may be more innovative. Power distance can be associated with command signals. A command center with a high power distance value will respond promptly to a command signal, while in a command center with a low power distance value this signal may not always be acted on or be present.

6.5 Using Cultural Constraints

Cultural constraints help a designer determine classes of similar feasible organizations by setting specific conditions that limit the number of various types of interactions between decision makers. Cultural constraints are simply represented as interactional constraint statements. Four types of interactions have previously been defined (information sharing represented by matrix **F**, control represented by matrix **G**, result sharing represented by matrix **H**, and command represented by matrix **C**). The upper bounds, lower bounds and constants of an interactional constraint statement can take a value between 0 or the number of fixed-type interactions allowed by user-defined requirements (whichever is higher) and the maximum number of interactions allowed by user-defined requirements for a given problem, and are formulated using a group's cultural score. An approach for determining the values of these constraints has been developed by Olmez [57]. The constraints are obtained using a linear regression on the four dimensions to determine the change in the range of the number of each type of interaction that is allowed.

$$dY = c + \alpha(PDI) + \beta(UAI) + \gamma(MAS) + \delta(IND)$$

where Y is #F or #G or #H or #C

Example:

$$\#F \leq 2, \#G = 0, 1 \leq \#H \leq 3, \#C = 3$$

The methodology to obtain the solution space given a set of user-defined constraints and cultural constraints using an extended lattice algorithm called C-Lattice is presented next.

C-Lattice Algorithm: The Lattice Algorithm allows the automatic generation of candidate structures based on a set of user and structural constraints. If the cultural constraints can be included in the problem statement in a manner similar to the structural constraints, then the lattice structure of the solution space will be preserved and an extended version of the Lattice algorithm may be used to generate structures that satisfy the additional cultural attributes. Since the cultural constraints impose limits on the number of interactions between the decision makers, they are placing additional structural constraints on the solution space. Hence the constraints R1 to R4 specified in [52] can be extended to include the cultural constraints R5 to R8. For example, for the cultural constraint statement give earlier, they become:

- R5: The number of **F** type interactions must be between 0 and 2
- R6: The number of **G** type interactions must equal 0
- R7: The number of **H** type interactions must lie between 1 and 3
- R8: The number of **C** type interactions must equal 3.

The flowchart in Fig. 6.3 explains the generation of the culturally constrained solution space.

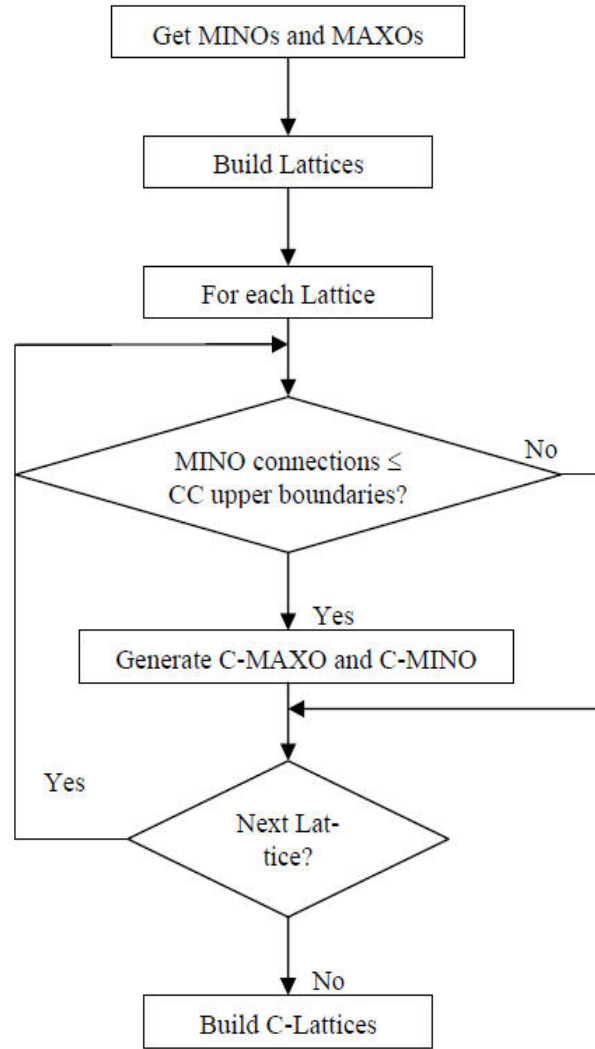


Fig. 6.3 Flowchart for culturally constrained solution space

MAXOs and MINOs are generated using the same algorithm described in [52]. The “Build Lattices” step checks if a MINO is contained within a MAXO. If it is, then the MINO is connected to that MAXO and forms a lattice. For each lattice in the solution space, we check the MINO to see if it violates the cultural boundaries. For example, if the number of **F** type interactions in the MINO is two and the maximum allowable by the cultural constraints is only one, then the MINO does not satisfy the cultural attributes and since the MINO is the minimally connected structure in that lattice, no other structure will satisfy the constraints. Hence the lattice can be discarded. If the MINO does pass the boundary test, then simple paths are added to it to satisfy the cultural constraints R5 to R8. The corresponding minimally connected organization(s) is now called the C-MINO(s) (culturally bound MINO). Similarly, by subtracting simple paths from the MAXO, C-MAXO(s) can be reached. The step “Build C-Lattices” connects the C-MINOs to the C-MAXOs. The advantage of using this approach is that the designer does not have to know the cultural attributes at the start of the analysis. He can add them at a later stage. This also enables him to study the same organization structure under different cultures. Also pre-

viously designed organization structures can now be analyzed in new light using cultural attributes.

6.6 Adversarial Modeling Using CAESAR III

The design approach and the algorithm are illustrated using a hypothetical example of an adversarial organization. The simulations were performed using a new application called CAESAR III developed in System Architectures Lab at GMU. CAESAR III is used for the design of information processing and decision making organizations at the operational and tactical levels; it takes into consideration cultural differences as required by the designer.

The scenario reads as follows: Intelligence from the field has informed Blue that the adversary (RedD) has organized a force to conduct operations in a distinct part (a province) of the Area of Responsibility. Intelligence has also indicated that the leadership consists of six persons with the command structure as shown in Fig. 6.4. The Field Intelligence Officers have different areas of responsibility.

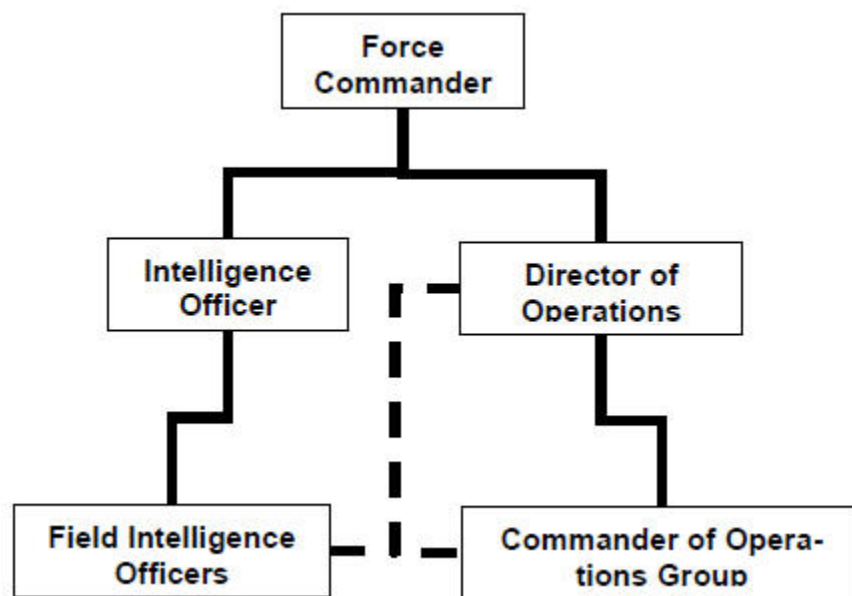


Fig. 6.4 Command Relationship Chart for Red

The cultural constraints for the two countries are also known.

TABLE 6.1 Cultural Constraints

	#F	#G	#H	#C
Blue	2	0	1-3	2-3
Red	2-4	0	1-5	2-4

Given the scenario and the cultural attributes of Red and Blue, can one infer the possible organizational structure of the Red Force and its information exchanges so that Blue can focus its ISR assets to the right targets?

Based on the command relationship chart, one can deduce the number of decision makers (six in this case) and also specify the interactions between them;

- The Field Intelligence Officers interact with the environment and send their Situation Assessment to the Intelligence Officer.
- The Intelligence Officer fuses this information and sends his Assessment to the Force Commander.
- Based on the information received, the Force Commander directs the Director of Operations to develop a Course of Action
- The Director of Operations in turn directs the Commander of Operations to develop a plan based on the COA and execute it.
- The variable links have been introduced into the problem based on the type of interactions that usually exist in command and control organizations. They may or may not exist in the Red group. Cultural attributes will be used to determine probable links.

This can be represented in block diagram form as shown in Fig. 6.5. This information can also be represented in matrices form as shown below where ‘1’ represents a fixed type interaction and ‘x’ represents a variable type interaction (Fig. 6.6).

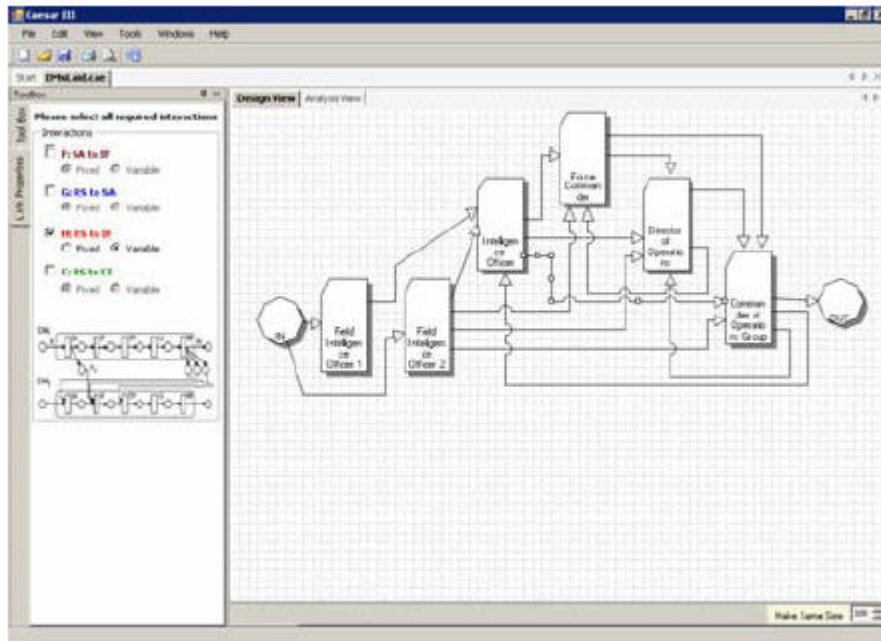


Fig. 6.5 Block Diagram of the Organization as seen in the CAESAR III GUI

The resulting universal net is shown in Fig. 6.7. Running the lattice algorithm without introducing the cultural attributes at this point helps design all feasible organizational structures that

meet the specific constraints of the problem. The resulting solution space has a single lattice bounded by one MINO and one MAXO. Figure 6.8 shows the partially expanded solution space.

Applying Red's cultural attributes to the solution space places further constraints on the number of allowable interactions and helps determine the (plausible) organizational structures that Red may be employing. The resulting solution consists of one MINO and 3 MAXOs and is shown in Fig. 6.9.

$$e = [1 \ 1 \ 0 \ 0 \ 0 \ 0] \quad s = [0 \ 0 \ 0 \ 0 \ 0 \ 1]$$

$$F = \begin{bmatrix} 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & x & x & x \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad G = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$H = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & x & x \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & x & 0 & 0 \\ 0 & 0 & x & 0 & x & 0 \end{bmatrix} \quad C = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & x \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

Fig. 6.6 Matrix representation of the design problem

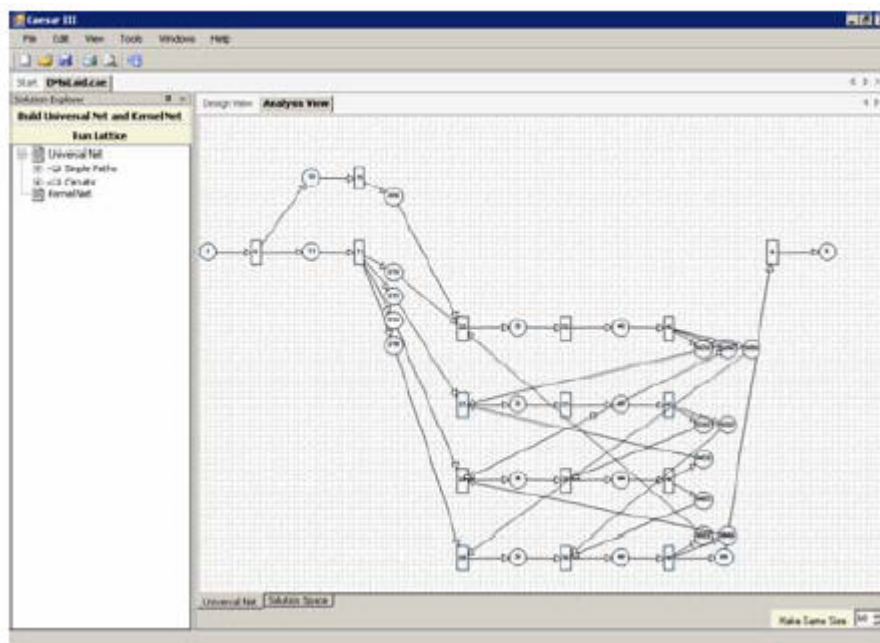


Fig. 6.7 Universal Net

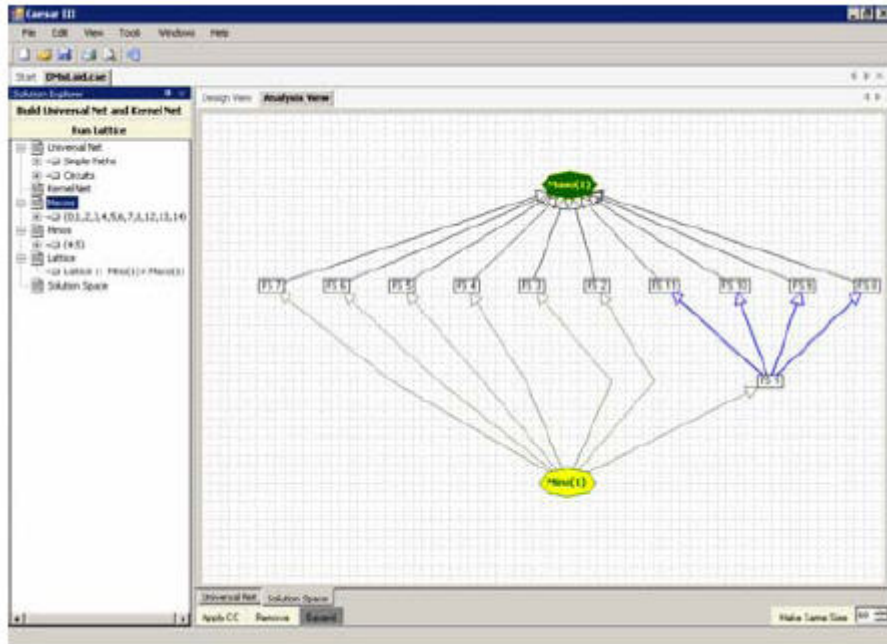


Fig. 6.8 Partially expanded solution space

The C-MAXOs and the C-MINOs lie within the MAXOs and the MINOs, i.e., the culturally bound solution space is contained in the un-constrained solution space.




Fig. 6.9 Culturally Constrained Solution Space for Red

An expanded lattice is shown in Fig. 6.10. All the structures that lie between a C-MINO and a C-MAXO satisfy the cultural constraints. The actual Petri nets corresponding to the CMINO and C-MAXOs are shown in Figs. 6.11 to 6.14.




Fig. 6.10 Expanded Lattice Structure from C-MINO(1) to CMAXO(1) for Red

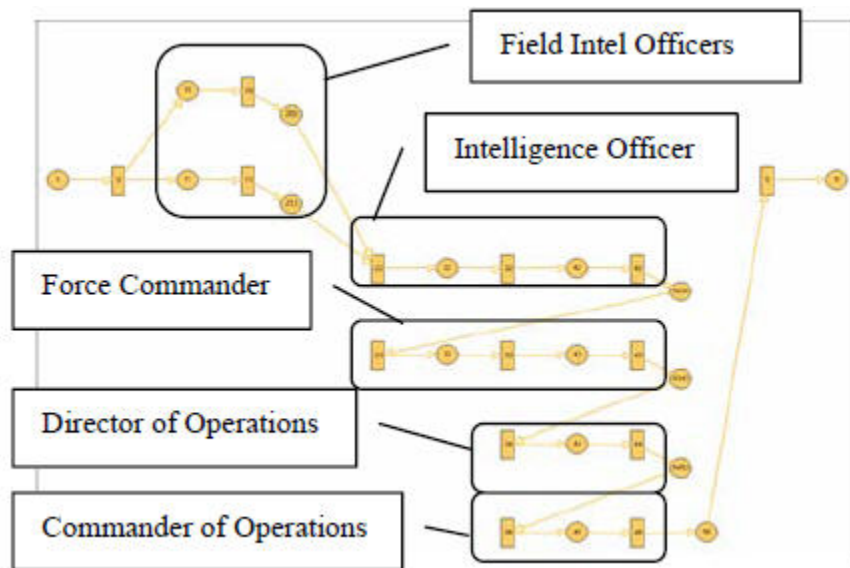


Fig. 6.11 C-MINO(1) for Red




Fig. 6.12 C-MAXO(1) for Red




Fig. 6.13 C-MAXO(2) for Red




Fig. 6.14 C-MAXO(3) for Red

Applying Blue's cultural attributes to the original problem results in only one C-MINO and one C-MAXO. The corresponding expanded lattice is as shown in Fig. 6.15.

The actual Petri net corresponding to the C-MAXO is shown in Figure 6.16. The C-MINO for Blue is the same as the C-MINO for Red.

Since the constrained solution space for Red has only one C-MINO, which is connected to all the three C-MAXOs, the C-MINO represents the set of interactions that must be present in all the structures that satisfy the cultural attributes of Red. Further analysis of this structure can help identify the high value ISR targets. In cases where there are more than one MINOs, identifying the interactions that are common to all the C-MINOs will indicate which areas to target for ISR activities.




Fig. 6.15 Expanded Lattice Structure from C-MINO(1) to CMAXO(1) for Blue




Fig. 6.16 C-MAXO(1) for Blue

Looking at the solution spaces for the two cases, it is easy to see that the cultural attributes do play a role in the final structure of the decision-making organizations and can provide valuable insight into possible structures that may be used by an adversary.

6.7 Coalition Modeling Using CAESAR III

The computational approach for the design of adversary organizations can also be applied to coalition operations. This is illustrated using a hypothetical example in which an emergency situation in an island nation requires rapid humanitarian assistance and disaster relief as well as securing military assets. The alternative architecture designs and the associated simulations to evaluate performance were carried out using CAESAR III.

The scenario depicts a situation in which anarchy has risen on an island due to a recent earthquake that caused substantial damage. The infrastructure and many of the government buildings are destroyed in the island's capital. The US maintains a ground station that receives data from space assets. It is concerned about the rising tensions, as there has been opposition to its presence on the island. As a result, US decides to send an Expeditionary Strike Group (ESG) to the island to: (1) provide timely Humanitarian Aid/ Disaster Relief (HA/DR) to three sectors of the island; and (2) counteract the effects of any hostile attacks which impede the normal operation of the HA/DR mission and the security of the ground station. As the ESG is away for the first critical day of the operation, countries A and B offer help to support the mission and agree to take part in a Coalition Force that would be commanded remotely by the US ESG commander. It is assumed that, close to the island, both countries hold different elements for an ESG compatible Coalition Force, which can be deployed in a matter of hours, while the ESG rushes to the island.

A team of five decision-making units carries out the HA/DR mission. The team is organized in the divisional structure and each unit under the team has its sub-organizations and staff to perform the tasks allocated to it. The five units are:

- (1) ESGC: Commander;
- (2) MEUC-Commander of the Marine Expeditionary Unit;
- (3) ACE-Air Combat Element with its Commander and sub-organizations;
- (4) GCE-Ground Combat Element with its Commander and sub-organizations; and
- (5) CSSE-Combat Service Support Element with its Commander and sub-organizations.

It is assumed that country A can provide support as ACE, GCE and CSSE while country B can only provide support as GCE and CSSE. The roles of ESGC and MEUC remain with the US. The countries are able to provide rapid assistance in coordination with each other and the design question becomes the allocation of different tasks to partners in this ad-hoc coalition.

This is a multi-level design problem in which interactions between different decision making units need to be determined both at the higher level (Level-1) as well as at the lower level (Level-2). The top level interactions correspond to interactions between culturally homogenous subunits, while the bottom level design problem consists of designing the internal structure of these homogenous subunits based on a defined set of interactional constraints and culture. Based on the structure of the ESG, one can impose user constraints to design the level-1 organization.

Figure 6.17 shows the block diagram of this organization as designed in CAESAR III; the matrices describing the interactions are shown in Fig. 6.18.




Fig. 6.17 Level-1 organizational block diagram.

$$e = [x \ 1 \ 1 \ 1 \ 0] s = [0 \ 0 \ 1 \ 1 \ 1]$$

$$F = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & x & 0 & x \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad G = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$H = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad C = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

Fig. 6.18 Matrix Representation corresponding to Fig. 6.17

Figure 6.19 shows the result of running the lattice algorithm on level-1 organization. The solution space contains one MINO, Fig. 6.20, and one MAXO, Fig. 6.21. The designer can pick a structure from this space and use it to design the sub-organizations at level-2.




Fig. 6.19. Solution space for Level-1 organization design as seen in CAESAR III




Fig. 6.20 MINO of Level-1 design


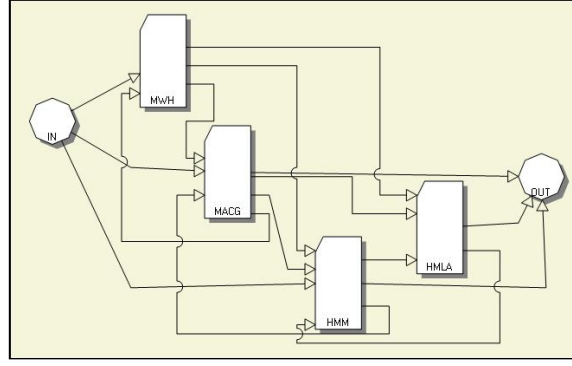


Fig. 6.21 MAXO of Level-1 design

Level-1 design is free of cultural constraints. However Level-2 design uses the C-Lattice algorithm to include cultural attributes to form the various coalition options. The sub-organizations of ACE, GCE and CSSE are designed using CAESAR III. Figures 6.22, 6.23 and 6.24 show the respective block diagrams along with the matrices specifying the user constraints. Since the US always performs the roles of ESGC and MEUC, these sub-organizations are not decomposed further.




$$e = [x \ 1 \ 1 \ 0] s = [0 \ x \ 1 \ 1]$$

$$F = \begin{bmatrix} 0 & 0 & 0 & 0 \\ x & 0 & 0 & x \\ 0 & x & 0 & x \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad G = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$H = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & x & 0 & x \\ 0 & 0 & x & 0 \end{bmatrix} \quad C = \begin{bmatrix} 0 & 1 & x & x \\ 0 & 0 & 1 & x \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

Fig. 6.22 Block diagram and matrix representation for ACE




$$e = [1 \ x \ x \ x] s = [0 \ 1 \ 1 \ 1]$$

$$F = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad G = \begin{bmatrix} 0 & x & x & x \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$H = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & x & x \\ 0 & 0 & 0 & x \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad C = \begin{bmatrix} 0 & x & x & x \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

Fig. 6.23 Block diagram and matrix representation for GCE



$$e = [1 \ 1 \ 0 \ 0 \ 0] s = [0 \ 0 \ 1 \ 1 \ x]$$

$$F = \begin{bmatrix} 0 & x & 1 & 0 & 0 \\ x & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad G = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$H = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & x \\ 0 & 0 & 0 & 0 & x \\ 0 & 0 & x & x & 0 \end{bmatrix} \quad C = \begin{bmatrix} 0 & 1 & 0 & x & x \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

Fig. 6.24 Block diagram and matrix representation for CSSE

Table 6.2 gives the Hofstede's scores for US, Country A and Country B. Using a multiple linear regression model, these scores are converted into limits to be placed on allowable interactions based on culture. These are imposed as additional structural constraints on the solution space of the sub-organizations. The cultural constraints for the three sub-organizations are shown in tables 6.3, 6.4 and 6.5. Maximum indicates the limit placed on the number of interactions by user constraints.

TABLE 6.2 Hofstede's scores for the three countries

Country	PDI	IND	MAS	UAI
US	40	91	62	46
A	38	80	14	53
B	66	37	45	85

TABLE 6.3 Cultural Constraints corresponding to ACE

Country	#F	#G	#H	#C
Maximum	$0 \leq F \leq 4$	0	$0 \leq H \leq 3$	$2 \leq C \leq 5$
US	$3 \leq F \leq 4$	0	$2 \leq H \leq 3$	3
A	2	0	$2 \leq H \leq 3$	3
B	2	0	1	$4 \leq C \leq 5$

TABLE 6.4 Cultural Constraints corresponding to GCE

Country	#F	#G	#H	#C
Maximum	0	$0 \leq G \leq 3$	$0 \leq H \leq 3$	$0 \leq C \leq 3$
US	0	2	$2 \leq H \leq 3$	2
A	0	2	$2 \leq H \leq 3$	1
B	0	$2 \leq G \leq 3$	2	$2 \leq C \leq 3$

TABLE 6.5 Cultural Constraints corresponding to CSSE

Country	#F	#G	#H	#C
Maximum	$1 \leq F \leq 3$	0	$0 \leq H \leq 4$	$3 \leq C \leq 5$
US	$2 \leq F \leq 4$	0	$3 \leq H \leq 4$	3
A	2	0	$3 \leq H \leq 4$	3
B	2	0	2	$4 \leq C \leq 5$

Using the C-Lattice algorithm, the solution space for each sub-organization is computed for each culture and a suitable structure is selected by the user. These structures are then used to form the different coalition options and analyze the performance. In view of the limited space, the complete solution spaces are not shown here. Figures 6.25-6.27 show the structures selected by the user for each country for CSSE. A similar approach can be use to select different structures to be used for ACE and GCE.




Fig. 6.25 GCE structure selected for US




Fig. 6.26 GCE structure selected for Country A




Fig. 6.27 GCE structure selected for Country B

Once the structure is selected, CAESAR III has the functionality of exporting it as a Colored Petri net to *CPN Tools* where it can be simulated to analyze performance. For the given scenario, based on the availability of support from the two countries, eight coalition options are possible, excluding the homogeneous option of all US. The five sub-organizations are combined together using Level-1 MINO and the eight options were simulated to study performance in terms of tasks served. The following assumptions are made. Each process (transition) needs 50 units of processing time. Each additional incoming link increases this time by 50 units. The reasoning is that the additional input(s) will require more processing. Hence, structures that have more interactions will take more time to process the tasks, which will affect the overall performance. Figure 6.28 shows the results of this analysis for all combinations. The x-axis shows the percentage of tasks **un-served**.

Based on these results, US-US-US-B-A performs best. Most options with country B in the CSSE role perform badly. This is because country B needs a high number of command relationships and the structure of CSSE allows for this to occur, thereby increasing the processing delay. User constraints on GCE allow for very similar cultural constraints for all countries and hence changing the ordering in this role does not change the performance very much. Similar results were obtained when the coalition options were simulated using a Level-1 MAXO organization.




Fig. 6.28 Percent of tasks un-served for coalition options.

A previously developed methodology for the computational design of information processing and decision making organizations has been enhanced to include cultural constraints that affect the choice of organizational structures. While the Hofstede cultural dimensions have been used, other cultural metrics can be used to derive the cultural constraints. A simple example illustrates

the approach for designing coalition organizations and analysing their performance. The results indicate that culture does affect the structure and working of organizations thereby affecting the overall performance. This could aid in the allocation of different tasks to partners in an ad-hoc coalition.

6.8 Conclusion

A previously developed methodology for the computational design of information processing and decision making organizations has been enhanced to include cultural constraints that affect the choice of organizational structures. While the Hofstede cultural dimensions have been used, other cultural metrics can be used to derive the cultural constraints R5 to R8. Two examples illustrate the approach: one for adversary organizations and one for coalition organizations. The results indicate that culture does affect the structure and working of organizations thereby affecting the overall performance.

Chapter 7

Extracting Adversarial Relations from Texts

Kathleen M. Carley

7.1 Introduction

There is a need to identify texts and extract from them various information about the adversary, their interactions, activities, beliefs, resources and so on. We used in this project a rapid ethnographic assessment procedure that moved from data to model using a semi-automated text analysis process. Key data used was newspaper reports. Over the course of the project this procedure became increasingly automated and the ability to identify agents and their activities improved dramatically. Central to this process is the AutoMap tool. AutoMap is based on network text analysis and so converts texts to networks of relations. We found it useful to first extract the semantic network and then the meta-network composed of agents, resources, expertise, locations, activities, beliefs and organizations. We note that beliefs are the most difficult to extract.

Data mining is commonly used to identify and extract entities. Named Entity Recognition is used to classify items such as people or locations [142]. Machine Learning is used widely to aid in the classification. Aspects of the data that give clues as to classification category are word length, part-of-speech, and external sources such as gazetteers and ontologies. Some algorithms parse at the shallow level of words only, and other parse deeply with machine understanding of part of speech and sentence semantics. Spatiotemporal knowledge discovery techniques are described by Roddick and Lees [143], and location techniques by Battenfield et al. [144].

Reliance on a gazetteer may improve the computer's ability to recognize locations. Gazetteers differ in scope, coverage and balance, accuracy, and entry specificity. Choice of gazetteer influences match results. The gazetteer can supply additional background knowledge that is helpful in data analysis. Some researchers use existing gazetteers such as the National Geospatial Intelligence Agency gazetteer¹ or GeoNames,² while others generate them automatically [145] or derive them from Wikipedia [146]. Semantic technologies have been used to identify network data in texts before [147], [148]. Workshops such as the Data Mining WebKDD/SNAKDD 2007 [149] and conference presentations [150] have been devoted specifically to mining data for social network analysis.

¹ National Geospatial Intelligence Agency gazetteer for download at
<http://earthinfo.nga.mil/gns/html/>

² <http://www.geonames.org>

7.1.1 Network Text Analysis (NTA)

Network Text Analysis is a set of methodologies for converting texts to graphs based on the theory that language and knowledge can be modeled as networks of words and relations such that meaning is inherent in the structure of that network. NTA encodes links among words to construct a network of linkages. Specifically, this method analyzes the existence, frequencies, and covariance of terms and themes, thus subsuming classical Content Analysis.

7.1.2 Semantic Network Analysis

In map analysis, a concept is a single idea, or ideational kernel, represented by one or more words. Concepts are equivalent to nodes in Social Network Analysis (SNA). The link between two concepts is referred to as a statement, which corresponds with an edge in SNA. The relation between two concepts can differ in strength, directionality, and type. The union of all statements per texts forms a semantic map. Maps are equivalent to networks

7.2 AutoMap

Texts, e.g., newspaper articles, blogs, and the stories told by people, are a key source of cultural and ethnographic information. AutoMap [151] is a text mining tool that enables the extraction of information from texts using Network Text Analysis methods. AutoMap supports the extraction of several types of data from unstructured texts. The type of data that can be extracted includes: content analytic data (words and frequencies), semantic network data (the network of concepts), meta-network data (the cross classification of concepts into their ontological category such as people, places and things and the connections among these classified concepts), and sentiment data (attitudes, beliefs). Each of these modes assumes the foregoing.

Coding in AutoMap is computer-assisted; the software applies a set of coding rules specified by the user in order to code the texts as networks of concepts. Coding texts as maps focuses the user on investigating meaning among texts by finding relationships among words and themes. The coding rules in AutoMap involve text pre-processing, statement formation, and post-processing which together form the coding scheme. AutoMap exists as part of a text mining suite that includes a series of pre-processors for cleaning the raw texts so that they can be processed and a set of post-processor that employ semantic inferencing to improve the coding and deduce missing information. These pre-processors include such sub-tools as a .pdf to .txt converters, non-printing character removal, and limited types of de-duplication. Text pre-processing condenses data into concepts, which capture the features of the texts relevant to the user. Statement formation rules determine how to link concepts into statements. The postprocessors include such tools procedures that link to gazetteers and augment the coding with latitude and longitude, belief inference procedures, and data secondary data cleaning tools. In addition there are a series of support tools for creating, maintaining, and editing delete lists and thesauri. AutoMap exports data in DyNetML and can be used interoperably with *ORA.

AutoMap is focused around the idea that meaning is carried in the way in which concepts are linked [148]. Concepts are words or phrases that represent a single ideational kernel; e.g., hope or United_States_of_America are both concepts. To identify the concepts, non-content bearing words are often deleted and thesauri are used to map alternative spellings and phrasing into a single concept. Syntactic clues are used to define connections among concepts leading to stronger linkages being built among concepts within the same phrase, than in the same sentence, than in the same paragraph. In its simplest form, a semantic network is built by building a network

where two concepts are linked just in case they are within so many words of each other or occur in the same sentence. Ontological thesauri that map concepts into categories are then used to cross-classify concepts into agents, organizations, knowledge, resources, locations, beliefs, tasks and events. This cross-classification results in a set of networks – i.e., a meta-network [152].

AutoMap is first used to extract entities (the nodes), then links, then to cross-classify entities into ontological categories. Entity extraction involves locating and classifying terms that represent instances of entity classes of the meta-network that deviate from the classical set of entities in text data. Unlike traditional text mining, which focuses only on named entities (people, places, organizations), we also extract more fuzzy entities, such as tasks (e.g. signing a contract) and resources (e.g. vehicles), which are not necessarily referred to by a name. The following excerpt from an UN News Service (New York) article released on 12-28-2004 illustrates the EE task:

Jan Pronk, the Special Representative of Secretary-General Kofi Annan to Sudan, today called for the immediate return of the vehicles to World Food Programme (WFP) and NGOs.

The underlined concepts are the entities in the meta-network. The quality and accuracy of the extracted network depends on the quality of the entities extracted. AutoMap uses a combination of sub-models to extract these entities. These sub-models include utilization of thesauri and the use of Conditional Random Fields for entity identification. Conditional Random Fields allow for modeling the relationship among y_i and y_{i-1} as a Markov Random Field (MRF) that is conditioned on x . MRF are a general framework for representing undirected, graphical models. In CRF, the conditional distribution of an entity sequence y given an observation sequence (string of text data) x is computed as the normalized product of potential functions M_i [153], [154].

The resulting entity extraction process using Conditional Random Fields consists of two steps. First, the Conditional Random Field is used to locate the terms that are relevant entities. These terms are then marked as being a part of a relevant entity. Second, the Conditional Random Field is used to classify the identified relevant entities. In order to do this, consecutive words that have been identified as belonging to entities are merged into one concept. This concept is represented as a concatenation of the consecutive entity words.

When using AutoMap to identify adversarial networks, the following features were particularly useful: anaphora resolution, deletion of stop words, thesaurus generalization, and meta-network thesauri (ontological cross-classification). In general, a windowing technique was used for placing links and links were placed among entities occurring within a window defined as two contiguous sentences. Finally, gazetteers were used to add latitude and longitude for locations terms.

7.2.1 Anaphora resolution

Anaphora resolution identifies the social entities that pronouns refer to. Co-reference resolution identifies multiple instances of unique real-world entities that multiple text phrases reference. The application of these preprocessing steps in the process of extracting relational data from unstructured text data can impact the entity frequency count, identity of entities and of the identification of relations between entities. It is not uncommon for these steps to modify 15 percent of the edges.

7.2.2 Deletion of stop words

Stop words are those words whose presence has little content of value to the analysis. Often, words such as a, an, the, to, for will fall in this category. Lists of common stop words exist in the machine learning community. We used these and augmented them with a set of commonly unused concepts in assessing adversarial relations. We have found it efficacious to remove most articles, prepositions, numbers, terms referring to temporal indicators such as days of the week, and terms referring to intensity such as more or fewer.

These stop words are collected into a delete list. These concepts are then removed before additional work on thesaurus construction is done. In general, you should create a cut-off limit (e.g. a word needs to be used at least three times. Concepts used less than that would be placed in the Delete List.

7.2.3 Thesauri generalization

One of the key issues in assessing texts is that different words are used to describe the same thing. For people, we might think of these alternatives as aliases. Thesauri are generally used to take multiple concepts, in different forms, and compile them under one key concept. The purpose of a generalization thesaurus is to cluster together all those concepts that refer to the same entity effectively forming a set of coding rules for translating those concepts into the general term. This generalization process can be used for aliases, to remove alternative ending, decrease the impacts of plurals, and combine concepts where differences in nuance are not relevant to the analysis.

Standard stemmers, which reduce words to their base such as farming and farmed to farm tend to over generalize and do not retain part-of-speech distinction. This we prefer to use specially design stemmers that preserve part-of-speech thus enabling auto-identification of tasks/activities and generic actors. These specialized stemmers are part of the automatically constructed generalization thesauri.

For adversarial reasoning, the key effort in thesaurus construction needs to go into the construction of alias files for organizations and people. These tend to be specific to adversarial group when referring to names entities and so specific people and groups. In contrast, other thesauri referring to activities or “generic” people, e.g., farmers, can be used across studies.

7.2.4 Meta-network thesauri

The meta-network is an ontological categorization of nodes in this case concepts into the who, what, when, where, how why needed to assess groups [152]. The meta-network is a multi-mode, multiplex model that reifies these entity classes as: agent, knowledge, resource, task, event, organization, location, belief, time.

Instance of an entity class can have attributes, e.g. the attribute of agent *John* might be *age*, 42 and *gender*, *male*. The relations among the elements within and across any entity classes form certain types of networks. For example, a social network is composed of relations among agents, and a membership network consists of connections among agents and organizations. The meta-network model allows for analyzing socio-cultural systems as a whole or in terms of one or more of the networks contained in the model. This ontological schema has been used to empirically assess power, vulnerability, and organizational change in a diversity of contexts such as situational awareness in distributed work teams, email communication in business corporations and counter terrorism [155], [156], [157].

7.5 Data to Model Processing

Each text is processed to remove noise and clean the text, to combine multi-word concepts into a single concept, to normalize the concepts into a reduced vocabulary, and to categorize concepts into the meta-network ontology [152]. The meta-network ontology includes agents, organizations, locations, events, knowledge, resources, and tasks (i.e. activities).

Initial cleaning of the texts involves reformatting as well as a generic cleaning. The generic activities include preprocessing used to correct the text. Examples include typo correction, the expansion of contractions and abbreviations. Pronoun resolution should be done and unidentified pronouns removed. Identification of compound concepts is done by applying a list of concept-changing n-grams. While typically the use of an n-gram is to identify words that are most commonly used together, in this context an n-gram is a multi-word concept whose definition changes when the concepts are reviewed individually versus as a single compound entity. Examples of concept-changing n-grams are "first aid" and "black market".

Concepts are segmented into specific and general types. The specific concepts identify instances of items, such as George W. Bush for agent, UNICEF for organization, and Pittsburgh for location. General concepts include soldier (agent), tank (resource), and base (location). Many of the general concepts in the ontology can be pre-established. Some minor adaptation needs to be done based on the domain as "front" is different for a military domain as opposed to "front" when speaking about weather forecasting. Some specific entities can be pre-established from existing lists such as a list of all countries or major cities, or a list of world leaders.

The processing material requires project-based modifications beyond what can be pre-established. Project-based specific entities can be found by reviewing all proper nouns identified in the corpus by applying part-of-speech analysis. For convenience, proper nouns adjacent to one another can be listed in an n-gram form as many project-based specifics are compound concepts. Alternatively, the approach to finding specific compound concepts would involve the generation of all n-gram possibilities, which is prohibitively large for human review beyond bigrams (n-grams with N=2). The list of possible concepts of interest can be culled by removing all concepts already placed in an ontological category. Using the pre-established material significantly reduces the amount of human involvement.

A base thesaurus is formed from the pre-existing ontologically categorized concepts and augmented with project-based material. The current pre-existing generics material consists of 22,455 entries. The current pre-existing base material (including general and specific entities) consist of 150,749 entries. A number of scenarios were examined including the following:

1. A scenario driven deterrence assessment: Uses a corpus of 27,000 text files from news sources and government websites. The project-based thesaurus added only an additional 962 entries. The resulting meta-networks contained 7,605 entities.
2. A military multi-actor experiment. Uses a corpus of 3,100 text files from news sources, web sites, and communication logs. The project-based thesaurus added only an addition 500 entries.
3. Open-source information on the Sudan. Uses a corpus of 71,000 text files from news sources, web sites, books, as well as additional information from a wide variety of collected information by scholar experts. The project-based thesaurus includes 38,552 location listings extracted from a gazetteer leaving 16,001 unique entries.

The use of pre-existing thesauri reduced significantly the amount of work needed to extract a meaningful meta-network. The cleaning of the text, extraction of proper nouns and the subsequent removal of pre-existing items for human review, and the generation of the meta-network using pre-existing and project-based thesauri, are all examples of workflows. The workflows are a sequence of common steps used to perform a task, often using different inputs or different thesauri. These web services can be composed into domain-specific workflows that can be used by analysts to automate and manage common sequences of operations. The workflows automate the task management allowing the analyst to focus on the domain and not on keeping track of the individual steps to be taken. By sharing workflows, a consistent approach for data-to-model is established. When advances are made, the workflow can be easily adapted to the new capability and the data-to-model processing re-run.

7.6 Limitations and Next Steps

The global learning of features along with their corresponding weights comes at a price: Training the identifier and classifier while using a reasonable iteration rate for the gradient takes a very long time. This limitation can be addressed to some degree by using more powerful hardware, especially by using more memory. Furthermore, an ability to add, change, or remove labels from the used ontology is essential to having a flexible yet robust learning and research process. While the meta-network has many labels of interest, it is likely that the model may be altered as it evolves in the future.

Conditional Random Fields enable us to detect relevant entities along with their corresponding weights without having to have any preliminary or initial guess about what some of those features might be for a particular data set or domain. This means we can let the computer do all the work as long as we provide it with some labeled training data. However, such uninformed global learning approach comes at a price: Additionally, other techniques for improved entity extraction should be considered. These include things such as improved anaphor resolution, entity inference for beliefs and events, and attribute extraction for concepts such as automated cross classification of resources and activities by DIME/PMESII areas.

However key improvements will require improved link identification. Extracting network ties, or relations between entities, is substantially harder than entity recognition. State-of-the-art systems perform less well on this task than on the recognition task. Most research on relation extraction assumes that the entities have been identified correctly. Main methods for extracting relations between entities are to discover verb relations [158], construct concept graphs based on rules [159], or use proximity to find relations within a sentence using a “word window” [160]. These techniques however need to be augmented with syntactic hierarchical parsing; i.e., placing links within clause, then sentence, then paragraph. In addition, machine learning techniques may also be used for improved link identification.

Chapter 8

Inferring and Assessing Informal Organizational Structures from an Observed Dynamic Network of an Organization

Il-Chul Moon, Kathleen M. Carley, Alexander H. Levis

8.1 Introduction

In today's world there are many organizations or groups that are organized virtually or covertly. Open source project teams, teams in massive multi-player on-line games, and terrorist organizations are just a few examples. For these organizations, what is known is what can be observed. What can be observed are the networks connecting individuals, resources, and activities across many lines and types of communications? Clearly there are many types of relations in this observed structure not all of which are necessarily work related. For these organizations, the organizational chart, the workflow, the formal structure is likely not to be known a priori. Indeed, it is unlikely that there is a formal structure in the sense of a declaration by the organization about who reports to whom and is doing what. Nevertheless, it is likely that the operational structure of the organization, who shares information with whom, resolves issues, etc. is embedded in the observed structure. If we could infer this operational structure from the observed structure we would have an improved understanding of how work is done in these groups, their strengths, and their vulnerabilities.

We propose and approach for inferring the operational structure from the observed structure. The observed and the operational structure are likely to have distinct profiles, e.g., key personnel and clusters of individuals. This is because the operational is focused only on work related activities whereas the observed is a concatenation of all activities, a snapshot of human endeavors. We illustrate the efficacy of this approach using data collected on a real-world, terrorist organization. The proposed approach expands the horizon of organizational analysis by enabling researchers to identify and assess these operational structures.

Understanding an organization's structure is critical when we attempt to understand, intervene in, or manage the organization [161]. However, organizational structures in the real world often differ from their recognized formal structure [162], and sometimes its membership conceals the formal structure with various types of social interactions and communications [163]. Furthermore, when we observe the actual social interactions among the members of the group, the observed social-network data are often noisy, and contain misleading and uncertain links [164]. The following two scenarios exemplify the impending confusion about the identification of an organizational structure.

***Scenario 1:** TF is an employee of a global investment bank in Hong Kong. In the formal organizational chart, he reports administratively to the bank's financial division director, who is also the head of the Hong Kong branch. However, because of his assignment to work on a global project, TF also reports to two senior project managers who manage the project from their offices in New York and London. His corporate email activity includes not only personal-activity re-*

porting emails to the local administrative manager and project status reporting to the two project managers, but also includes information sharing emails to work colleagues.

Scenario 2: IC is a developer in a software development team which is informally organized. He has frequent email contact with the users and other team members, including one team leader with whom he often reports his progress. Because there is no formal team structure, there is no membership boundary in the team, so the team involvement is determined by consensus from the active members who are reporting bugs and developing programs.

In these scenarios, we identify three different types of organizations that vary in their boundaries and explicitness. Firstly, the organizational chart unequivocally outlines the formal hierarchical structure, but the employees have another hierarchical reporting structure that is not shown in the formal chart. Secondly, his email account shows his contacts, regardless of the contacts' importance or the nature of the relations, so the uncovered email transaction structure from his account contains people with critical work relationships and ones with insignificant relationships at the same time. This second organizational structure is a *social network* in this paper. The third structure, our definition of *decision making organizational structure* in this paper, is a social structure including only relevant personnel, or three formal or informal bosses, and work relationships, or reports to the bosses in terms of completing the organization's goal. These different organizational structures can be also seen in diverse organizations, i.e. grass-roots organizations, self-organizing clubs, startup companies, terrorist networks, military command and control structures, etc. This paper uses a terrorist network as a test dataset.

We focus on the differences in analysis approaches regarding the two above organizational structures: *meta-network* (an extended version of *social network*) and *decision making structure*. Meta-network is a network representation of a complex organizational structure. Its dataset is gathered from email transactions, survey from group members, observations on social interactions, and etc. Meta-network analysts concentrate on finding key personnel, i.e. which boss is more important in Scenario 1. Or, they find clusters, i.e. clusters of developers of open-source development team in Scenario 2. Decision making structure is an organization structure design whose members, or decision makers, interacts with each other in various purposes over the course of decision making. The structure is from organizational charts, survey, or subject-matter experts of the organization. Decision making structure analysis uncovers the information and response transmissions in members' cognitive processes while a decision is made, i.e. when TF's report weigh in the formal or informal bosses' decision making processes in Scenario 1, to what extent IC and his discussion partner share the information and when in Scenario 2.

Considering the above two perspectives, we need the third approach that combines the two. We can combine the approaches in many ways, i.e. regarding a critical organizational structure as a decision making structure and applying social network analysis to the structure (applying social network analysis to a decision making structure). Or, we can see the meta-network as a decision making structure and estimate the cognitive processes of members of the network (applying decision making structure analysis to a dynamic network). In this paper, we introduce one approach combination. First, we extract the decision making organizational structure from an observed meta-network of a target organization. For instance, we extract the only relevant people in the decision making processes among TF's contacts in Scenario 1. This extraction is done by considering the work relationships among the members of the group and the work flow of the organizational objective. Next, we analyze the extracted decision making structure with the social network analysis approach. For example, among the three bosses and TF in Scenario 1, we

identify the most important personnel in terms of information delivery, situation cognition, linking to others, by utilizing social network metrics. Then, we can see the different key personnel lists and clustered members between the original meta-network and the extracted decision making structure. These differences imply that the analysis result can be richer if we investigate not only the existing meta-network, but also the inferred structures from it.

The workers segregate and create clusters socially based on the work flow rather than their formal structure. If this is true, an analyst may find out how well the formal structure supports the current work practices by comparing the formal and informal structures. As another example, Rabasa et al. [165] think that Al-Qaeda operatives may be embedded in a social network of a community including civilians and operatives at the same time. Although they co-exist in the social network, it is certain that management activities occur among the operatives. The decision making structure extraction will reduce or limit the relevant personnel in the social network, will help set the scope of investigations, and produce various analysis results from different decision making structure viewpoints. Finally, this work is an effort linking two different disciplines, social network analysis, and decision making structure analysis. Meta-networks have been gathered from various terrorist networks and military organizations, but these have not been used frequently in the decision making structure analysis domain because the interpretations of the meta-network and decision making structure are different. With the proposed framework, we successfully extract a decision making structure from a meta-network, so that we can use the existing meta-network datasets in further decision making structure analyses.

8.2 Background

Our framework is presented in two steps: (a) inferring a decision making structure from a meta-network, and (b) analyzing the extracted structure with social network analysis metrics and algorithms. Thus, the theories behind our approach are twofold. First, we explain the complex nature of a meta-network and how we exploit the complex organizational structure in inferring its decision making structure. Second, we describe the used social network analysis metrics and algorithms.

8.2.1 Inferring a decision making structure from a complex system of an organization

The organizations of interest in this paper exhibit the characteristics of a complex system. According to Morel and Ramanujam [166], there are two commonly observed characteristics of a complex system: a large number of interacting elements and emergent properties. First, a corporate organizational structure consists of a large number of interacting elements such as workers, information, expertise, and resources [167]. These elements should be assigned and distributed properly to perform tasks, and such assignments and distribution relationships are the organizational structure of the corporation. Similarly, a terrorist network is a collection of heterogeneous entities interacting with and assigned to each other. Though a traditionally terrorist network was regarded as a simple terrorist-to-terrorist network [168], [169], recent observations and analyses [170], [171] assert that the terrorist network includes bomb materials, reconnaissance on targets, as well as terrorists.

Second, the organizations of interest have emergent properties. A *synthetic organization* [172] is an organization established after a major event, such as a disaster. The organization emerges around formally designated offices by linking NGOs and relevant groups to the offices. The organization self-organizes the work relationships and seeks a better structure over the course of the event. This emerging structure concept can also be applied to corporate and terror-

ist network domains. Employees of a corporation have their superiors and take orders from them, as in a hierarchical organization, but they also keep and follow work relationships in practice. Also, it is often seen that a task-force team emerges before or after important events [173]. This task-force team shows the emergent properties of the organizational structure in a corporation. Additionally, terrorist networks frequently show the emergent properties by adapting their structures to situations [174], [175].

If the organizations in focus are complex, we should find a decision making structure by considering the various types of interacting elements and the adaptive nature of the structure. At the same time, since the traditional organizational structure is defined as a structure managing individuals in an organization, the found structure should contain people-to-people relationships. Thus, we focus on developing a model that takes the complex nature into consideration and generates a set of work relationships among the individuals.

CAESAR III [176], [96] is a model that we regarded as a base of our developed model. Originally, it was used to analyze the cognition processes of multiple decision makers. The individual cognitive processes are structured as a network of various types of links that differ in terms of inputs and outputs of the cognitions. Thus, the model is similar to our approach. Therefore, our major effort in this paper is inferring the links of cognitive processes among individuals from a meta-network.

8.2.2 Assessing vulnerabilities and criticalities of the organizational structure

There have been a number of approaches in evaluating the organizational structures. For instance, traditional management science developed qualitative evaluation criteria [177]. However, though these qualitative examinations are insightful, the qualitative approaches have problems. They are not scalable to large and complex organizations, nor applicable to various disciplines, and nor designed to assess the complex representation of a meta-network. Therefore, in this paper, we will use a quantitative model.

Social network analysis has been one of the most useful tools in analyzing organizational structures, i.e. corporate structures and terrorist networks [168], [178]. It is able to find key personnel [186] and embedded clusters. Also, it assesses the characteristics, such as degree of centralization and levels of hierarchy, of the organizations. Gabbay and Leenders [179] link the social network analysis to the management of social capital of a corporation. Also, Reagans and Zuckerman [180] investigate the performances of various corporate R&D teams with social network analysis. This analysis is used not only in the corporate domains, but also in the counterterrorism field, and Krebs [168] visualized the terrorist network responsible for the 9/11 attacks and calculated the social network centrality metrics of terrorists.

In this work, we follow the basic approach of social network analysis, which involves calculating the social network metrics and finding key entities in the structure. However, we are different from the traditional social network analysis in two ways. One way is that we analyze both the original meta-network and inferred decision making structure. The other way is that we use a couple of metrics, cognitive demand and communication [181] - which are not common in social networks, but insightful in examining a complex organization. Furthermore, we use QAP and MRQAP analysis techniques. These techniques have been used to correlate two networks and regress one network against another. We correlate the inferred structures to the original structure to examine to what extent the extracted ones are embedded in the original ones.

8.3 Dataset

Throughout this paper, we use a dataset collected from the 1998 U.S. Embassy bombing incident in Kenya. As the organizations of interests exhibit complex organizations, we use a meta-network format [187] to represent and analyze the target organization. Meta-network is an extended version of a social network, including various types of nodes and heterogeneous links, which follow the nature of a complex system. Initially, this dataset is from a network text analysis [182] on open-source documents, but later, the soundness and realism of the dataset were verified by human analysts. This meta-network dataset is appropriate for this analysis for a couple of reasons. First, it has a directed terrorist-to-terrorist network required for inferring a *Command Interpretation* structure, which will be explained later, included in the expected decision making structure. Second, it has a detailed task network. With inputs from human analysts, the dataset has a detailed task procedure of the incident, so it is particularly appropriate when we extract a decision making structure for the completion of a certain task.

As our framework starts with a meta-network, the initial input dataset is a collection of terrorists, information and resources for the bombing and related tasks. Fig 8.1 is the visualization of the meta-network of the Kenya case. Also, we visualized two sub-networks, the terrorist social network in Fig. 8.2 and the task precedence network in Fig. 8.3. The basic statistics of this network is listed in Table 8.1 and Table 8.2. For each of the sub-networks, there is an interpretation for the links. For instance, the link in a social network represents that the two terrorists interacted or communicated with each other, and the link in a task assignment network shows that the terrorist was assigned to completion of the linked task.

TABLE 8.1 The meta-network of the dataset, a terrorist group responsible for 1998 U.S. embassy bombing in Kenya. The numbers in the cells are the densities of the adjacency matrices.

	Terrorist	Expertise	Resource	Task
Terrorist (17 terrorists)	Social Network (0.147)	Information Distribution Network (0.095)	Resource Dis- tribution Net- work (0.088)	Task Assign- ment Network (0.126)
Expertise (8 bits)		Not used	Not used	Required Ex- pertise Network (0.048)
Resource (8 resources)			Not used	Required Re- source Network (0.076)
Task (13 tasks)				Task Prece- dence Network (0.121)

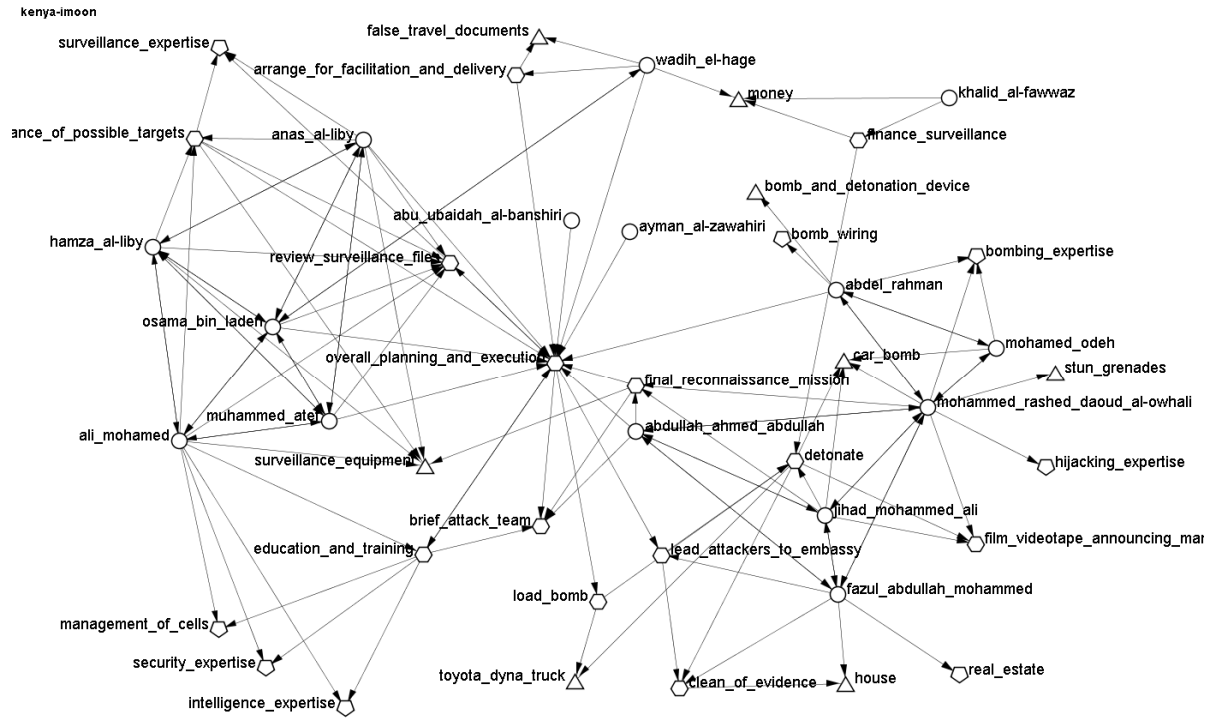


Fig. 8.1 The visualization of the meta-matrix of the terrorist group responsible for the 1988 U.S. embassy bombing in Kenya

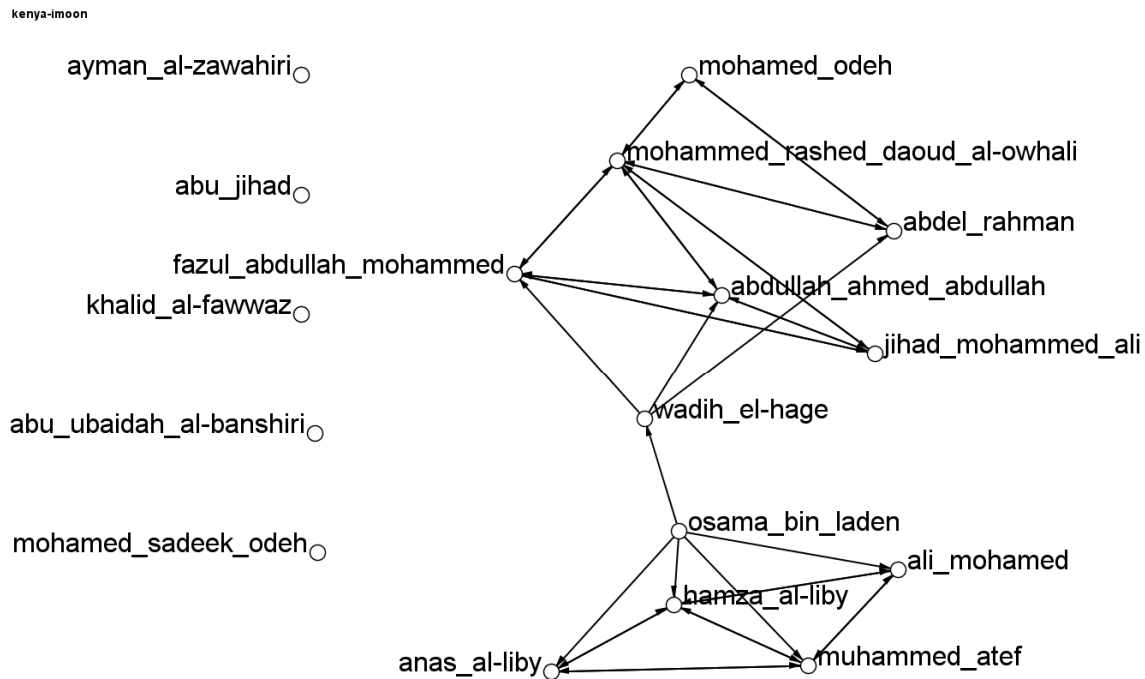
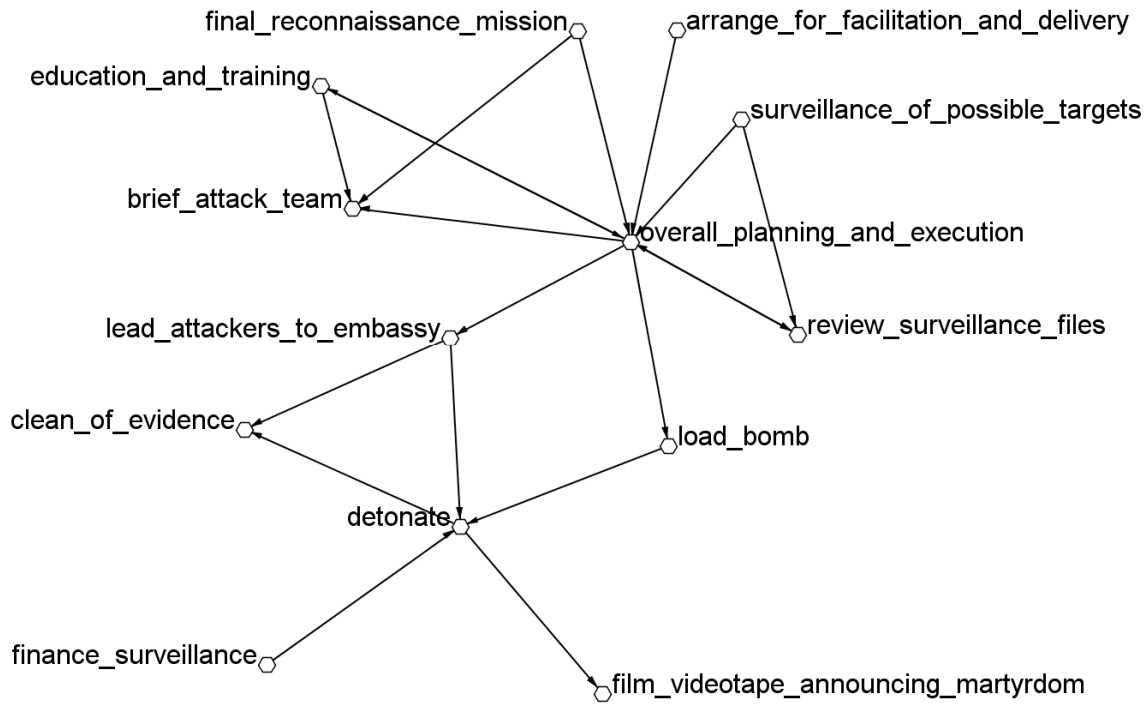


Fig. 8.2 The terrorist social network in the meta-matrix



powered by ORA, CASOS Center @ CMU

Fig. 8.3 The task network in the meta-matrix

TABLE 8.2 A table of descriptive statistics for the metrics. This table includes means, standard deviations, and a cross-correlation table.

	Mean	Std. Dev.	Total Degree Centrality	Betweenness Centrality	Eigenvec- tor Cen- trality	Cogni- tive Demand	Communica- tion
Total Degree Centrality	0.0921	0.0478	1.0000	0.7411	0.4880	0.9113	0.4030
Betweenness Centrality	0.0065	0.0073	0.7411	1.0000	-0.0650	0.8384	0.2870
Eigenvector Centrality	0.0337	0.0204	0.4880	-0.0650	1.0000	0.3087	0.3802
Cognitive Demand	0.0681	0.0393	0.9113	0.8384	0.3087	1.0000	0.3929
Communica- tion	0.6964	0.1897	0.4030	0.2870	0.3802	0.3929	1.0000

8.4 Method

Our framework is about extracting a decision making structure from the meta-network of an organization as well as analyzing and comparing the extracted structure and the original meta-network. In this section, we introduce how to infer a potential decision making structure in the first stage and network metrics in the second stage.

While the analysis procedures are largely in two steps, there are five detailed stages in this analysis framework. The extraction requires three stages. First, we obtain a target organization to analyze and its task of interest. Second, we identify the sub-task network by including only relevant tasks to the completion of the task of interest, and this leads to limiting the personnel involved. Third, the target organization is examined from three perspectives: information sharing, result sharing, and command interpretation. Each of the examinations generates a decision making structure corresponding to the perspective.

The analysis and comparison are done in two steps. First, we compare the extracted structure to the original network. Additionally, we estimate to what extent we can recreate the original structure with the extracted ones. These comparisons show the effectiveness and the usefulness of the extraction overall, since we expect the extracted structure to be based on the meta-network, but not be exactly the same structure. Second, we evaluate the network metrics of individuals, identify the key personnel, and see the differences between the key personnel list from the original and the extracted structures.

This framework is also designed to convert the meta-network into an input dataset for CAESAR III model, a decision making structure analysis framework. While we discuss and experiment inferring a structure for CAESAR III from a meta-network, we do not utilize CAESAR III to analyze the extracted model from its viewpoint. Our evaluation analysis is limited to social network approaches.

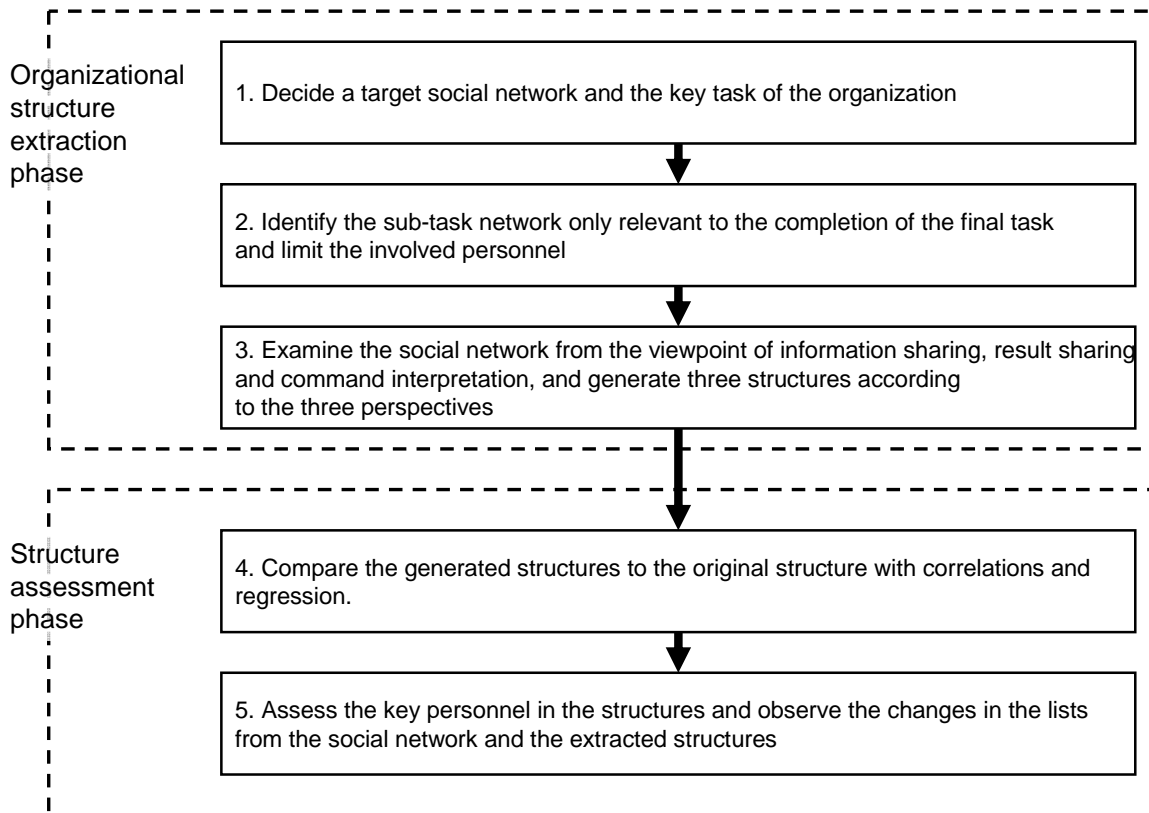


Fig. 8.4 The procedure of the introduced analysis framework

8.4.1 Extracting a decision making structure from a meta-network

The scope of the decision making structure is limited by focusing on a single task execution. This way restricts the number of individuals who make up the extracted structure and makes the others as the outside collaborators. As the number of individuals of interests decreases, we can focus on the investigation of the specific task performance and keep the generated structure recognizable to human analysts. Also, in the management science community, these selected individuals are regarded as decision makers, so this limitation differentiates between a social agent and a decision maker in the structure.

After selecting the decision makers, we infer the various management relations by utilizing the social network as well as the task assignment, the information, and the resource distribution networks. For instance, when two members are connected with a communication path and one has expertise required for the other, the shortest path may be the information sharing path in terms of management relationships. With similar methods, in addition to the information sharing relationships, we infer result sharing and command interpretation relationships. These are originated from three different structural links in the CAESAR III model. In the model, information sharing, result sharing, and command interpretation links are different in their timings of message arrival. Information sharing messages are delivered after the sender is aware of the situation and before the receiver performs the information fusion. Result sharing is done after the sender's response selection. Command interpretation occurs before the receiver's response selection. The information fusion, response selection, and command interpretation are the cognitive processes defined in CAESAR III.

Limiting task network and finding decision makers

Since the decision making structure in this paper is task-oriented, our framework aims to extract a structure responsible for completing a certain final task. This task is a user-defined parameter. With the given final task, we can retrace a sub-task network from a meta-network by following the prerequisite tasks repeatedly, starting from the final task. For example, in Fig 8.4, the final task is *overall planning and execution*; then, its sub-prerequisite tasks are *surveillance of possible targets*, *final reconnaissance mission* and *arrange for facilitation and delivery*. These four tasks consist of the sub-task network for extraction, and the 12 terrorists assigned to those tasks are the decision makers of this task-oriented decision making structure.

After limiting the involved decision makers, we aggregate the uninvolved agents as an outside organization. It is typical to see a decision making structure interacting with outside organizations. If we configure a task-based sub-decision making structure, some of the individuals will be excluded, since they are not doing the tasks in the sub-task network. However, it is still possible that the excluded ones hold required resources or information, and this will require communications between the selected decision makers of the extracted structure and the outside organization, which is the group of the excluded individuals. Thus, finding assigned decision makers doesn't just limit the personnel of the decision making structure, but also specifies the boundary decision makers interacting with outside entities. In this example, we have a total 17 terrorists, and 12 terrorists are selected as decision makers. Thus, the other 5 terrorists form the outside organization of this decision making structure.

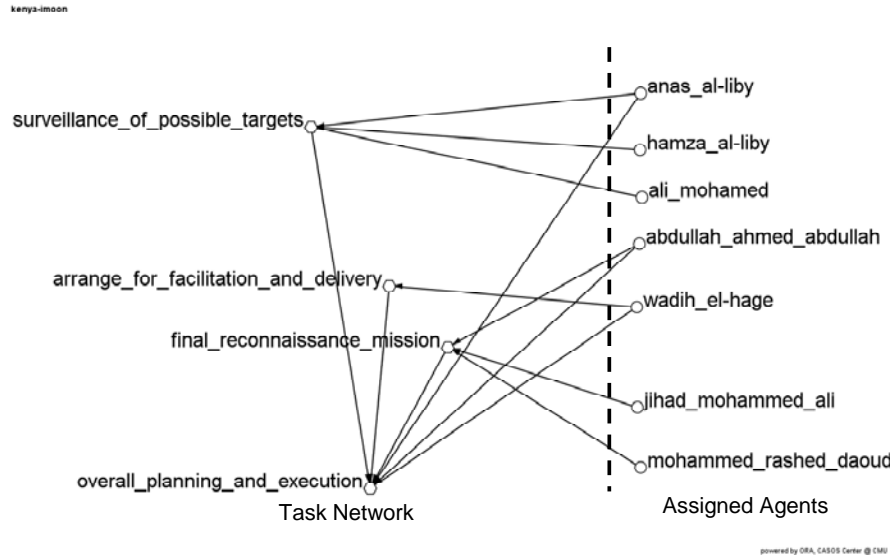


Fig. 8.5 The partial visualization of the task precedence network (task-to-task) and the task assignment network (terrorist-to-task). The dashed line represents the separation of the task network and the assigned agents. When users set up *overall planning and execution* as a final task for the extraction, the visualized tasks and the individuals are the components of the sub-task network, and the accompanying decision makers, respectively.

Information sharing structure

In a meta-network, a piece of information, or expertise, is represented as a knowledge node. Thus, we assume that producing information is represented as a link from an agent node to a knowledge node. Also, we infer that one decision maker will acquire information through an information sharing path if 1) he needs the information to perform his assigned tasks, 2) he does not have the information, and 3) the information sharing path is the shortest path from the nearest decision maker holding the information for him. Figure 8.5 describes the case of information sharing links. According to the sub-network in the figure, *Ali Mohamed* is assigned to *surveillance of possible targets*, which requires *surveillance expertise*. However, *surveillance expertise* is not available to *Ali Mohammed*, but available to *Anas Al-Liby*. Then, *Ali Mohamed* finds shortest paths possible to *Anas Al-Liby*, and he finds the shortest paths with two social links going through *Osama bin Laden*, *Hamza Al-Liby* or *Muhammad Atef*. Then, the links in these three shortest paths will be the information sharing links.

Result sharing structure

Result Sharing (RS) is communication from a decision maker finishing his assigned task to a decision maker with a task that required the previously done task. For instance, there is a RS communication from a terrorist who finished *surveillance of possible targets* to a terrorist who will perform *overall planning and execution*. Figure 8.6 shows the above two tasks and their assigned agents. *Surveillance of possible targets* has three assigned agents, and *overall planning and execution* has eight agents. Then, there will be 21 result sharing links originating from the three agents to the seven agents, excluding the agent who is assigned to the next task and already knows the results of the previous task.

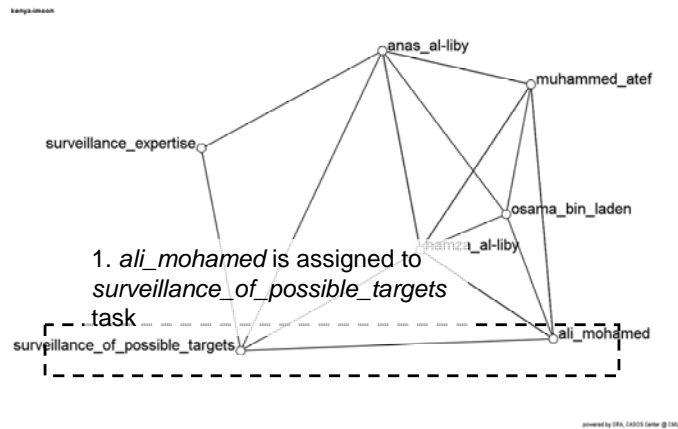


Fig. 8.6a A partial visualization explaining the formation of information sharing links: First step, *Ali Mohamed* is assigned to *surveillance of possible targets*.

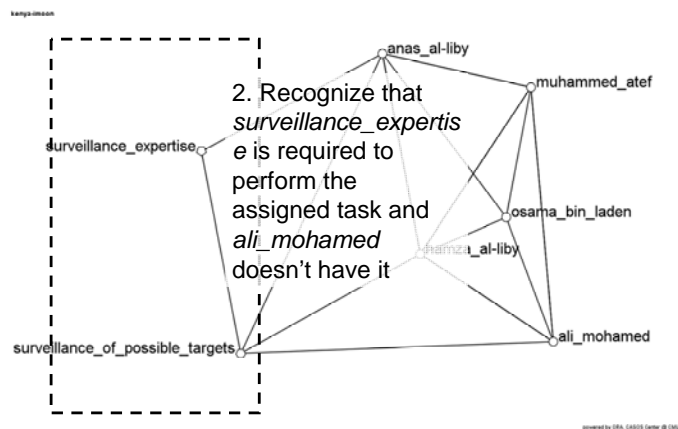


Fig. 8.6b Second step, *Ali Mohamed* requires *surveillance expertise* to perform his assigned task, but he does not have it.

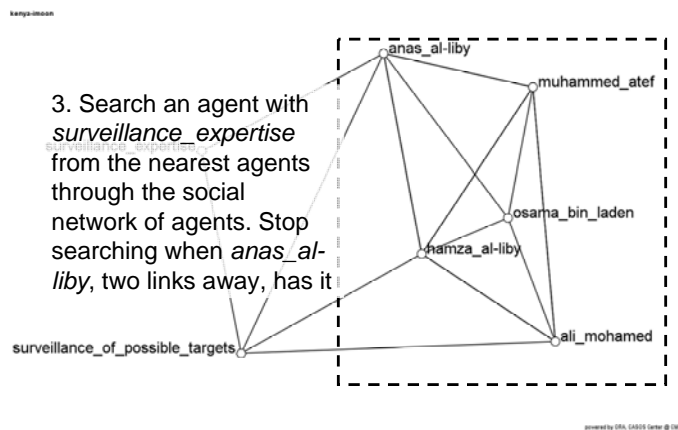


Fig. 8.6c Third step, the organization searches an agent with *surveillance expertise* from the agents near to *Ali Mohamed*. It finds an agent two social links away, *Anas Al-Liby*.

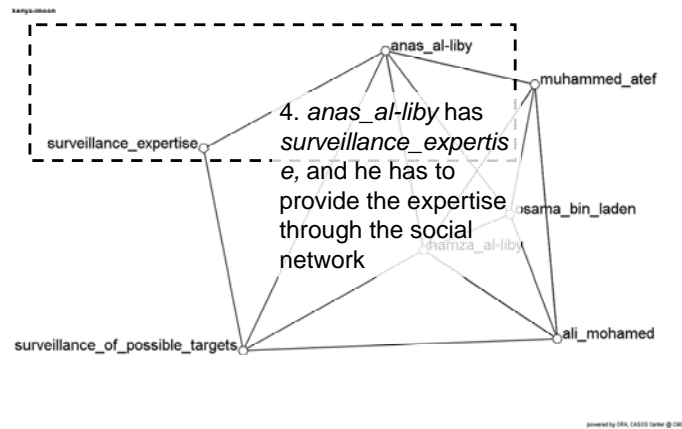


Fig. 8.6d Fourth step, *Anas Al-Liby* has the required expertise and has to deliver the expertise through the social links.

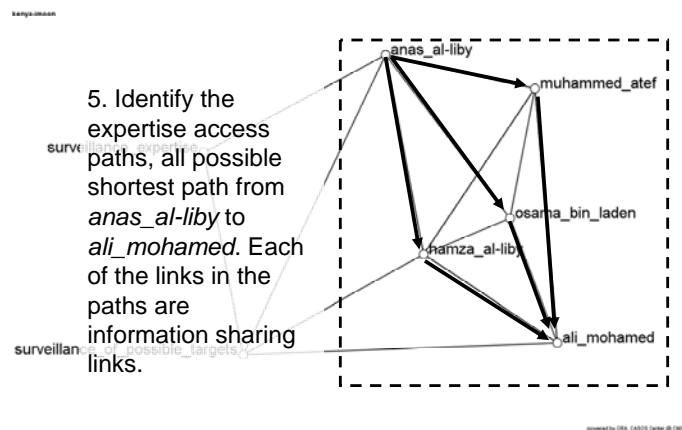
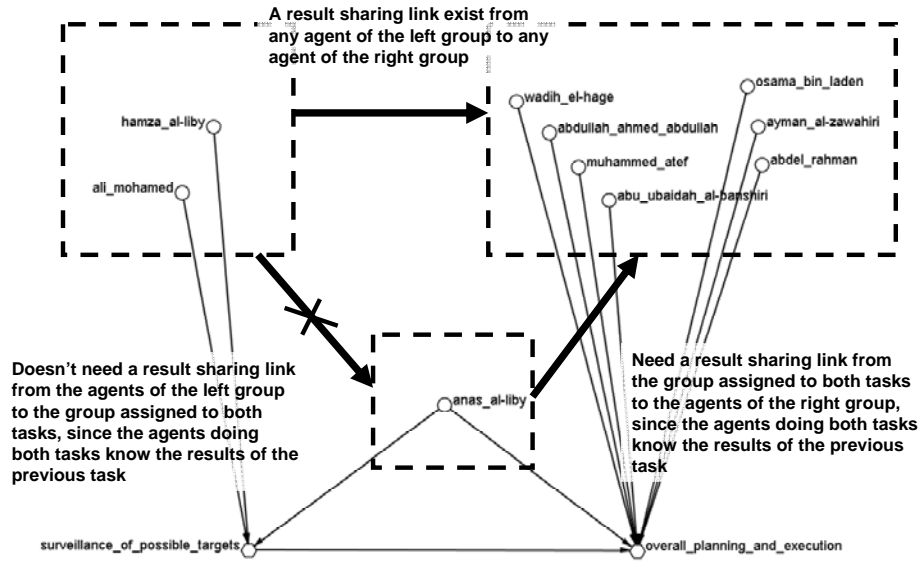


Fig. 8.6e Fifth step, there are three possible shortest paths from *Anas Al-Liby* to *Ali Mohamed*. These paths are information sharing links.

8.4.2 Assessing a network structure with measures

The original meta-network and the inferred decision making structures are all in the meta-matrix format. Therefore, we apply network analysis metrics to assess the criticality of individuals in a network. The metrics are five: *Degree centrality*, *Betweenness centrality*, *Eigenvector centrality*, *Cognitive demand*, and *Communication*. The detailed interpretation is in Table 8.3.



powered by ORA, CASOS Center @ CMU

Fig. 8.7 A partial visualization of two tasks and ten assigned agents. This precedence task relation will result in 21 result sharing links between the agents doing the prior task and the agents performing the next task. One agent who is doing both does not need any result sharing link.

TABLE 8.3 Three traditional centrality metrics and two dynamic network metrics used to assess the criticalities of individuals in the structure

Name	Interpretation	Reference
Degree Centrality	Number of in-coming and out-going links from a node, Degree of direct influence to others	Freeman [184]
Betweenness Centrality	Number of shortest paths passing a node, Degree of information flow control	Freeman [184]
Eigenvector Centrality	Calculates the eigenvector of the largest positive eigenvalue of the adjacency matrix, Degree of connections to the high-scoring nodes	Bonacich [185]
Cognitive Demand	Measures the total amount of effort expended by each agent to do his/her tasks, calculation details are elaborated below.	Carley [181]
Communication	Measures the communication need of agents to complete their assigned tasks, calculation details are elaborated below.	Carley [181]

8.5. Results

The described decision making structure extraction scheme is applied to the U.S. embassy bombing in Kenya case, and the task of interest was *detonation*. Next, we regress the decision making structures against the original meta-network structure to find which decision making structure is embedded in the observed network and to what extent. After estimating the overall correlation level between the original and the extracted structures, we describe and visualize the extracted structure. Next, we calculate five network metrics on the original meta-network and three different management networks. Comparisons on the calculated metrics provide an insight into who stands out in different settings and why. Also, we identify the clusters based on the factor analysis of the metrics of the four networks.

8.5.1 Initial result and descriptive statistics

Figure 8.7 is the visualization of the extracted decision making structures for the detonation task, and the image is generated by ORA [185]. The collection of these extracted networks is an input dataset for the CAESAR III model, and subsequent cognitive process analysis in decision making structure can be done with the model. However, we leave the analysis as our future work in this paper. Whereas the original meta-network has 17 members, the extracted structure has only 14. The removed members are not related to the task network of detonation. The topologies of the structures are different. First, the information sharing structure is somewhat similar to the person-to-person network of the meta-network. The inference of the information sharing is done by trimming the links not included in the information passage. Therefore, the base of the information sharing is the person-to-person network (social network), so the inferred network resembles the social network. Second, the result sharing network is very different from the social network. The result sharing is inferred from the task dependency network and task assignment network. Due to the difference between the result sharing structure and the social network, this organization may suffer from the delivery of information about the completion of prerequisites during the task execution period. Finally, the command interpretation structure only includes three individuals. In the original social network, most of the individuals are linked as a circle with directed links. Therefore, the inference on the command interpretation is not clear for most of the members. However, *Osama bin Laden*, *Wadih el-Hage* and *Abdel Rahman* show a clear hierarchy in the social network. We do not believe that the actual command interpretation is as sparse as the inferred structure, but from the observed structure, there is no clear way to infer the hierarchy of the other members.

8.5.2 Embedded decision making structures in an observed meta-network

We analyze how the extracted decision making structure was embedded in the observed meta-network and to what extent. We use the QAP/MRQAP technique to compare and to regress the extracted decision making structures to the original network. This is a statistical analysis to support the qualitative findings of Section 5.1. If the meta-network implies such decision making structures, the correlation and the R-square of the regression result will be high. Table 4 displays the result of QAP correlations between each of the extracted structures and the meta-network. Information sharing is very highly correlated with the original structure. This high correlation is from the heuristic of the extraction. When we extract the information sharing links, we just trim the existing links, not add ones. However, the high correlation also tells us that there were not many trimmed links, which implies that the observed social links served well as information diffusion paths. The low correlation between the result sharing structure and the meta-network is

coming from many additions of links. This means that the network does not adequately support informing the result of the prerequisite tasks to the individuals doing subsequent tasks.

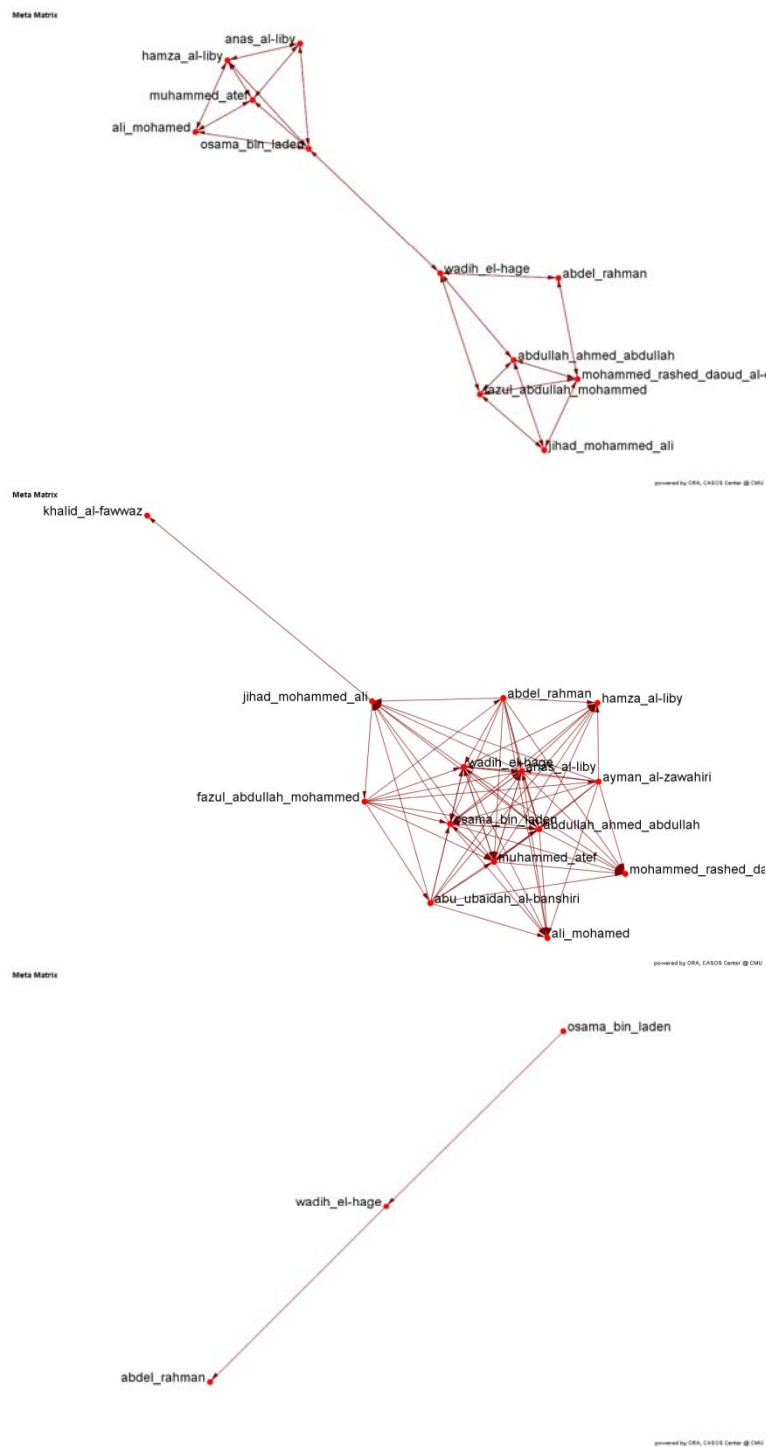


Fig. 8.8 Three extracted decision making structures. (Top) Information sharing, (Middle) Result sharing, (Bottom) Command interpretation

TABLE 8.4 A table of QAP correlation and other distance metrics between the original structure and the extracted decision making structures. (IS=Information Sharing, RS=Result Sharing, CI=Command Interpretation)

	CI	IS	RS
Correlation	0.2204	0.8181	0.1399
Significance	0.0200	0.0000	0.0510
Hamming Dis- tance	34.0000	12.0000	81.0000
Euclidean Dis- tance	5.8310	3.4641	9.0000

The MRQAP analysis in Table 8.5, between the extracted structures as independent variables and the meta-network as a dependent variable, results in a high R-squared value, 0.6759. This is a very high value considering the R-squared is usually very low in MRQAP analyses. As the previous correlation indicates, the information sharing structure was the biggest contributor in estimating the link existence in the meta-network. The levels of standard coefficients of the command interpretation and the result sharing structures are similar. However, the result sharing structure was more significant than the command interpretation while the information sharing was far more significant than the other two. From this MRQAP result, we can see that the original meta-network can be explained by the decision making structures and it embeds those structures. However, the result sharing and the command interpretation are not as well represented as the information sharing.

TABLE 8.5 A table of MRQAP regression results. The dependent network is the observed meta-network, and the independent networks are the extracted meta-network. (R-Squared = 0.6759)

Variable	Coef	Std.Coef	Sig.Y- Perm	Sig.Dekker
Constant	0.0288	0.0000		
CI	0.2080	0.0524	0.2250	0.0400
IS	0.7876	0.8232	0.0000	0.0000
RS	-0.0487	-0.0648	0.1790	0.0600

8.5.3 Personnel with different levels of importance in structures

Table 8.6 shows the top three individuals in the four structures (original meta-network, information sharing, result sharing, and command interpretation) and by using five metrics (degree centrality, betweenness centrality, eigenvector centrality, cognitive demand, and communication). When observing the importance of individuals in the extracted structures, *Osama bin Laden* stands out in the information sharing aspect. In the original observed network, he ranked seventh in degree centrality, ninth in cognitive demand, and has zero betweenness centrality, though he ranked second in eigenvector centrality. However, the information sharing network ranks him second in betweenness centrality. Also, *Wadih el-Hage* is an individual with high importance in the extracted structures. He is not ranked in the top three with any metrics of the original network. However, he is ranked second (RS) and third (CI) in degree centrality; first (IS and CI) in

betweenness centrality; and third (CI) in eigenvector centrality, etc. Actually, *Wadih el-Hage*, whose alias is *the Manager*, was actively engaged in and even managed this terrorist attack. While the human analyst and the network text analyzer generated a meta-network not reflecting his importance, our inference and the meta-network including expertise, resources, and tasks are able to find his importance in the organizational structure. These over- or under-estimations on the criticality of personnel can be found from the metrics of other individuals, i.e. *Anas al-Liby*.

TABLE 8.6: A table of top three individuals from five metrics and four structures (ORI=original meta-network, IS=Information Sharing, RS=Result Sharing, CI=Command Interpretation)

Measure	Structure	Rank 1	Rank 2	Rank 3
Total Degree Centrality	ORI	Mohammed Rashed Daoud al-Owhali	Ali Mohamed	Fazul Abdullah Mohammed
	IS	Mohammed Rashed Daoud al-Owhali	Ali Mohamed	Fazul Abdullah Mohammed
	RS	Anas Al-Liby	Wadih el-Hage	Abdullah Ahmed Abdullah
	CI	Ali Mohamed	Mohammed Rashed Daoud al-Owhali	Wadih el-Hage
Betweenness Centrality	ORI	Mohammed Rashed Daoud al-Owhali	Fazul Abdullah Mohammed	Abdel Rahman
	IS	Wadih el-Hage	Osama Bin Laden	Fazul Abdullah Mohammed
	RS	Jihad Mohammed Ali	Fazul Abdullah Mohammed	Ali Mohamed
	CI	Wadih el-Hage	Abdel Rahman	Osama Bin Laden
Eigenvector Centrality	ORI	Anas Al-Liby	Osama Bin Laden	Ali Mohamed
	IS	Anas Al-Liby	Osama Bin Laden	Ali Mohamed
	RS	Anas Al-Liby	Abdullah Ahmed Abdullah	Osama Bin Laden
	CI	Anas Al-Liby	Ali Mohamed	Wadih el-Hage
Cognitive Demand	ORI	Mohammed Rashed Daoud al-Owhali	Ali Mohamed	Abdel Rahman
	IS	Mohammed Rashed Daoud al-Owhali	Ali Mohamed	Abdel Rahman
	RS	Abdel Rahman	Mohammed Rashed Daoud al-Owhali	Anas Al-Liby
	CI	Mohammed Rashed Daoud al-Owhali	Ali Mohamed	Abdel Rahman
Communication	ORI	Abdel Rahman	Mohammed Rashed Daoud al-Owhali	Jihad Mohammed Ali
	IS	Jihad Mohammed Ali	Muhammed Atef	Wadih el-Hage
	RS	Jihad Mohammed Ali	Muhammed Atef	Wadih el-Hage
	CI	Jihad Mohammed Ali	Muhammed Atef	Wadih el-Hage

Figure 8.9 shows that the difference of metric evaluation results across the original meta-network and decision making structures. Specifically, we subtract a metric value of a meta-network from the value of an decision making structure. Overall, the differences of the metrics are big, which indicates the inference estimated the levels of individuals' importance quite differently. However, the difference in betweenness centralities from the original network and decision making structures are quite similar except for a few individuals.

Osama bin Laden (A0) and *Wadih el-Hage* (A2) show extreme underestimations in betweenness centrality of the original network compared to that of the information sharing network. When we remember that *betweenness centrality* is specialized in the information diffusion passage and the information sharing network is an inferred information flow network from a meta-network, those two are the key personnel in diffusing information pieces in this network. Also, *Abu Ubaidah Al-banshiri* (A11) is somewhat underestimated in the result sharing structure. He has a big positive difference in degree centralities, eigenvector centralities, and cognitive demand, which means that he has higher value in result sharing compared to the original observation.

TABLE 8.7 I.D. assignments to individuals. I.D.s will be used to distinguish individuals in the later tables. We used some abbreviations for names (Fazul= Fazul Abdullah Mohammed, Jihad= Jihad Mohammed Ali, Banshiri= Abu Ubaidah al-Banshiri)

Name	Osama Bin Laden	Muhammed Atef	Wadih el-Hage	Ayman Al-Zawahiri	Anas Al-Liby	Abdel Rahman	Fazul	Al-Owhali
ID	A0	A1	A2	A3	A4	A5	A6	A7
Name	Ali Mohamed	Hamza Al-Liby	Khalid Al-Fawwaz	Banshiri	Abdullah Ahmed	Abdullah	Jihad	
ID	A8	A9	A10	A11	A12	A13		

8.5.4 Personnel clusters with similar characteristics

Since we have four structures and five metrics for each structure, we cannot visualize or cluster the individuals without dimensionality reduction. Therefore, we use principal component analysis (PCA) to project the individuals in two dimensions with highest variances. Table 8.8 shows the coefficients to generate the two components corresponding to the two dimensions, and Fig. 8.10 is the projection of the individuals on a two dimensional scatter plot. The clusters in the plots are member profiles according to the criticality. For instance, there may be a group of people with high betweenness and low degree centrality, and PCA will put those individuals close to each other. We apply this analysis to the two structure sets: the original network and the collection of the three inferred structures. Thus, we can distinguish the different member profiles coming from the original dataset and the inferred dataset. Before the interpretation, it should be noted that we disregarded *Khalid al-Fawwaz* (A10) because he is an extreme outlier in PCA. In the original and the inferred networks, he was the only one who had all the necessary resources to execute his assigned task. This makes him unique in the membership profile and disrupts the overall visualization of PCA. Therefore, we perform the PCA without him, but he, himself, forms a cluster whose profile is '*Completely supported to perform his task in terms of provided resource and expertise*'.

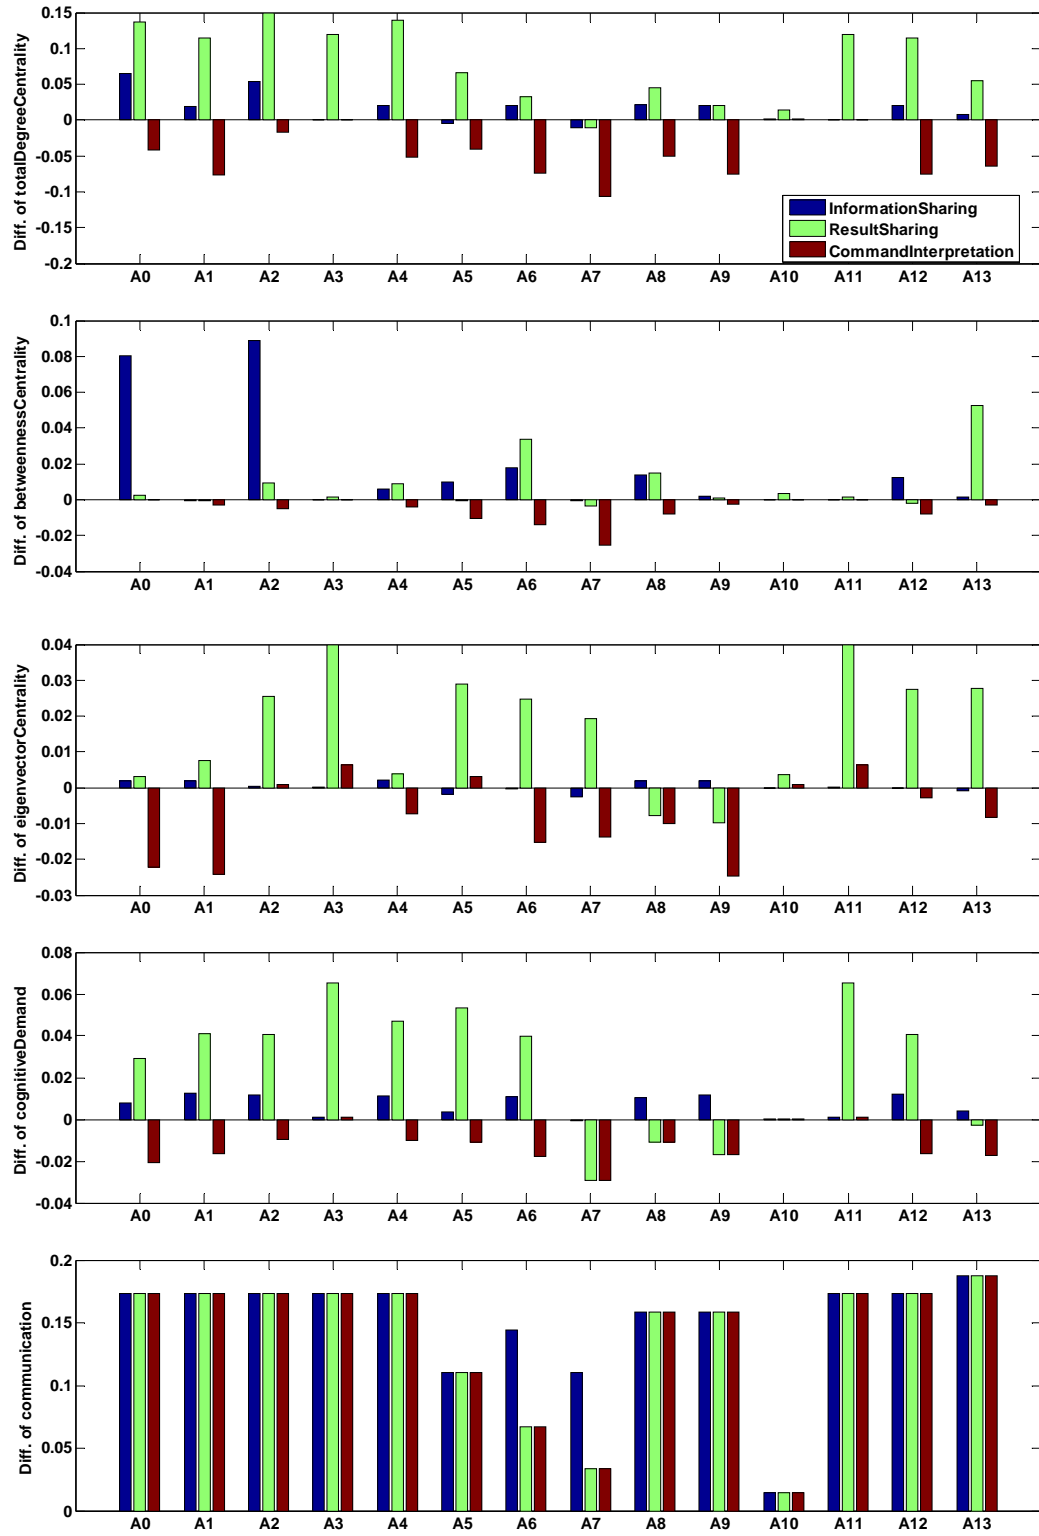


Fig. 8.9 Charts displaying the difference of metrics between a meta-network and extracted structures

According to Table 8.8, we have four sets of coefficients: two principal components for the original and the inferred. In the original, the high first principal component value implies *having more connections to other personnel, resources, and tasks* because it has high coefficients in degree centrality and cognitive demand. The high second principal component value means *high demand in communication to complete the assigned tasks* because it has high coefficient in communication. In the inferred structures, the meaning of the first principal component, *low demand in communication to complete the assigned tasks*, is similar to the opposite of the second principal component of the original, and that of the second component, *having more connections to other elements*, is similar to the first component in the original.

TABLE 8.8 Coefficients of two principal components from the original structure (top) and the extracted structures (bottom)

	Structure	Prin. Comp. 1	Prin. Comp. 2
Total Degree Centrality	ORI	0.6473	-0.4217
Betweenness Centrality	ORI	0.0920	-0.0309
Eigenvector Centrality	ORI	0.0513	-0.1513
Cognitive Demand	ORI	0.6088	-0.1763
Communication	ORI	0.4463	0.8759
	Structure	Prin. Comp. 1	Prin. Comp. 2
Total Degree Centrality	IS	0.1672	0.5702
	RS	-0.0848	0.3977
	CI	0.0571	0.2502
Betweenness Centrality	IS	0.0214	0.2144
	RS	0.0424	0.0664
	CI	-0.0018	0.0067
Eigenvector Centrality	IS	0.0103	0.1234
	RS	-0.0286	0.0297
	CI	-0.0241	0.0805
Cognitive Demand	IS	0.1114	0.4353
	RS	0.0282	0.2082
	CI	0.0783	0.3402
Communication	IS	-0.4370	0.1346
	RS	-0.6114	0.0850
	CI	-0.6114	0.0850

Figure 8.10 displays the clusters of individuals in the projection of the two principal components of the two structures. The original structure suggests five member profiles: *many connections to organizational elements and medium communication demand to complete their tasks* (A6, A8, A9); *medium connections and medium communication demand* (A0, A2, A1, A4, A12); *less connections and medium communication demand* (A5, A13); *less connections and high communication demand* (A3, A11); and *medium connections and low communication demand* (A7). The inferred structures provide four profiles: *medium or less connections and low communication demand* (A0, A1, A2, A4, A5, A7, A8, A10, A12, A13); *medium connections and medium communication demand* (A9); *high connections and medium communication demand* (A6); and *less connections and high communication demand* (A3, A11). These profiles tell the groups of indi-

viduals well supported in communication to complete their tasks and the groups, which are not. Also, it specifies the groups communicating frequently with other parts of the organizations and groups not communicating that frequently. *Al Zawahiri* and *Banshiri* were grouped in the same cluster in both structures. They were suffering from sparse communications to others and high communication needs to complete their tasks.




Fig. 8.10 Two projections of metrics of individuals using two principal components. The left is using only the original structure, and the right is from only the extracted structures.

8.6. Conclusion

This paper demonstrates what can be achieved by integrating social network analysis and decision making structure analysis. Social network analysis has been a prominent tool in investigating the structure of an organization. However, it is also susceptible to errors embedded in the given network structure. Therefore, reorganizing the links is required to perform analysis correctly. This reorganization is often done by human analysts. We expect to reduce such efforts by utilizing the introduced methods.

Furthermore, the method produces a set of different decision making structures that differ from each other in their natures. For instance, information sharing is a different relation compared to result sharing or command interpretation. When we only used a social network analysis, often the links are single-mode, meaning that the links are not differentiable. Therefore, the above method will enable analysts to think about the different types of links among the same entity types, and the analysts can reason more deeply by asking questions such as why these two agents have a command interpretation without any result sharing.

From the organizational structure perspective, a decision making structure and a meta-network are both network structures. Therefore, the analysis methods are interchangeable to some extent. For instance, we can apply social network metrics to both structures. This interoperability or interchangeability makes the analysis more comprehensive. For instance, we have different sets of critical personnel by analyzing various management relations and an original meta-

network. We are not certain which set contains the true personnel of interests, but we can suggest a package of results to human analysts.

Future work on this integration will include two major components. First, we should strengthen the decision making structure extraction heuristics. Currently, the information sharing extraction generates a dense network that is not common in the management science field. Also, we have a too sparse command interpretation that we believe are more in the organization. Therefore, we develop the existing method further or validate the current model by showing that the dense information sharing and the sparse command interpretation are legitimate. Second, we need to include more decision making structure oriented analysis methods in the framework. The result in this paper only comes from social network analysis, though it used the decision making structure for the analysis input. There are several decision making structure analysis methods, e.g. generating a set of feasible decision making structures under certain cultural constraints. In spite of these incomplete developments, this framework still shows its value by showing 1) the trimming process of a noisy meta-network, 2) different criticality analysis results from the extracted decision making structure, and 3) the opening of a unified organization analysis framework integrating social network analysis and decision making structure analysis.

Chapter 9

Simulating the Adversary: Agent-Based Dynamic-Network Modeling

Kathleen M. Carley, Geoff Morgan, Jesse St. Charles, Brian Hirshman

9.1 Social Network Simulation

Social network simulation (SNS) is an emergent area of research that combines social network analysis and simulation, typically agent-based simulation. This area is often referred to as dynamic network analysis as much of the focus of the combined modeling approach is on how networks evolve, change, and adapt. Additionally SNS has a focus on how individual and group learning and behavior is impacted by and impacts the changes in the networks in which the individuals are embedded. Frequently, in social network simulations, the social network and other networks, such as the knowledge network, and/or the individuals or “nodes” in the network are co-evolving as agents interact, learn, and engage in various activities. The need to address complex systems but produce realistic results means that these SNS are typically focusing on many types of networks simultaneously not just the social networks. An example of such a model might be one that explores how communicating new ideas via diverse social-media has differential impact on the movements of ideas and diseases through the population and response to the information and disease by the populace.

There are various types of social network simulation each has a unique perspective on the problem and each has its own collection of strengths and weaknesses. We begin with more formal approaches that rely heavily on statistics and mathematical formalisms and then move on to less formal bottom up approaches. System Dynamics is a top down, aggregate view of networks. Regression or econometric approaches like Quadratic Assignment Procedure provide a non-parametric approach to modeling dynamic social networks. More traditional parametric statistical approaches to SNS will use methods such as Expectation Maximization or Maximum-likelihood estimation to find the optimal (or near optimal) model parameters given the data. Finally, agent based SNS provides an intuitive bottom up approach for investigating social systems. Regardless of the method used for social network simulation, there are unique sets of challenges around validation, analysis, prediction, and computational efficiency that are common to all.

9.1.1 System Dynamics

System Dynamics supports top down reasoning about complex systems. Basic variables, system level mechanisms and the relations between them are modeled. System dynamics uses stocks, flows, and feedback loops to describe system behavior but because of its top-down, aggregate perspective it is less useful at the individual level. If we were studying information diffusion in a social network setting, a system dynamics approach might have a stock of people who have the knowledge, and stock of people who don’t have the knowledge, with knowledge flowing between them at some rate which is dependent on the percent of the population that already has the information, the density of the social network, and other graph-level network metrics. The approach is perhaps accurate in the aggregate but we lose the subtlety and nuance that explicitly

representing complex networks of people provides. For most social network simulation needs, the system dynamics approach is not the modeling framework of choice and is only used to talk about overall change in the structural parameters of networks such as the change in density but does not produce specific new networks of who is interacting with whom.

9.1.2 Statistical Network Generation

Both parametric and non-parametric statistical methods have been applied to learning and inferring models of social networks.

For parametric approaches: Random Graph Models provide a statistical, data driven mechanism for social network simulation. These models are derivative from graph theory or have been observed in real world networks. Networks are generated randomly using edge generation functions. Edges in social network do not exist or not exists; rather they have some probability of existence. This probability is modeled by a named parametric distribution like, Poisson, Exponential, or Power Law. Each of these distributions comes with a set of simplifying assumptions that may or may not be appropriate for the phenomenon being modeled. Optimal parameters for these edge models are empirically derived using expectation maximization or maximum likelihood estimation. However, due to the complexity of networks, the state space of these systems is massive making direct solving of the likelihood function intractable. This requires clever heuristic approximations to find near-optimal parameter values rather than the most optimal parameter values. Random graph models have been used to simulate collaboration and affiliation models. It is worth noting that these statistical approaches are typically aimed at the simulation of topological formation and have much less complexity and much less to say from a sociological viewpoint than their agent based counterparts.

For statistical models that have a stronger sociological basis we look toward P^* models, otherwise known as Exponential Random Graph Models (ERGM). These models are based upon Markov random graphs and represent a logistic regression of the network parameterized by various network statistics such as reciprocity, transitivity, centralization, connectedness and others. Using a pseudo-likelihood function, P^* models are fit to observed networks. This model can then be sampled to produce simulations of the observed social network. Tools like Sienna, developed by Tom Snijders, can be used to fit ERGMs to data.

There are also non-parametric approaches to social network simulation; one example is the Multi-Regression Quadratic Assignment Procedure, or MRQAP. MRQAP uses multiple samples from the social network being studied to perform a regression analysis of dyadic information that is correlated. Since properties of transience, reciprocity, and homophily are generally assumed to exist within social networks, most dyadic links have significant correlation with one another. This autocorrelation would normally be a significant issue for regression analysis but MRQAP uses a permutation procedure to account for the autocorrelation. MRQAP produces a regression model of social relationships that can be useful for running hypothesis tests on networks; this method can be significantly affected by bias learned in the model.

9.1.3 Agent Based Models

One of the most commonly used and intuitive approaches to SNS is Agent Based Models (ABM) (see *agent-based models*). ABMs employ a bottom up approach in which a set of heterogeneous agents, their behavioral properties, the “rules” of interaction, the environment and the interaction

topology that the agent populates is explicitly modeled. Complex social behavior emerges from simple individual level processes. In ABMs many computational entities, with varying levels of cognitive complexity, interact with one another in a manner similar to the real world entities they represent. These agents are simplified versions of their real life counterparts (e.g., ants, people, robots, or groups), only retaining elements salient to the phenomena being studied. Agents interact in a virtual world and can be constrained and enabled by the network position they occupy.

In most ABMs the topology of the virtual world is a simple 2-D grid and agents form “networks” as they occupy the same or neighboring spaces or the agent’s network is prescribed as the set of other agents within so many spaces of ego. Networks generated from grid-based interactions or defined in terms of grid-nearness tend not to have the same properties as true social networks; i.e., the distribution of ties, the method of tie formation and dissolution, and the relation of ties to physical space are not realistic. Most ABM toolkits support this type of grid-based modeling of the social topology.

There is, however, a growing interest in and a growing number of ABMs where the agents exist and move in a socio-demographic or network topology rather than a grid topology. An example here is the Construct model. In these models the agents occupy a social network position defined in terms of which other agents ego can interact with. In other words, rather than physical adjacency, social adjacency is used. This network topology may be static or dynamic. This latter type of model where agents exist in dynamic social networks rather than on grids is where most research on SNS is focusing. This is the approach we found to be most valuable for modeling the adversary and it is embodied in Construct.

9.1.4 Relational Sources of Complexity in ABM SNS

Social network simulation has a deliberate and expected preoccupation with relational information. The space in which people interact is a social one, there may be geographic motivations for communication, but these considerations merely temper and constrain the social space. As such, for the virtual spaces in the ABM SNS, agents interact in a social space where every agent is potentially adjacent to all others.

Using populations of virtual individuals, network effects emerge from both intended and unintended interaction among agents connected by ties of varying strength. The strength of the tie between two agents is defined in terms of frequency of interaction or strength of social tie or degree of similarity. Agent behavior, when the network is dynamic, can change the strength of this tie. Moreover, these ties may be hierarchically organized; e.g., two agents may interact and that interaction may be a work-based interaction and that interaction may be further characterized as interaction vis a specific task.

The mechanisms that drive interactions in an ABM SNS are typically based in social theory. Theories of human interaction such as homophily, transitivity, reciprocity are coupled with basic or sophisticated cognitive abilities. Theories of interaction take into account both social, emotional, and cognitive processes. Hence the cognitive load on the agent to determine whom to interact with when can be quite high increasing the demand both for more storage and more computational processing power.

Another source of complexity is the over-lapping social circles. Since the agents are socially embedded, the social environment itself may be characterized in multiple ways. For example, agents might be tied by different forms of similarity – age, gender, ethnicity, attendance at com-

mon events or co-location, shared resources, shared knowledge or beliefs, various role based relations – kinship, mentorship, leadership, and through diverse media – face-to-face, email, web. The result is that the agents occupy a multi-dimensional topology.

Finally, agents are not just embedded in social networks. Rather, they are connected in trails of who was where when doing what with what information or resources and for what effect. That is underlying any SNS that deals with socio-cognitive actors (hence not the simple statistical models) there will be an ecology of interlinked networks. This is referred to as the Meta-Network – a multi-mode, multi-link, multi-level network of networks at multiple points in time. For most ABM SNS the classes of nodes will include: Who (people, teams, organizations); What (tasks, events); How (knowledge, resources); Why (beliefs); and Where (locations). By formalizing these entities we are able to explicitly get at unique relationships between them implicit in multimodal data. A network of people-to-organizations is an affiliation network, while a network for knowledge-to-tasks is a Needs network, and agent-to-agent networks are the familiar social network formalization. A meta-network approach allows the developer to more fully represent and formalize relationships present in the real world that drive social interaction. For example, if a person is driven to interact with another person because they have to complete a particular task, and this particular task requires they know something specific but they don't have this knowledge then they have to go to a resource (book) or another person to gain the knowledge required.

These five key sources of complexity are completely connected network (full adjacency), hierarchical interaction, cognitive load, high dimensionality of and overlap in the social space, and meta-network considerations. These factors dramatically increase the complexity of social network simulation over many traditional agent based simulations. These factors also reduce the size of populations that can be simulated and increase the computational resources needed to simulate the system. ABMs have been used to model incredibly large populations, e.g., millions of agents. Parallelizing activity makes this possible. However, when accurate network representations are added as in the SNS models rather than just deriving the network from grid-based interactions, standard approaches to parallelization are no longer possible.

9.1.5 Common Research Challenges

Two core challenges are reuse and validation. Reuse is the process of taking an existing model and with no change to the internal processes reuse the model with different input data to address a new situation. An example would be to use a model of information diffusion to first explore how best to communicate medical information to effect change in smoking behavior and then reuse it to explore how to intervene in the social network to effect world leader's understanding of global climate change. Currently, most models are one-off model and require sufficient rebuilding and extension for new problems. SNS models, however, are a potential exception. In SNS models, these models can be built to take as input one or more real-world networks. The SNS models can then be used, on any network data set, to identify the probability of alternative futures and the impact of various interventions. A core advance in this area has been the development of support technologies to generate networks from socio-demographic data, such as census data, as import to ABM SNS (see e.g., the work on BioWar). For adversarial modeling we enabled reuse by augmenting Construct so that it could take meta-network data directly from ORA. This is described in the later modeling chapter and was used with the Indo-Pak scenario.

Validation of any socio-cultural simulation is difficult. The core reasons are that these models violate all the assumptions that underlie validation theory due to being comprised of agents

that learn. In the SNS area, the challenges are further compounded by the lack of spatio-temporal network data and by the fact that human lab experiments are inappropriate as network effects do not show up without groups greater than 5.

For ABMs the key validation approach is to do validation in parts and to validate each mechanism separately. The hope is that by validating the pieces, some confidence is bestowed to the whole. The problem, is little is known about the conditions under which this is true for a complex non-linear system.

Docking and model-to-model comparison is a key validation strategy. This process involves showing that for two or more models, common inputs produce common outputs. This allows a simulation that has not been formally validated to gain validation from an older simulation which has been validated and sheds light on the elements of the models that are robust. This model-to-model approach is part of the multi-modeling approach used in this MURI.

For the statistical models like ERGM, P* and MRQAP the models have been “trained” on real data. In this case, validation is the process of seeing whether the predictions hold in the future or in other time periods. For these models generalizability is more of a concern if the models learned are ported to other reasonably equivalent systems.

9.1.6 Applications

Common uses of SNS range from theoretical investigation to applied analysis and prediction tasks. Researchers can explore the ramifications of sociological principles like homophily and transitivity: are these mechanisms sufficient to produce real networks that we observe? What are the properties of analysis methods and are they robust? In applied settings SNS can assist in predicting how a network will evolve and help analyze the dynamic equilibriums that might arise. Key application areas are the spread of disease, information diffusion, belief formation and diffusion, and activity contagion. SNS are critical for understanding the impact of various interventions where social influence is expected to play a role. Here we use them to assess adversarial groups. In particular we used agent-based dynamic-network models, specifically Construct.

9.2 Agent-Based Dynamic-Network Models

Agent-based modeling is a simulation technique which relies on the capabilities of individual actors, called agents, in order to model a global behavior. In an agent-based model (ABMs) complex system level behavior emerges from the local action of, and interaction among, a large number of heterogeneous agents. The relationships between agents, the social and spatial topology in which agents are embedded, and the logic that guides agent behavior play a crucial role in determining the overall behavior of the system. Global outcomes emerge as heterogeneous agents interact and engage in various local activities.

There are a number of advantages to investigating a research problem by building or extending an agent-based model. All simulation techniques, including agent-based modeling, are key tools for theory development as they force researchers to encode their assumptions when writing models and to question previously hidden assumptions in theories. This process allows a researcher or policymaker to realize the limitations of a particular theory or solution, or conversely to develop extensions of a theory into a new domain or to develop a solution that is more robust. When building an agent-based model, the simulation designer will have full control over what types of data will be gathered, and can be modified relatively easily if follow-up virtual experi-

ments are performed. The data gathered will not be subject to the kinds of cognitive or methodological biases found in empirical research. Virtual experiments performed using agent-based models may be more ethical than those using people, especially if the experiment requires radical or harmful reorganization of the actors involved. The size of agent-based virtual experiments can also be much larger than those performed using traditional human subjects, and the marginal cost of adding an extra actor or even an entire replication can be trivial. Simulation can also be used to examine the same starting condition multiple times, allowing the researcher to perform a ‘what-if’ analysis as random changes build up and cause the simulated population to evolve differently. Simulation can also be used predicatively in order to forecast what would happen to a specific initial condition; when run multiple times, broad trends may be detected and outlying cases and their causes potentially identified. Lastly, by leading the researcher to think about the kinds of local rules that lead to global patterns or by forcing the researcher to confront the unintended consequences of seemingly individual rules, agent-based modeling can help a researcher understand the link between individual and social behavior.

Agent-based modeling, like other technique, has its strengths and limitations. These models are particularly valuable for comparing, contrasting and combining theories about how individuals act and so serve as a virtual world for developing theory by both exploring theory interactions as well as generating and testing hypotheses. They are valuable when there are not strong empirical regularities relating the past to the future as they allow discovery of the space of possibilities. ABMs are tools for gaining intuition about how individual differences can have systemic global consequences. Finally, ABMs enable experimental protocols to be examined and the likely consequences estimated using virtual experiments when the same experiment is too complex, costly, technologically infeasible or unethical to run in the real-world. Agent-based modeling also has a number of weaknesses. These models often have a vast number of parameters and so must be run a large number of times in order to appropriately explore the parameter space. This can create analytic difficulties. Validation, as will be discussed, may be difficult. Many ABMs are built with rules specific to a narrow domain and so have to be significantly rebuilt to be used in a different domain. ABMs can require vast quantities of computational resources, particularly if very high fidelity agents are used.

ABMs are distinct from other mathematical or modeling techniques such as closed-form solutions, discrete event simulations, and system dynamics models. While ABMs are agent focused, the other techniques are population focused. Closed-form solutions are mathematical transformations which attempt to find an exact (and optimal) solution to a particular problem when expressed mathematically; while such solutions may be found for certain simple problems, they are often not applicable for the complex and often inexact problems that agent-based models are used to address. In contrast ABMs are concerned with the process and not on some optimal or final state. Discrete-event simulations focus the design of the model around events, usually organizing the simulation around an event queue; however, these events need not be generated by actors themselves. ABMs can take as input event sequences but add individual rules of behavior to respond to such events. System dynamics models focus on aggregate behaviors in a society, and as such attempt to express the number of agents who have a particular trait without completely specifying the agents themselves. Both system dynamic models and ABMs are complex system models. The key difference is that the logic for social change is that system dynamic models are top-down whereas ABMs are bottom up. From an environment perspective in an ABM, the environment, such as the social network, is represented explicitly; in contrast, the other models represent the environment using summary statistics such as density. As a result, only

an ABM can explore the explicit flow of ideas, beliefs, influence, trust, disease, money etc. though the network as agents interact and get as output the specific network and which agent has what when.

9.2.1 Agents and Their Environment (and Social Network)

ABMs vary in how the environment is represented. This could be as simple as a single dimension or array and so ego interacts with those other agents that are within so many squares left or right of ego. This is the case in Kaufman's NK model. Traditionally, however, the environment was a grid and the agents interacted with other agents in and/or could move to those squares that surrounded them. Most early studies explored the relative impact of von Neuman (squares left, right, up, down of ego) or Moore (eight squares around ego) or extended Moore neighborhoods (squares within some distance of ego). In these traditional approaches the structure of the social network is directly tied to the physical position of the agents. Examples of such models are the game of life, the original Schelling segregation model and the more recent SugarScape models developed by Epstein and Axtell. In general, it is difficult to get realistic social networks in this representation of the environment. Further, as early results showed, unless the grid is bent into a torus, the resultant social behavior is largely dictated by "edge effects"; i.e., restrictions on activity caused by being at the edge of the physical grid.

More advanced models place agents in a socio-demographic space and separate the physical and the social space. In such models, very few have explicitly modeled the social network. Increasingly, however, researchers are incorporating more realistic network representations, such as small-world, scale-free, or other types of network generators. The most advanced of these models are the dynamic-network ABMs in which the networks and the agents co-evolve (the first model of this type was Construct). In some cases, the models are instantiated with networks that are actually derived from real data. These models will often generate or import an appropriate graph before the simulation agents are initialized, and then assign each agent to a graph position when the simulation starts. Other models use a social network gathered from empirical studies. These networks have the advantage of being as realistic as possible, but may potentially bias the simulation results due to the structure and nature of the particular social network gathered (see *social network simulation*). Correctly specifying the topology of a social network in an agent-based model has important implications for the conclusions drawn. In modeling the adversary it is valuable to use the social network of the adversarial group.

The quality of the social network modeling can have important effects on simulation outcomes. For instance, in the Construct model, the social network topology has a non-linear effect on knowledge and belief diffusion rates in the system. Construct uses sophisticated agents that have the ability to interact and choose partners with which to exchange knowledge and belief. A stylized meta-network, which specifies the pattern of potential partners with which an agent can interact, can be imposed to limit the form of the evolved networks. We use Construct to model the adversary. Our results indicate that the most effective type of intervention depends on how the adversary is structured; e.g., Al Qaeda and Hamas have different structures and the same intervention, such as isolation of the top leader, in the two cases can lead to performance decrements in one and performance improvements in the other.

9.2.2 Trade-offs

When building an ABM, particularly an agent-based dynamic-network model, researchers should be aware of the key trade-offs. One important trade off is between simplicity and realism. Simple

models, such as Schelling's segregation model, attempt to use a specific principle to describe an important trend in human or social behavior. By keeping the principle narrow, the modeler seeks to illustrate how a particular phenomenon has important explanatory power. Such models are extremely valuable for engaging systematic thinking in an area and for making key points to an audience. However, the results generated by using such models, however, tend to be quite fragile and can change radically as new types of agents, alternative environments, or additional interaction logics are added. In contrast, more expressive models are more veridical and by capturing greater realism are capable of explaining a wider swath of human socio-cultural behavior. The more expressive emulative models often employ multiple modules, as well as an extremely large number of parameters, in order to increase their accuracy and predictive power. This increase in power and fidelity, however, comes at the cost of ease of explanation, time to generate results, and time to analyze model results. Another important trade-off is between the sophistication of each agent and the number of agents in the model. In general, the more sophisticated the cognitive model the fewer the number of agents represented. Models with a larger number of agents typically employ simpler agents with fewer rules such as in the artificial life simulations; in contrast, models with only a few agents typically employ quite sophisticated cognitive agents capable of actually doing tasks, e.g., flying planes, as in tac-air-SOAR. The reason is simple: processing and run-time constraints are such that increasing either the number or the cognitive sophistication of the agents increases computational costs. There are two source of complexity in the agent model: cognitive and social. The more sophisticated the cognitive model the more the ABM can be used to explore behavior on specific tasks, such as buying groceries. The more sophisticated the agent's social model, the more it can be used to address issues of socio-cultural change and information diffusion. Both cognitive and social complexity increase computational processing costs. Historically, ABM designers with more emulative models have either worked with a few (less than a hundred) very realistic cognitive agents, or a moderate number (less than 25,000) of very realistic social agents that are moderately cognitively realistic, or millions of agents that are both cognitively and socially simplistic. A further trade-off occurs between overall model sophistication and speed. Not only will more complex models take longer to run due to their more complex computer logic, but they will also take a substantial amount of time to code, debug, and process results. Simple models, on the other hand, will run faster but may be more limited in their output. As a result, the more sophisticated the model, the more likely it is built, maintained, and extended by a team whereas the simple model may be built by a single research.

9.2.3 Validation and Verification

Validation, or the alignment between the model's behavior and actual empirical data, is a major concern and is a criticism often levied against simulation models of socio-cultural systems. Though models are often criticized for insufficient validation, the type, scope, extent, and precision of validation depends on the data available, the type of model built, and the expected use of the model's predictions. More validation is not always better; extremely basic models are rarely validated as their purpose is illustration, and some models need not be validated at all. On the other hand, emulative models rarely can be validated using a single case scenario and consequently the researcher needs to fuse data from a wide variety of sources – often at different time scales and collected for diverse purposes – to obtain a “good enough” dataset for validation.

ABMs of socio-cultural systems present special challenges to validation and analysis. One cannot naively assume that if the basic model of a single agent is validated then the aggregate model is valid, as interaction effects may lead to very different behavior. Typically models are

validated at either the individual agent or the collective level but not both. When compared to engineering models of physical systems, these socio-cultural ABMs have more variables, high covariance among variables, discontinuities in variables, and non-stationary processes, interaction effects and temporal variations in the relations among variables often due to learning. As such, the nature of socio-cultural ABMs violates basic assumptions about the nature of simulation models that underlie the traditional formal approaches to analysis and validation developed in engineering and the physical sciences. This means that a new science of validation is needed, and that socio-cultural ABMs should not be used to predict the future but to describe the space of future possibilities. Given these complexities, new approaches to validation in this area have emerged: validation by parts (validating individual sub-modules), validation of inputs, validation of processes, and validation by docking. The process of docking two models, whereby the results of one model are compared to that of another, enables a greater understanding of what factors make the model results robust and identifies common failings. Moreover, common results from divergent models, as we found in the Indo-Pak scenario enhance the likelihood of the overall finding.

9.3 Construct

For modeling the adversary we extended and used the Construct model. Key extensions were enabling reusability by allowing the model to be instantiated by ORA, geo-spatial diffusion modeling, and multi-intervention analysis. Over the course of the MURI a large number of studies were done. These included: analysis of basic adversarial forms, impact of alternative COA related to both agent isolation and knowledge promulgation, and assessment of best approach for effecting belief change, i.e., “winning the minds and hearts of the adversary.”

Construct 3.5 – hereafter referred to as simply Construct – is a multi-agent dynamic-network simulation model for examining the co-evolution of agents and the socio-cultural environment [109], [188]. Using Construct, one can examine the evolution of networks and the processes by which information moves around a social network [189], [190]. Construct captures dynamic behaviors in groups, organizations and populations with different cultural and technological configurations [191]. In Construct, groups and organizations are complex systems. The variability of human, technological and organizational factors among such systems are captured through heterogeneity in information processing capabilities, knowledge, and resources. Multiple nonlinearities in the system generate complex temporal behavior on the part of the agents.

Construct is the embodiment of constructivism, a mega-theory which states that the socio-cultural environment is continually being constructed and reconstructed through individual cycles of action, adaptation and motivation. Many social science theories and findings are part of the constructivist theoretical approach including structuration theory [192], social information processing theory [193], symbolic interactionism [194], [195], social influence theory [196], cognitive dissonance [197], and social comparison [198]. In addition a number of cognitive processes are embedded such as transactive memory [199].

There are three key features of Construct 3.5 that make it ideal for our purposes. First, the experiment designer has complete control over which sub-agent models are used for interaction over the course of a run. Second, Construct contains a suite of agent models which enable diverse socio-technical conditions to be modeled. Third, general agent characteristics can be easily configured *a priori* using empirical data or they can be based on hypothetical data. To use Con-

struct, as we did in this work, the researcher specifies both the relevant agents [200] and the social and knowledge networks [201].

While additional information about the Construct interaction model can be found elsewhere (e.g., [109], [190]), the core Construct agent dynamics are as follows.

Each time period of the simulation, agents take a variety of actions including initiating an interaction, responding, sending messages, engaging in tasks, updating beliefs. For each agent, these action tend to occur cyclically except for responding to media and making decisions which may occur off cycle – see Fig. 9.1. Exactly which actions an agent can take, and how many can be done simultaneously, depends on the agent’s socio-cognitive nature. Each action takes a certain amount of time, typically a time period. Human agents use their preference for homophily or expertise, their transactive memory of other agents’ knowledge, their beliefs, their socio-demographic characteristics, their availability, and their recommendations from others in order to rank the importance of interacting other agents in their social network. Based on this ranking, the human agent may choose to initiate communication with one or more other agents. The type of agent – human, web-page, etc – will determine whether the agent can initiate interaction and what are the agent’s information processing characteristics.

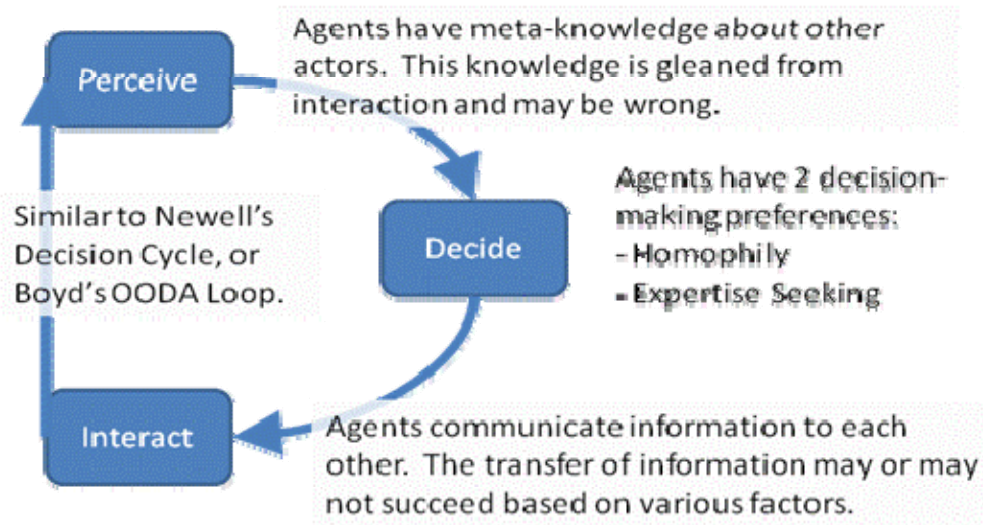


Fig. 9.1 Cycle of Agent Activity

If two agents are able to interact and communicate, then both agents will prepare a message and send it to the other. A message is a set of memes [202], and so consists of one or more instance of any or all of the following: knowledge, beliefs, transactive memory about the knowledge of third parties, or transactive memory about the beliefs of third parties. Once prepared, the message is communicated to the interaction partner, where it may be modified, misinterpreted, or ignored based on the socio-cognitive properties of the receiver. After receiving a message, processing it, and possibly learning from it, both parties may modify their beliefs or make any relevant decisions. This process then repeats for each agent during each time period.

All agents operate in the same “time frame” meaning that interventions and/or interrupts can occur at a particular time and all agents can respond to it – e.g., a news add can come out at time

period 3, and agents will respond during that time period and other periods when the interrupt or intervention is active. Statistics, outputs, and decision information are gathered relative to these the time periods, as well as at the end of the simulation.

Although Construct was originally developed as a pure lock-stepped model with each agent interacting each time period and then updating their memory, that is no longer the case. As of version 2.5, Construct includes event driven mechanisms, variable duration interaction, and fixed as well as mutable agent characteristics. In addition to these programming changes to update Construct with newer simulation technology, additional work has performed done on validating the core mechanisms independent of the exact technological mechanism in a variety of settings. Finally, the current version is multi-threaded.

The fundamental mechanisms in Construct have been scientifically validated [188], [203], [204], [205]. It has been used to explain group mobilization [188], the impact of leadership [206], and the impact of the printing press [207]. Directly germane to the current study, Construct has been used to compare and contrast different educational media by socio-demographic feature [208] and the impact media and opinion leaders in real cities [209].

9.3.1 Agents

Agents are decision-makers with varying information processing, socio-demographic, and access constraints and as such may or may not be human [152]. Within Construct, agents go about their business interacting, communicating and learning each time period, as described in Fig. 9.1. As agents learn or acquire information, they may change their preferred interaction partners and modify what they are likely to communicate. These factors, in turn, influence what types of decisions are made by each agent. A variety of factors influence who agents select as interaction partners, what they communicate with that partner, how much and how they communicate, whether they learn anything from that partner, and the accuracy and sustainability of that learning. Such factors include the agent's socio-demographic characteristics, information processing characteristics, proximity, and current position in the social and knowledge networks. The agent model has been described in depth in other venues (e.g., [188], [200], [201]; thus, we concentrate here on both a high level description and details of those components used for the simulations reported.

Within Construct, agents both influence and are influenced by others. Agents who have influence over others can use that influence to escalate or de-escalate activity at a societal level by communicating information and/or beliefs. Social influence – as derives from shared attributes such as socio-demographic factors, shared knowledge, beliefs, and proximity – co-evolves with the spread of knowledge and beliefs [109]. Consequently, in more heterogeneous populations where the lines of differentiation line up the chance of self-reinforcing beliefs at the group level is greater [210]. Factors that are not influenced by the diffusion of information and beliefs include the agent's socio-demographic role (e.g., age, race, gender, level of education), the agent's basic cognitive limitations and information processing capabilities (e.g., likelihood of forgetting, risk taking, amount of information and beliefs that can be communicated or processed, and whether the agent has transactive memory), the size of their sphere of influence (at least in the short term), and factors that have resulted from socio-cognitive interactions (e.g., literacy, access to newspapers, radio and the internet).

Within Construct, agents develop likelihoods of interacting with others based on relative similarity (RS) and relative expertise (RE) [211], [200]. Relative similarity is a homophily based

mechanism [212], [109] and derives from the idea that individuals are more likely to interact if they have more in common. Homophily based interaction is a multi-causal phenomenon due to ease of communication, shared understandings, and comfort. The relative similarity of i and j , from i 's perspective, is characterized as

$$RS_{ij} = \frac{\sum_{k < K} (AK_{ik} * AK_{jk})}{\sum_{j < I} \sum_{k < K} (AK_{ik} * AK_{jk})}$$

where individual i 's relative similarity to j , is determined in terms of socio-demographics, knowledge, and belief items K in the agent-to-knowledge matrix AK .

Of important note: an individual is most relatively similar to itself, and each period will have a reasonably high probability of choosing to “interact with itself” and to avoid communicating with others. Just because an agent has the highest relative similarity with itself, however, does not mean that an agent will always interact with itself; indeed, due to the large number of other agents in the simulation, such avoidance of communication is relatively rare.

Relative expertise is a search based mechanism and derives from the idea that individuals are more likely to interact if one has information that the other wants. The relative expertise of j as judged by i is characterized as

$$\text{if } AK_{ik} = 0, \text{ then } X_{jk} = AK_{jk} \text{ else } X_{jk} = 0$$

$$RE_{ij} = \frac{\sum_{k < K} X_{jk}}{\sum_{j < I} \sum_{k < K} X_{jk}}$$

where individual i 's relative similarity to j , is determined in terms of socio-demographics, knowledge, and belief items K in the agent-to-knowledge matrix AK [213].

Agents are more likely to initiate interaction with another if they think the other has information they need and/or they are similar to them. However, there is a curvilinear relation between this familiarity and expertise; to wit, as agents initially increase in similarity (homophily) they are more likely to realize the other has expertise they need but as they increase still further in similarity they realize that the other is so similar there is no specialized expertise.

The researcher needs to specify the strength of each of these factors for agent-agent interaction. Herein, we set all human agents to use both logics and to at any time create a combined probability of interaction that is based on 60% similarity and 40% expertise. In both cases, individuals are giving and receiving information and the overall tendency to give versus receive is about 60/40 as identified by Valente, Poppe and Merritt [214].

When setting up a virtual experiment in Construct the researcher needs to specify multiple parameters for each agent. This is often facilitated by the use of agent classes to parameterize multiple agents simultaneously. Specifically, the agent needs to specify the number of agents in each of the classes of agents in a virtual experiment, the distribution of socio-demographic parameters for the agents of that class, the distribution of cognitive factors for each class, the sphere of influence for that class, and the access constraints for that class. There are many other factors that can be varied, such as the rate of forgetting. However, we have found that for modeling the adversary the items listed are the core variables that need to be defined.

9.3.2 Agent Classes

In this study, we find it helpful to think in terms of two meta-classes of agents – human agents and media agents (which may or may not be human). Each time period, human agents may interact with other members of the general human population or with a media agent. Sometimes, it is useful to further break the general public into subgroups such as red, green and blue, or terrorists, harboring population and US forces.

There are two classes of human agents: the general public, who will make decisions, and the opinion leader, who can help sway decisions. In the experiments performed, the opinion leader attempts to get the general public to act in one way while the media are designed to thwart it.

In this experiment, we consider five classes of media agents: newspaper advertisements (ad), publically accessible web sites (web), centers that have people in them that provide assistance when someone comes in physically or calls in via phone (call), radio advertisements (radio), and letters sent via postal mail (mail). Media agents differ from each other in terms of the time periods they are active and the length of the messages they send. All agents can communicate facts or beliefs, but the particular set transmitted depends on their knowledge or belief at the time. All media agents are passive – they cannot initiate communication with a human agent. Instead, they provide information only when the human agent selects to go to, listen to, or read the information available through the media.

These particular media agents were chosen because they represent distinct forms of access to information. You might ask why we did not use television when it is so prevalent. The reason is that, within the characteristics we were varying television and radio ads are identical. Thus one can think of radio as radio/television ads.

The number of each type of agent, their activity level and length of messages sent needs to be defined. Table 9.1 provides an example.

The initial knowledge and beliefs held by each of the media agents and the general public at the beginning of the simulation need to be defined. See Table 9.2 for an example. Note the user can specify one or more beliefs and for each define the distribution of knowledge. Over the course of the simulation, the general public, i.e., the agents representing humans, learns; however, the knowledge of the opinion leader and the media remains constant. The number of facts in each category, specified in the left-hand column of the table, is proportioned based on subject-matter expert's views of the relative amount of time it takes for the overall meta-concept – such as know-how for a task – to diffuse.

In order for a human agent to make a decision, an agent must recognize that the activity exists, must have sufficient know-how knowledge, and hold a positive view of one of the two beliefs. In order to have sufficient know-how information, agents must learn at least three of the six know how facts; considering that agents do not start with any of this information, they must learn it from, ultimately, the opinion leader or media. Additionally, we have modeled two beliefs here – one where the true belief is that one shouldn't engage in the activity (believe not right), and one that is neutral as to whether there is some benefit to engaging in the activity (believe worth doing). In order to make the decision, agents must hold at least as many positive beliefs as negative beliefs, or they must be subject to social influence from their peers which convinces them that the decision is a good one. Each activity modeled would have a set of beliefs associated with it.

TABLE 9.1 A table illustrating how a user can characterize different classes of agents by specifying their number, activity, and message capabilities.

Class	Number	Periodicity	Active Time Periods	Message Length
Humans	3000	Continuous	104	1 fact, belief, transactive fact or belief, or social information
Opinion Leader	1	Periodic every other time period	52	1 fact, belief, or transactive fact or belief, or social information
Ad	1	New news ads are periodic	26	1-2 facts or beliefs
Web	1	Periodic access every other time	52	4 facts or beliefs
Call	1	Periodic access every fourth time	26	3 facts or beliefs
Radio	1	New ads are periodic	1 (per ad)	1-2 facts or beliefs
Mail	26	Periodic new mail	6	3 facts or beliefs

TABLE 9.2 A table illustrating how a user can characterize a population by differentially distributing information and beliefs across classes of agents.

Information and Beliefs	General Human Population	Opinion Leader	Media Agents				
			Ad	Web	Call	Radio	Mail
Activity exists (1 fact)	0%	100%	100%	100%	100%	100%	100%
Activity know-how (6 facts)	0%	100%	10%	33%	10%	10%	10%
Believe right (3 facts)	1%	100%	0%	0%	0%	0%	0%
Believe not right (4 facts)	5%	0%	33%	100%	100%	33%	33%
Believe worth doing (3 facts)	1%	100%	0%	0%	0%	0%	0%
Believe not worth doing (3 facts)	5%	0%	33%	100%	100%	33%	33%
General knowledge (500 facts)	20%	20%	10%	2%	5%	10%	10%

The more facts per category – and hence the more complex the message – the longer it takes that category as a meta-concept to diffuse. However, it is important to note that all information related to the activity (twenty total facts) is small relative to the amount of simulated general

knowledge (five hundred facts) so that most of the time the general population will not be communicating facts about the activity. Furthermore, the ratio of positive and negative facts associated with the belief influences whether the “correct” belief is positive, negative or neutral. We have two beliefs here – one where the true belief is that one shouldn’t engage in the activity, and one that is neutral as to whether there is some benefit to engaging in the activity. Finally, the amount of information associated with activity know-how and with any one belief is comparable so that both spread in a comparable amount of time. The more complex the know-how and the more complex the belief the longer that information will take to spread and the lower the fraction of the population that will have the expertise or belief at any one time.

The advertisement is meant to provide a small amount of knowledge and belief while also containing a large amount of general knowledge information to encourage agents to examine it. Such behavior is typical of articles or advertisements in newspapers. Advertisements only exist for a few time periods during the simulation, reflecting relative infrequent publication. Advertisements can be expected to have a small impact on a variety of agents due to the infrequent interaction and small message conveyed; however, they will be among the most common media that human agents access. Since the advertisement is in printed media, it can be subject to two different constraints: a cognitive constraint, literacy, and an access constraint, subscription access.

In contrast to the advertisement, the web site is designed to provide a large amount of belief information by proving a large number of reasons why the activity is inappropriate. In doing so, however, it could potentially be scraped for knowledge information, thus serving a purpose that is contrary to what the designers intended. For this reason, resources such as the website can be two-edged swords. Because the website is frequently available, it will be easily accessed; however, users accessing it may have literacy or internet access issues.

The information-call-center is designed to answer questions about the activity, based on requests for information from those members of the general population who contact the center. Because the information-center represents the actions of humans who work at the center it has associated with it more social knowledge than the web site. Unlike the web-site, though it may be difficult to get to the center as it requires physical movement and thus may not be as favorable an interaction partner to some agents.

The radio advertisement is very similar the print advertisement. It is designed to provide a small amount of information or beliefs but can reach a large number of agents in the general population. Unlike the advertisement, however, the radio advertisement is not affected by the literacy or access constraints as modeled in this experiment.

The postal mailing is designed to represent a piece of mail containing information meant to deter at-risk agents from engaging in the activity in question. It, too, has the same information content as the advertisement, but the way the general population interacts with it is unique. Only some “human” agents receive mail. However, whether or not the “human” agent reads the mail is up to the individual agent. For the next six time periods, the mail message resides in the agent's “mailbox”. The general population agent then has a certain probability of checking their mail and learning the information in the letter. Agents who read the letter absorb some of the information contained in the letter.

To date, a large number of media have been modeled. In general, we tended to model media known to be used by the adversary and/or US forces.

9.3.3 Agent Socio-Demographics

In Construct, agents can have a set of non-evolving attributes that influence behavior. Herein, we consider those attributes to be socio-demographic characteristics. These attributes can be set based on census data, or based on other considerations. The researcher can in fact define any characteristics as agent attributes and then use these to effect interaction. The critical difference between attributes and knowledge/beliefs is that for an agent the attributes are fixed for the duration of the simulation; in contrast, the agent's knowledge and beliefs may change. Consequently, attribute based interaction tends to be stable, and variations in interaction are due to changes in knowledge and beliefs.

The user can define any attributes that make sense within the socio-cultural context being modeled. One possibility is to base this off of population demographics. The socio-demographic attributes are used to set the baseline interaction that exists independently of agent knowledge. The greater the overlap in agent socio-demographic attributes, the more likely the agents will interact, as part of the homophily effect [215]. Table 9.3 shows an example of attribute setting using a set of ubiquitous general aspects of human behavior.

Two classes of agents, general public human agents and opinion leaders, have these socio-demographic attributes. Media agents could be “targeted” so that they were aimed to “match” and so interact with humans with different attributes. Specifically, media were designed to target the agents who had either the lowest or second-lowest level of income and education. Thus, the opinion leaders and media would match any agent who had one of those two attribute values, but would not match any other agent. These attributes were oversampled in the human agent population relative to the general population of the United States in order to better understand the effects of the cognitive limitations, and information processing capabilities.

TABLE 9.3: A table illustrating how the user can differentiate agents by varying the socio-demographics.

Attribute	Number of Values	Values (% of Human Agent Population)
Age	5	0-29 (20%); 30-39 (20%); 40-49 (20%), 50-64 (20%), 65+ (20%)
Gender	2	Male (50%); Female (50%)
Race	5	White (60%); African-American (15%); Hispanic (10%); Asian (10%); Other (5%)
Income	6	0-15k (40%); 15k-30k (35%); 30k-50k (15%); 50k-80k (6%); 80k-120k (3%); 120k+ (1%)
Parent	2	Yes (50%); No (50%)
Education	4	Less than high school (40%); High school diploma (35%); College degree (30%); Graduate schooling (1%)

The correlation between these attributes is also an important consideration. Population level correlations could have been generated in one of three ways: 1) proportional to census data, 2) randomly, or 3) evenly. Results can vary dramatically with the socio-demographic distribution.

9.3.4 Agent Cognitive Limitations and Information Processing Capabilities

Construct agents are complex. Two core features of Construct agents are information processing capabilities and transactive memory. Agents are information processing decision makers and so have one or more of these capabilities: initiate interaction, send messages, receive messages, learn from messages. Agents can have both general and transactive memory [199]. An agent's general memory can contain both what information the agent knows – its facts – and what beliefs it agent holds. The transactive memory, on the other hand, contains the agent's understanding of third parties – who knows what and believes what. This transactive memory can be incorrect: the third party might not know the knowledge or hold the belief. Who knows what, as well as who knows who knows what, can be tracked by time period.

Agents make decisions as to whether or not to engage in activities based on their current knowledge and beliefs. These decisions require various information and beliefs, such as: information on how to do the activity, a belief that the agent should do the activity, and a belief that the activity is appropriate. The point here is that there is a mask on information and beliefs such that different information and beliefs are needed for different decisions. In addition, for this study, decisions are made at the final time period based on accumulated information and beliefs.

However, agents differ in their information processing constraints. In this study we use the following factors: amount of information and beliefs that can be communicated or processed at a time, whether the agent has transactive memory, and whether the agent can initiate interaction or learn. These factors are set differently for each agent class. In Table 9.4 illustrative distribution of cognitive capabilities per agent class are described. In Table 9.5 illustrative distribution of the information processing capabilities per agent class are described. Other possible factors that we can consider in the future are forgetting and risk-taking.

TABLE 9.4 A table illustrating how the user can differentiate agents based on constraints.

Factors	General Population Human	Opinion Leader	Media Agents				
			Ad	Web	Call	Radio	Mail
Access Constraints	Literacy, Web, Newspaper	None	None	None	None	None	None
Number of messages received and processed at the same time	1	None	None	None	None	None	None
Number of messages sent at same time	1	unlimited	unlimited	unlimited	unlimited	unlimited	unlimited

9.3.5 Networks

Construct is a multi-agent dynamic-network simulation system in which the agents are constrained and enabled by their position in a meta-network. A meta-network defines the set of relations among who, what, how why through a set of geo-temporal trails [152], [216]. As such, a meta-network is a multi-mode, multiplex, multi-level network. Consequently, in Construct,

agents are embedded in a large number of networks, including formal and informal relations among agents, relationships between agents and knowledge, and assignments of knowledge and beliefs to tasks. From a meta-network perspective the key entity classes in Construct are agents, knowledge or expertise, beliefs, and tasks. Thus the core networks are the social network among agents, the knowledge network (agents to knowledge), the beliefs network (agents to belief), the assignment network (agents to tasks), and the requirements network (knowledge + beliefs to tasks). Within this the social network can be further broken down in to a proximity based network, a socio-demographic based network and a knowledge/belief based network.

TABLE 9.5 A table illustrating how to define agent classes by varying the information processing capabilities of the agents in that class.

Factors	General Population Human	Opinion Leader	Media Agents				
			Ad	Web	Call	Radio	Mail
Initiate	Yes	Yes	No	No	No	No	No
Send	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Receive	Yes	Yes	No	No	No	No	No
Decide to take action	Yes	No	No	No	No	No	No
Learn	Yes	No	No	No	No	No	No
Change beliefs	Yes	No	No	No	No	No	No
Information Atrophy	No	No	No	No	No	No	Yes
Message Complexity	Very Low	Very Low	Low	High	Med	Low	Med
Supports multiple searches	Yes	Yes	Yes	Yes	No	No	Few

Additionally, Construct allows the experiment designer to set networks as fixed or dynamic during the simulation. Moreover, the initial topology of such networks can be specified. This enables the impact of topology to be studied at the same time as the impact of information-based media. In this study, parts of the social network are fixed based on demographics and parts dynamic based on changing expertise. The result is that the overall probability of interaction between each dyad is dynamic. In this paper, both knowledge networks and belief networks, were dynamic and would change over the course of the simulation.

The social network is of particular interest to this study. Empirical studies of social networks often form a network by asking individuals for the names of their interaction partners. The result is a snapshot of a network at a point in time as it is perceived. Based on this perspective it is tempting to think of networks as simple binary relations, two individuals either are or are not connected. Simulation makes it obvious that the idea of a network is more amorphous.

In Construct there are a number of ways to characterize the network of possible agent-agent interactions. All agents exist in a social network, and in this network the links among agents are probabilistic. These probabilities evolve over time, changing as agents increase in similarity and expertise. At any point in time, who is interacting with whom can be extracted in multiple ways:

as a moving average, as probabilities, as a number of interactions in one particular time period , and as whether an interaction occurred from the beginning of the simulation to that point.

In designing the simulation, the sphere of influence – the alters with which an ego’s probability of interaction is nonzero – is the set of others who the agent is likely to interact with. In the full Construct model this sphere can grow and shrink; however, in this study we leave it fixed. Agents with greater reach – such as the opinion leader – have a larger sphere of influence, while most human agents have a relatively small one. Constraints on information access, as will be described, can impact the effective size of an agent’s sphere of influence by making certain types of agents inaccessible. The size of the sphere of influence per agent class is described in Table 6; both the theoretical maximum is provided and the “in-practice” value determined from the experiments run. Note, the opinion leader is in every general human agent’s sphere of influence. Additionally, for each human agent whether or not a media agent is in the human agent’s sphere of influence depends on whether or not the human agent has an access constraint that prevents interaction.

In Table 9.6, demonstrates the distinction between the theoretical maximum size of the sphere as well as the in practice value. The network underlying the sphere of influence is designed exogenously by the experimenter prior to the start of the run. However, the actual sphere of influence in practice is the set of partners with whom the individual agent interacts due to homophily, expertise, or socio-demographic similarity. Since agents often do not interact with all of their potential partners, the effective size of the interaction sphere in practice is often much smaller than the theoretical maximum.

TABLE 9.6 A table illustrating the way in which the user can adapt the agent classes by specifying the size of the sphere of influence per class.

Factors	General Population Human	Opinion Leader	Media Agents				
			Ad	Web	Call	Radio	Mail
Sphere of Influence Theoretical Maximum	40±10	3000	3000	3000	3000	3000	3000
Sphere of Influence in Practice	25±10	250±50	300±75	100±50	66±20	250±100	150±50

For each pair of agents, the probability that they interact is a function of proximity, socio-demographics, knowledge, beliefs, and the interaction logic. Since the socio-demographics remained constant in this study, the overall probability of interaction contains both a fixed and a non-fixed component. Since these overall probabilities can change, we say that the social network is evolving as who actually interacts with whom will vary over the simulation run: the interaction partners of the early simulation periods will differ substantially from those of the later periods. This evolution can be observed in the changing likelihoods that the agent has for interacting with those in its sphere of influence; however, the size of the sphere of influence and the topology of the fixed portion of these probabilities do not change. Thus, as the probability of interaction increases for any pair of agents, that increase must come relative to that of other agents

in the interaction sphere and must mean that both agents are evolving to become relatively less similar to all other possible interaction partners.

Using Construct a number of different topologies for the fixed portions of these networks can be examined. They can be random [217], cellular [218], or small world [219], [220]. The accuracy of the simulated topology is extremely dependent on number of agents and the overall density. For example, with 10 agents and a density of 0.5 it is not possible to cannot get a cellular network – agents are too interconnected to exhibit cellular structures; similarly, a network with 100 agents and a density of 1.0 is not random, cellular or small world as everyone is uniformly connected to anyone. When the populations have more than 3000 agents, the selected densities and sphere of interaction sizes should be selected to ensure that the topologies examined are good representations of that topology; i.e., truly random, cellular or small world.

The key differences in the random, small-world and cellular network are clustering and the ratio of internal to external ties. In a random network the links are distributed independently and identically. These are the general fixed communication links. In the small-world network, each agent has a few links and a few agents have many links. In contrast, in a cellular network the agents are clustered in to a few cells and mostly communicate with other cell members while only one or two members per cell interact with anyone in another cell. The placement of these ties affects the diffusion of information throughout the society and has the potential to lead to different rates of diffusion among and between different agents.

9.3.6 Constraints on Information Access

Cognitive and social factors combine to determine the level of information access that individuals may have. We examine three different information access mechanisms: literacy, internet access, and newspaper readership [221]. Within Construct, these access mechanisms affect whether agents can interact with a specific media and get information through a specific forum. These mechanisms are implemented as “switches” that the researcher can enable or not, depending on the research question.

In Construct, agents can be literate or not, as set by an experimenter-controlled switch. The literacy mechanism affects all media that require reading printed material. This means that printed advertisement in newspapers, web site, and information sent in letters via the postal system are affected. When literacy as an information access parameter is enabled, illiterate agents can still access these media; however, they do not learn all the information and beliefs conveyed in the message and they may even mis-learn information. A small level of mis-learning is implemented as the literature on literacy shows that literacy is in part a matter of degree which often leads the illiterate individual to misinterpret what is being read. Literate agents are unaffected by enabling the literacy mechanism, and receive the full information from these media. When the mechanism is disabled, all agents receive the full information.

In Construct, agents can surf the web or not – and those that do have access to internet-based media. When the internet access constraint is enabled, agents lacking web access cannot read information posted on web sites at all. Agents with internet access can read such information, and use this information to affect subsequent interactions with other non-web agents. When the mechanism is disabled, all agents can read information from web sites.

In Construct, agents also have the ability to read newspapers and access the information contained in them. The newspaper access mechanism affects all media that require physical new-

sprint such advertisements in newspapers and specialized articles by opinion leaders. When newspaper access is enabled, agents lacking newspaper subscriptions cannot read articles published in the paper. Agents who are newspaper readers, though, can still read such information. When the mechanism is disabled, all agents can read information printed in newspapers.

It is important to note that these mechanisms interact. For example, if an agent is illiterate and has a newspaper subscription, that agent may read the news articles but do so with error. On the other hand, if an agent is literate but does not have access to the internet, they still cannot read web-pages (and the literacy parameter has no effect).

For each agent class, the researcher must exogenously specify whether or not an access constraint applies, and the probability that an agent in that class is constrained. In this study access constraints only apply to general public human agents. That is neither the opinion leader nor the media agents are constrained. In this study, the probability that an agent is illiterate, cannot access the web, or does not read a newspaper was derived from socio-demographic attributes and national averages. A series of formulas, one for each constraint, that determine the probability that the agent is constrained based on age and education were derived from national data (see [221] for details).

Chapter 10

Adversary Modeling –Applications of Dynamic Network Analysis

Kathleen M. Carley, Il-Chul Moon, Geoffrey Morgan, Michael Lanham

10.1 Introduction

This chapter focuses on the general utility of dynamic network analysis for adversary modeling and presents a brief survey of a various modeling efforts. The tools discussed in the previous three chapters, Automap, ORA, and Construct, help their users to perform, among other tasks, various kinds of **dynamic network analysis**.

Dynamic network analysis (DNA) focuses on the theory and design of dynamic networks among diverse entities, and the study of all phenomena emerging from, enabled by, or constrained by such networks. Entities include both actors, such as robots and humans, and artifacts, such as events or resources. DNA focuses on developing theory around the creation and maintenance of such networks, on developing plausible and useful models, on measuring changes in networks, and on evaluating these found networks.

Because DNA is indeed, *dynamic*, changes over time are a critical and novel aspect of the science that distinguishes it from the larger body of network science and social network analysis. A practitioner of DNA, with the appropriate data, can answer all five journalistic questions: Who?, What?, Why?, How?, and When?.

It has become, with state of the art tools and good data, routine to answer questions such as:

- Who are the Key Actors?
- What resources are critical to this network?
- What locations are critical to this network?
- Who are the leaders in this group?
- How likely are tasks given to this network to be performed well?

But these questions are mostly *static*, developing good methods of answering questions about the dynamics of the network is one of the principal research areas. Advances in this field are and would allow questions such as these to be answered:

- Has the group changed, and how?
- Has an individual's role changed?
- Are new groups forming?
- Are the sub-units different?

But answering these kinds of questions is difficult, not only because quantifying change is a hard problem, but because data is usually incomplete and almost always has some amount of uncertainty to it. Inferring missing links and nodes in the network is one of the key problems – particularly when studying adversarial and covert networks, where obfuscation is one of their key protective mechanisms. Multiple inference mechanisms exist.

Since adversarial networks are dynamic, DNA techniques need to provide for ways of assessing and forecasting change, and evaluating the impact of courses of action. Given the likelihood of missing data, DNA techniques need to be robust against potentially missing data. One way of increasing robustness is to focus on what group of nodes are critical as opposed to focusing on the exact rank ordering. Missing links can be inferred and then DNA techniques can be used to assess the impact of inferred links on the conclusions drawn.

We now present a few examples of applications illustrating various ways adversarial forces have been modeled with DNA techniques.

10.2 Predicting important changes: the organizational restructuring of Al-Qaeda

A key question in understanding the adversary is “when did they change?”. To address this question change detection tools for networks can be used. In this case, given a set of networks from a sequence of time periods, change in the group is detected by identifying shifts in trends for key network metrics such as betweenness.

We used AutoMap to extract a sequence of networks describing the connections among members of al-Qaeda from open source documents. A set of networks, one per year, from 1994 to 2004 were extracted. These networks were then examined with ORA and key metrics calculated including closeness, density, average betweenness, and the clustering coefficient. This data was then used to assess change. In Fig. 10.1 the Closeness CUSUM Statistic for this data is shown. This technique showed that although the evident signal change was in 2001, the change began in 1997.




Fig. 10.1 The closeness CUSUM statistic graph over time for Al-Qaeda

The strength of this technique is that identifies the historical change point. It is critical to note that this technique is for assessing historical data. It is not a predictive technique. Further, it is left to the analyst to identify what factors at the change point could have led to the change. In this case contributing factors may have been that bright star was cut short.

10.3 Assessing Destabilization From a Social and Task Based Perspective

In many domains where situations are dynamically changing, ‘what-if’ analysis is a critical question to prepare for the future. Some disciplines, such as intelligence, corporate management, military command and control, etc, have various threat scenarios and organizations wonder what will happen if the scenarios become realized. For example, from the perspective of destabilization analysis, an interesting question is what will happen to a terrorist organization if key terrorists are removed. Destabilization analysts want to know the deterioration of the adversarial organization’s performance and the organizational structure after their removal. Doing such what-if analysis is a way of evaluating alternative courses of action.

The ideal methods to answer these questions would be replicating the target domains and the organizations many times in the real world and testing the scenarios in the replicated environments. The organization science and social science communities approximate the above experiments through field studies and collection of experimental data in labs. However, these techniques are very expensive, and can potentially transgress into unethical or immoral areas. Adversarial organizations, in particular, are difficult to replicate in the real world. Generally, we have limited understanding about their organizations or their complex collective emergent behavior arising from their decentralized structures.

Overcoming these limitations has been one of the most significant benefits of using Agent Based Models (ABM) with large numbers of heterogeneous agents. The nature of ABMs have provided a nice analogy to human organizations and actors which has allowed policy domains, such as civil violence [222], the transportation of goods [223], [224], to experiment and use ABMs. For example, Bio-war [225] is a city scale ABM for examining the impact of various interventions in mitigating the impact of weaponized disease attacks or pandemics. Additionally, the growth of computing power allows the researcher or policy analyst to use ABMs to run multiple experiments for many times with less cost. Researchers interested in organizational theory and strategy have used ABMs to build and develop theories of organizational learning and performance [226], [227], [228].

Figure 10.2 reflects the generalized approach taken by researchers conducting Dynamic Network Analysis. The simulation analysis begins by selecting a target organization to simulate and proper parameter value selections. This is a specification applied to every simulation. In contrast to these general simulation specifications, a researcher can apply different simulation scenarios, i.e. by changing who to remove and when from a simulated organization. Each of these different simulation scenarios forms an experiment cell in a virtual experiment. A researcher then replicates each experiment cell with a coded simulation model. After the replications, the simulation model generates 1) organizational performance and general statistics and 2) detailed agent behavior records over the course of simulations. With regression analysis, analysis of variance, and simple visualizations the researcher(s) can review and analyze the performance values and log records. Of particular note is the ability to chain experiments together as simulation generates an estimated organizational structure(s) and element distributions at the end of the simulation run. The estimated organizational structure(s) can be fed back to the simulation model, and the simulation analysis cycle can start again.

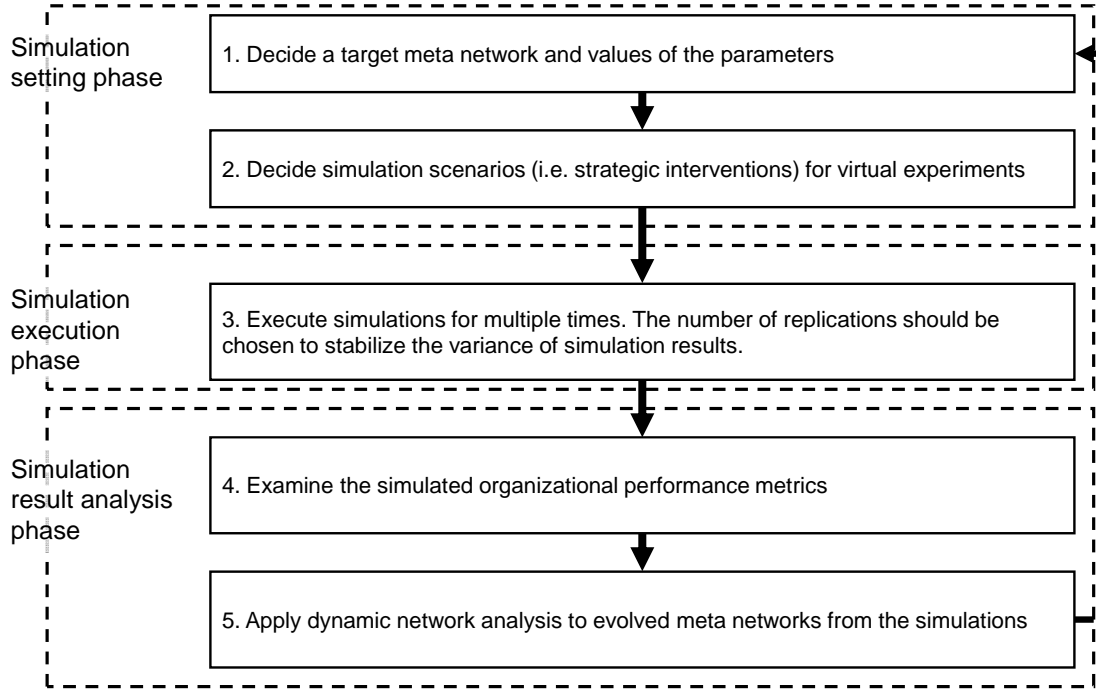


Fig. 10.2 An overall simulation analysis procedure

In the remainder of this section, we’ll review ‘what-if’ analyses of adversarial organizations under different possible intervention scenarios. The analyses are done by using a version of Construct modified to work directly from ORA using networks generated by the user and outputting the evolved organizational structure of the group being simulated and its performance characteristics. This version of Construct is referred to as JDyNet. Next we’ll review the collected structural datasets, the modeled the target organizations’ agent behavior, such as task performance, information diffusion, and resource passing. Finally, we’ll review the created hypotheses with dynamic network analysis from the viewpoint of terrorist removals, and how we turned the hypotheses into simulation scenarios. With the input organizational structure dataset, agent behavior model, and simulation scenarios, we gauge the impact of the intervention scenarios. In this analysis we use of the Tanzania and Kenya dataset reflecting publicly available information about the bombings of the American Embassies which have been coded as a meta-network connecting people, resources, expertise, and tasks.

10.3.1 Simulation Model Description

JDynet is the simulation model designed and used to estimate the collective behavior of adversaries throughout this chapter. JDynet takes a number of inputs which reflect an adversarial organizational structure and parameters. After a simulation run, JDynet produces an expected post-scenario organizational structure and various over-time organizational performance scores. During the simulation, JDynet calculates its internal status variables repeatedly and simulate the time flow. This analysis procedure incorporates inputs, outputs and simulation model internal variables. Table 10.1 summarize the relevant variables.

TABLE 10.1 This table contains a summary the input and output variables, and the associated parameters, for the JDyNet simulation runs with associated names and description.

Type	Name (Default value in the parenthesis)	Implication
Input	A networked organizational structure (a meta-network)	A network including agents, knowledge bits, tasks, and locations. The network represents the target domain's complex organizational structure.
	Simulation scenario	A sequence of agent removal specification. An element of sequence specifies the removal target agent and the removal timing.
Output	An evolved network organization (a meta-network)	A network organization with a recreated agent-to-agent (AA) network and an agent-to-location (AL) network, both of which reflect organizational element transfers, social interactions and geospatial relocations.
	Diffusion	A performance metric showing how fast information can diffuse across the network.
	Energy task accuracy	A performance metric showing how accurately information is distributed to agents who require it to complete their tasks.
	Binary task accuracy	A performance metric showing how accurately agents can classify their binarized assigned tasks with provided information
	Task completion	A performance metric displaying what percentage of the organization's tasks are completed
	Task completion speed	A performance metric displaying how quickly each of tasks can be completed on average. The inverse of the average task completion simulated time-step
	Mission completion speed	A performance metric displaying how quickly the entire task dependency network can be completed. The inverse of the mission completion simulated time-step
	Gantt chart	An estimated mission progress displayed in the Gantt chart format
Parameters	Number of time-step (5000)	The number of simulated time-steps
	Number of replications (30)	The number of replications to stabilize the outputs of this stochastic simulation
	Weights for requested element delivery (0.33), others' request passing (0.33), or the agent's request passing (0.33)	Only used in task performance agent interaction model. Weights for selecting an agent interaction purpose. An agent selects one purpose out of three, requested organizational element (expertise or resource) delivery, his required element request to others, or passing others' request to different others.
	Correct binary task accuracy threshold (0.5)	When calculating binary task accuracy, the agents have to make guesses on the unknown information. This number specifies the probability of the correct guess

Type	Name (Default value in the parenthesis)	Implication
	Interaction count for time-step (3)	An agent cannot interact with another agent after this maximum interaction count.
	Cognitive power for time-step (3)	An agent can only respond to the number of interactions specified by this parameter.
	Exchange success rate (0.75)	If an agent diffuses information or passes a resource to another agent, there is a success rate of such trials.
	Interaction social distance radius (1)	Interaction candidates are limited to agents who are within N social link radius from the interaction initiating agent.
	Task execution success rate (0.5)	When an agent performs a task, the agent can accomplish the task with this success rate. If the task is not ready (the ready state is elaborated later), an agent cannot perform the task.
	Exchange only required elements (true)	If this is true, agents only exchange expertise or resources only the receiving agent needs such elements.
	Treat resource as information (false)	If this is true, resources are duplicated when it is passed, so that the sending and receiving agents have the passed resource.
	Take over removed agent links (true)	If this is true, an agent recognizing that the interacting agent is removed can take over the target agent's various links to organizational elements, other agents and assigned tasks.
	Recognize that interaction partner is removed (0.1)	This is a success rate that an agent recognizes the interaction target agent is actually removed.
	Recover links from the removed agents (0.3)	After an agent recognizes another agent is removed, the agent can recover the links between the agent and the other agent with this probability
	Request decay time (7)	After this number of simulated time-steps, the organizational element request is removed.
	Transactive memory decay time (7)	After this number of simulated time-steps, an agent's transactive memory about other agents is removed.
	Maximum transactive memory element (30)	This is the maximum number of transactive memory about other agents' links

Subsequent sections introduce and explain where the simulation uses these inputs, outputs and parameters. The sections will also discuss the values selected, and interpretations of the values. Researchers design virtual experiments by varying these parameters or inputs to test hypothesis on the destabilization of a target organization. A human analyst 1) selects the most appropriate parameter values, 2) strategizes the agent removal sequence and 3) runs a number of simulations with the specifications. After the runs, the analyst drills down the organizational performance degradation and correlates the impact with his agent removal sequence.

Agent social behavior

JDynet agent behavior is largely in two parts: social interaction and task performance. An agent initiates social interactions to receive expertise or a resource from the interaction partner or to send a request for expertise or resource to the partner. An agent also executes a task that is ready for execution. A task is ready for execution if all the prerequisite tasks are done, and if the group of assigned agents has at least one required resource and expertise. More detailed descriptions are in the following sections. Fig. 10.3 shows the high-level agent behavior flow during the simulations.

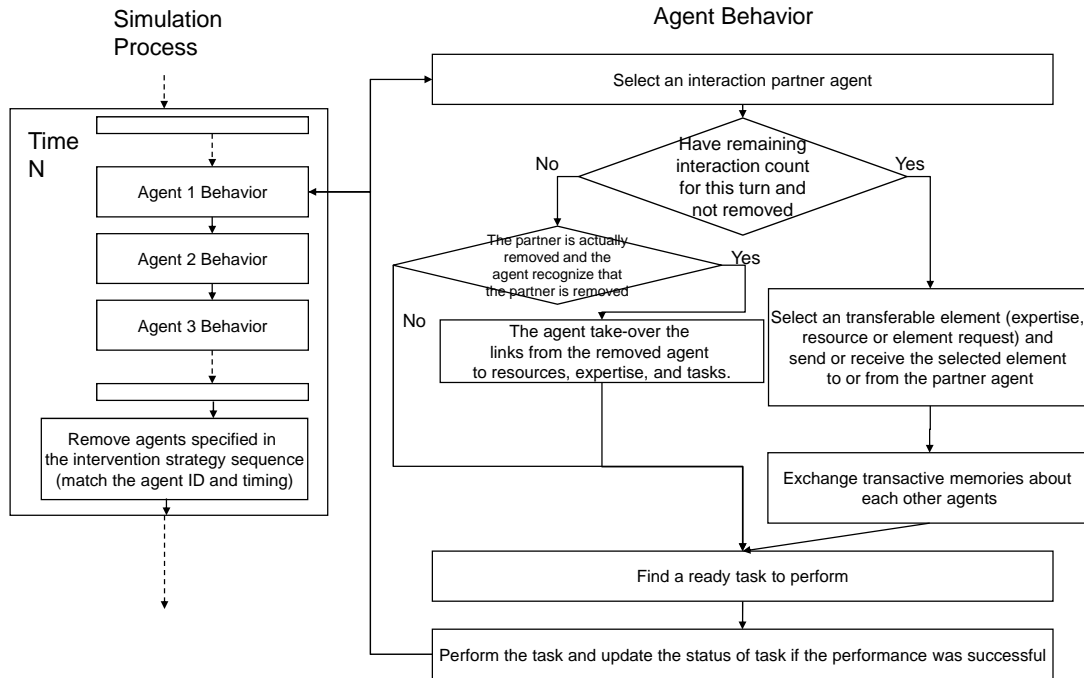


Fig. 10.3 High level agent behavior logic

Selecting an interaction partner agent

The model of agent behavior with respect to tasks draws on research in the operations research area and combines this with the socio-cognitive agent behavior model in Construct [8]. An agent only initiates interactions with others if they need to communicate with them to perform his assigned tasks. They may seek their own necessities, pass the past interaction partner's request for resource or expertise, or pass the acquired resources or expertise to the past partner who needs them.

Agents in this model select an agent as an interaction partner if he can give a necessity to them. If there is no agent that needs an organizational element (e.g. a resource, knowledge), the agents choose an interaction partner randomly. Additionally, an agent can pass expertise, a personal resource, as well as an element request that he or a past-interaction-partner initiated. This model illustrates how the agents will interact when they are goal-oriented. While the sociological model is appropriate for simulating the belief or ideology dispersion, this model is appropriate for simulating the organizational collective behavior to complete the tasks in their task network.

This task-completion oriented agent interaction is modeled as three different agent choice motivations below.

- 1) *Choosing a motivation for interactions*: An agent chooses one interaction motivation out of three motivations: requested element delivery, others' request passing, and the agent's request passing. This is a random weighted selection, and the weights are specified by an analyst, *Weights for requested element delivery, others' request passing, and the agent's request passing* in Table 10.1. After the choice of the motivation, the agents select an agent as the following partner choice mechanisms.
- 2) *The agent's request passing*: If an agent chose the agent's request passing motivation, the agent finds one required-but-not-acquired expertise or resource to perform his assigned tasks. Then, the agent searches an agent who has the required organizational element, and he initiates an interaction with the searched agent to receive the required element. If there is no agent with the required element, the agent interacts with a randomly chosen agent and leaves a request for element delivery. The possible interaction partners are limited as the sociological limit the interaction candidate set.
- 3) *Others' request passing*: If an agent chose the others' request passing motivation, the agent finds one requested element among the requests for element delivery from others. Rest of the selection procedure is identical to the agent's request passing motivation. The agent searches an agent who has the requested organizational element, and he initiates an interaction with the searched agent to receive the required element. If there is no agent with the required element, the agent interacts with a randomly chosen agent and leaves a request for element delivery. The possible interaction partners are limited as the sociological limit the interaction candidate set.
- 4) *Requested element delivery*: If an agent chose the requested element delivery motivation, the agent will find an agent who left a delivery request during the past interactions. The organizational element in the delivery request should be possessed by the agent. Then, the agent initiates an interaction with the found agent to send the organizational element that the interaction partner requested previously.

Transfer an organizational element or a delivery request

The effect of an interaction between two agents is either resource passing, expertise diffusion or delivery request. There are also two ways of modeling organizational elements transfer. The original Construct model did not differentiate a resource from expertise from the perspective of diffusion. The interaction sender's resource is duplicated and put in both sender's and receiver's possessions. Therefore, in the original Construct, interaction results in the diffusing of organizational elements, not the passing on of requests. Here, in JDyNet we explore how to extend Construct to include requests.

On the other hand, this suggested model provides a new way of producing interaction outcomes. First, a resource is not duplicated and just passed from the sender to the receiver. Second, an agent can leave a delivery request for expertise or a resource, so that the interaction partner can remember that the initiating agent needs such elements. Both ways of transferring an organizational element allows only one element transfer for a single interaction. If the simulation has already removed the agent from the possible interaction list (e.g. the agent is dead, unavailable for that 'turn'), or has exceeded the number of interactions specified as *maximum interaction*

count for time-step in Table 10.1, then the agent cannot transfer any of expertise or resources. By taking these diverse forms of interaction together complex sequences of actions by agents can be captured as is illustrated in Fig. 10.4.

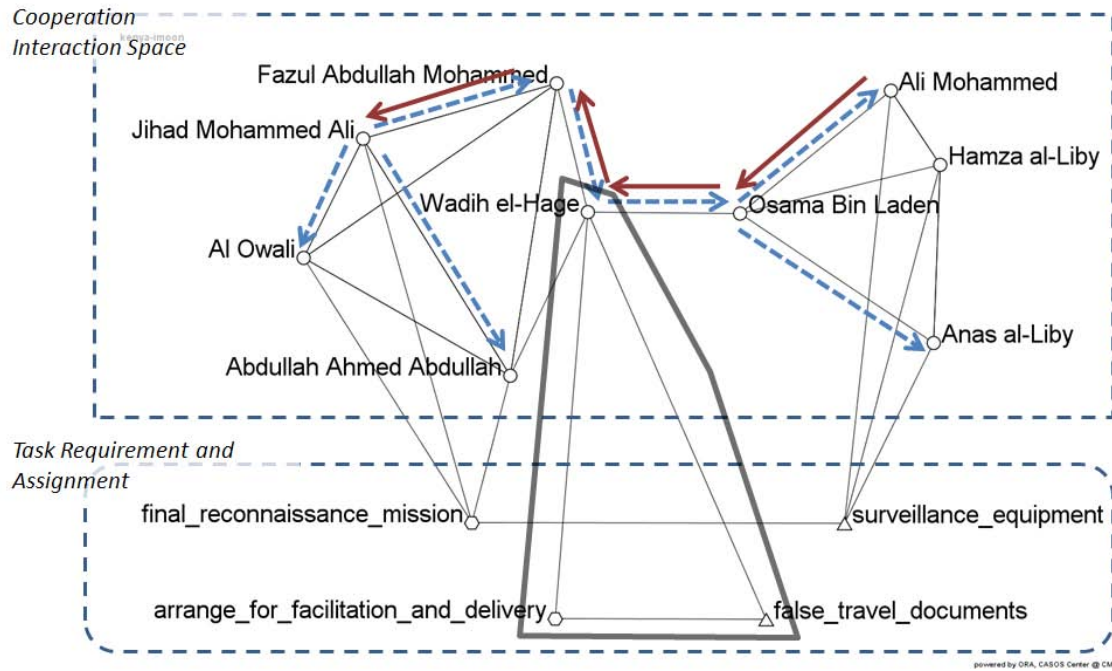


Fig. 10.4 An example of agent behavior during the simulation from the Kenya data. The dashed arrows are the organizational element (*surveillance equipment*) requests to the interaction partner agents. The solid arrows are the actual transfer of the *surveillance equipment*. The solid line polygon includes *Wadih el-Hage*; *arrange for facilitation and delivery*; and *false travel documents*. *Wadih el-Hage* can perform the *arrange for facilitation and delivery* task because he has the required resource, *false travel documents*.

Take-over the removed agent's expertise, resource, task and social contacts

If an agent initiates an interaction with an already-removed agent, the interaction-initiating-agent may recognize that the partner agent is not present. This recognition is turned on if an analyst makes *Take-over removed agent links* true. The recognition also depends on the random coin toss whose probability is specified as *Recognize that the interaction partner is removed* in Table 10.1.

After the coin toss, if the agent is allowed to take-over the removed agent's neighbor agents, resources, expertise and tasks, the agent creates a link to those legacies. However, to recover the links from the removed agents, the recovering agents should have prior knowledge about the existence of the link. This is modeled from the transactive memory. Each of the agents has transactive memory storing the perceived link information of other agents. After the link recoveries, he updates the agent-to-agent network space, so that the other agents wouldn't make more interactions to the removed agent. Each agent can communicate transactive information in addition, or instead of information that they directly know. The movement of transactive information is shown in Fig. 10.5.

This take-over mechanism simulates the resilience of an adversarial organization. The adversaries will reassign agents to resources, expertise and tasks, to compensate the removed agents.

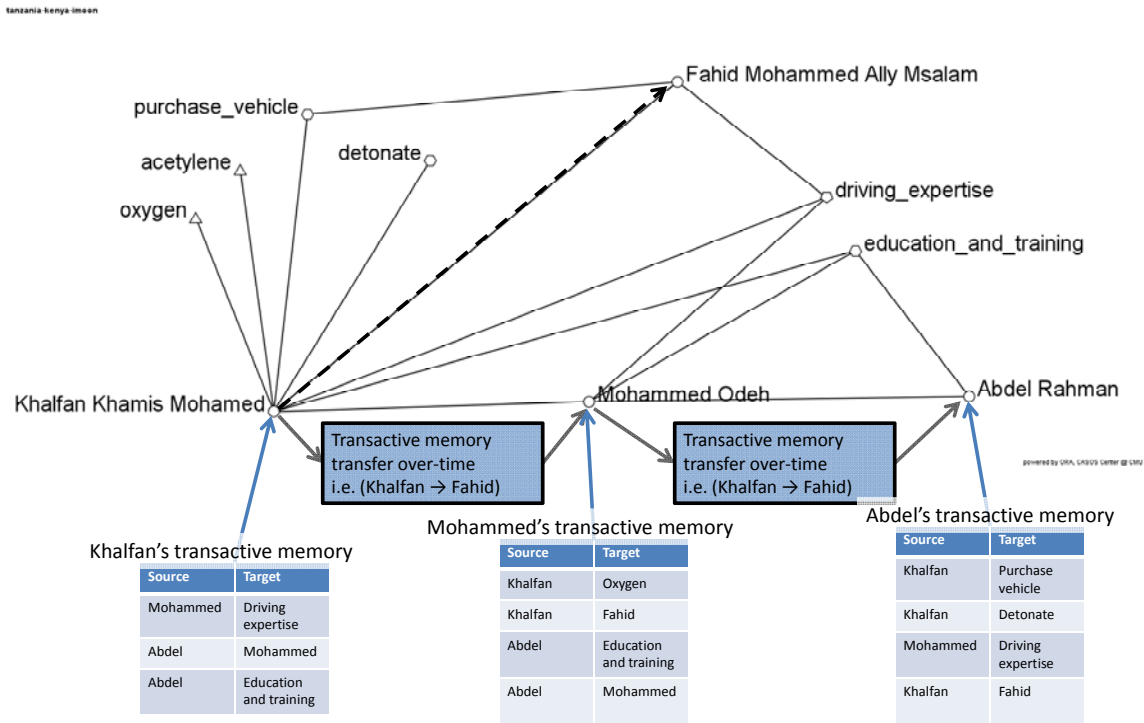


Fig. 10.5 A illustrative example of transactive memory transfer. A link information, such as *Khalfan Khamis Mohamed is linked to Fahid Mohammed Ally Msalam*, can be transferred through the interaction network among agents. The transferred transactive memory is stacked in the received agent's transactive memory repository. After transactive memory decay time-steps, the decayed transactive memory element is removed.

Perform an assigned task

An agent performs a task if the task is ready for execution. There are three statuses for a task according to the resource and expertise distribution over the course of simulation.

- 1) *Not Ready*: A task is not ready if its prerequisite tasks are not completed. The prerequisite tasks are defined in the task dependency network of the input meta-network.
- 2) *Ready*: A task is ready if its prerequisites are done. However, this ready status does not guarantee the task completion. The group of assigned agents has at least one piece of each required element to the task, but the group may not be fully equipped with expertise and resources to perform the task. From ready status, an assigned agent can perform the task by coin-tossing with *Task execution success rate probability*. If the required resources and expertise are not acquired by the group of assigned agents, the task performance will always fail.
- 3) *Done*: A task is done if the group of assigned agents has the required expertise and resources, at least one piece, and if one of the assigned agents performed the task by successfully coin-tossing whose probability is specified in *Task execution success rate*.

10.3.2 Virtual Experiment Design

Virtual experiment design for destabilization analysis consists of two parts. First, an analyst needs to specify the simulation model parameters, such as the number of simulation time-steps, interaction ways (either sociology-oriented interaction or operations-research oriented interaction), weights for interaction methods, etc. Second, an analyst needs to compose a simulation scenario: who to remove and when. Fundamentally, an analyst can determine the values for the parameters specified in Table 10.1 with his qualitative insights into a target organization.

The presented virtual experiment in this section varies simulation scenarios in three ways: removed agent selection scheme; number of removals; and removal timings. The permutation of these three factors and values are listed in Table 10.2. There are sixty-four (64) different virtual experiment cells that each have a different simulation environment. For instance, the experiment cells with larger *intervention size* remove more terrorists over the course of simulations. The experiment cells with *later intervention timing* removes agents in the relatively late phases of simulations. Also, the experiment cells have diverse intervention target selection schemes according to *removal target selection scheme*. This is the manipulation of simulation scenario. Further analyses can be done by changing the simulation parameters, but such experiments are not done in this case. Human analysts can alter a virtual experiment by altering the default value that used in this report and listed in Tables 10.1 and 10.2

TABLE 10.2 A table describing the design of a virtual experiment assessing the impact of diverse courses of action for targeting difference adversaries. For each cell shown there would be 15 replications and 2500 simulation time steps.

Name	Value	Implication
Removal target selection scheme	Degree, Betweenness, Eigenvector centralities and Cognitive Demand (4 cases)	Agents with high network values are considered critical, and their removal is critical to the organizations. This is how we pick target agents to remove.
Intervention size	1, 5, 9, and 12 agent removals (removing 10%, 30%, 50% and 70% of agents, 4 cases)	The intervention size specifies how many agents to remove with this intervention.
Intervention timing	125, 250, 500, and 1000 time-step (removing at after 5%, 10%, 20% and 40% timeflow, 4 cases)	The intervention happens at a specific stage of simulation period.
Total virtual experiment cells	64 cells (4x4x4 cases)	

The following dynamic network analysis measures decided which agents to remove over the course of simulations: Degree, Betweenness, and Eigenvector centralities and Cognitive Demand. Without such analysis, a human analyst would need some other heuristic to reduce the quantity of possible simulation scenarios. Interpretations of the metrics listed above are in Table 10.3.

TABLE 10.3 Dynamic network metrics used to determine the target agents to remove

Name		Interpretation	Reference
Degree	Cen- trality	Number of in-coming and out-going links from a node, Degree of direct influence to others	[229]
Betweenness		Number of shortest paths passing a node, Degree of information flow control	[229]
Eigenvector		Calculates the eigenvector of the largest positive eigenvalue of the adjacency matrix, Degree of connections to the high-scoring nodes	[230]
Cognitive	De- mand	Measures the total amount of effort expended by each agent to do his/her tasks, calculation details are elaborated below.	[231]

Remove an agent specified in the intervention sequence

JDynet helps a researcher or analyst assess the impact of intervention strategies, or an agent removal sequence. In Table 10.1, there is an input, *Simulation scenario*. Simulation scenario is a sequence of agent removal specifications. An agent removal specification displays the target agent to be removed and when the target will be removed in the simulation time.

At the end of every time-step, JDynet goes through the agent lists and finds an agent that should be removed at the time-step. If an agent is removed, then the agent cannot make any actions, either social interactions, organizational element transfers, or task performances.

Performance measures

There are four performance metrics to assess the change of the organization: *Diffusion*; *Energy Task Accuracy*; *Binary Task Accuracy*; and *Task Completion*. The performance metrics support the evaluation of performance of the evolving organization over time. Evaluation is also necessary by comparing evolved performance values to those of a non-intervention case (baseline). The metrics are shown below.

- 1) Diffusion: Diffusion measures the dispersion of expertise and resources across the agents.

$$(Diffusion) = \frac{\sum_{i=0}^A \sum_{j=0}^K AK_{ij}}{K \times A}$$

- 2) Energy Task Accuracy: Diffusion only considers who knows or has what. Whereas, energy task accuracy calculates the extent to which the agents have the knowledge they need to do the tasks they are assigned. This is done by introducing the agent-to-task (AT) and knowledge-to-task (KT) network in the formula.

$$ETA = \frac{1}{T} \sum_{t=0}^T \frac{\sum_{k=0}^K (KT_{kt} \times \sum_{a=0}^A AK_{ak})}{\sum_{a=0}^A AT_{at} \times \sum_{k=0}^K KT_{kt}}$$

- 3) **Binary Task Accuracy:** Binary task accuracy measures the agents' binarized, assigned task classification capability with the current information and resource availability. A task classification is performed by classifying N organizational elements required to perform the task. An agent always classifies an organizational element correctly if he has the element. If the agent does not have an element, he can guess the correct answer with 50% of chance. Therefore, if an agent has M (<N) required elements, then he has to guess (N-M) elements to get the result of the binarized task. The task performance is 1 if the agent classifies more than 50% of required elements correctly.
- 4) **Task Completion:** Task completion measures the number of completed task over the course of the previous simulation period. A task is completed if the task's status is in *done* status as explained in the previous section. Task completion is a simple ratio calculated from

$$\frac{(\text{number of completed tasks})}{(\text{number of tasks in the organizational structure})}.$$

- 5) **Task Completion Speed:** A task duration is the simulated time length between the task's ready status to done status. Then, each task's speed is determined by inversing the task duration. I average each of the task speeds and calculate the organization level task completion speed.
- 6) **Mission Completion Speed:** Mission completion speed is the inverse of the number of simulation time-steps over the course of the task dependency network completion. The task dependency network completion means the entire task network completion by completing individual tasks one by one.

10.3.3 Results

We ran the above virtual experiments with the Tanzania and Kenya embassy bombing case. The first sub-section examines the agent removal impacts toward organizational performances. The second sub-section examines the delayed task completion timing caused by the agent removals. The third sub-section enumerated the key individuals over the course of simulations. The last section of the results are the visualization the agents' collective behavior during the simulations.

Impact to performance measures

After running four different simulation scenarios for each of 64 virtual experiment cells, we get 64 simulation results. Table 10.4 shows the regression on the simulation settings to the organizational performance metrics. This regression is done by using the two continuous virtual experiment factors (timing and size) and one factor (removal selection scheme) with four categories. There are four categories in the removed agent selection metrics. The four categories' representation is through assigning 1 if the simulation used the metric, and 0 if it did not. According to the regression result, earlier interventions (smaller intervention timing value) and larger interventions (larger interventions size) are preferable in reducing the performance. In terms of the removal target selection, removing top Degree Centrality terrorists can reduce the mission execution speed, the task execution speed, the binary task accuracy and the level of diffusion. Similar trends are seen in the case of removing top Eigenvector Centrality terrorists.

TABLE 10.4 A table showing the standardized coefficients for regression to the six organizational performance metrics at the end time using the virtual experiment settings (treating removed agent selection scheme with four categorical values) (N=64 cases) (* for $P < 0.05$)

Standardized Coefficient	Mission Speed	Task Speed	BTA	ETA	Diffusion	Task Completion
Intervention Timing	0.143*	0.205	-0.021	0.081*	-0.070*	0.369*
Intervention Size	-0.808*	-0.063	-0.263*	-0.978*	0.977*	-0.817*
Degree Cent.	-0.013	-0.192	-0.121	0.044	-0.043	0.033
Betweenness Cent.	0.195	0.074	-0.166	0.041	-0.042	0.132
Eigenvector Cent.	0.014	-0.096	-0.312*	0.018	-0.014	-0.003
Cognitive Demand	0.000	0.000	0.000	0.000	0.000	0.000
Adjusted R-Square	0.679	0.005	0.046	0.958	0.954	0.797

Table 10.5 shows another regression analysis used to investigate the characteristics of removed agents. For this regression, researchers compiled the average network metrics of removed agents and virtual experiment settings. Again, the *intervention size* has significant influence over the mission speed, energy task accuracy, diffusion and task completion. If the intervention size gets larger, the above metrics get smaller. The *intervention timing* affects somewhat influence over mission speed, binary task accuracy and task completion. If the intervention timing gets earlier, the damage gets larger (and actual performance values decrease). The regression indicates that *Eigenvector centrality* has higher influence and important metrics in predicting the simulated organizational performance. For example, if we choose to remove low *Eigenvector centrality* agents, we can lower the mission execution speed, task execution speed, binary task accuracy, energy task accuracy, and task completion. If we choose to remove high *Degree centrality* agents, we can reduce the mission execution speed, task execution speed, binary task accuracy, energy task accuracy and task completion levels.

One thing should be noticed is the high R-square values. In typical cases, agent based social models do not produce high R-square values because inherent randomness and complex agent behavioral models. In contrast, the presented operations research based model shows high R square values in the linear model.

While this is an overall result of the 64 different virtual experiment settings, we present the results grouped by their first factors: *target selection scheme*, *intervention size*, and *intervention timing*. Fig. 10.5 is the over-time organizational performance evolution of the virtual experiment cell by the first factors.

TABLE 10.5 A table showing the standardized coefficients for regression to the six organizational performance metrics at the end time using the calculated metrics of removed agents (N=64 cases) (* for $P < 0.05$)

Standardized Coefficient	Mission Speed	Task Speed	BTA	ETA	Diffusion	Task Completion
Intervention Timing	0.143*	0.205	-0.021	0.081*	-0.070*	0.369*
Intervention Size	-2.198*	-0.121	-0.857	-0.806*	0.845*	-1.555*
Degree Cent.	-0.251	-0.985	-2.197*	-0.473*	0.542*	-0.614
Betweenness Cent.	0.137	-0.146	0.465	-0.151*	0.142*	0.050
Eigenvector Cent.	1.043*	1.389	2.587*	0.551*	-0.645*	1.104*
Cognitive Demand	0.608*	-0.319	-0.441	-0.184*	0.181*	0.219
Adjusted R-Square	0.814	-0.018	0.102	0.984	0.982	0.851

Binary task accuracy converges to the evolved state quickly because I used the modified version of binary task accuracy by averaging the values from the start time up to the measure calculation time. The energy task accuracy and diffusion charts exhibit big drops at the intervention timing: 125, 250, 500 and 1000. On the other hand, the task completion chart shows gradual damages over the course of simulations. If an agent is removed while the agent is not needed right now to execute current tasks, then the agent's removal does not decrease the performance right away. When the agent is needed, the baseline case can perform without serious problems, but the removal cases are damaged when the time comes. In terms of Energy Task Accuracy, large intervention leaves constant and permanent damages while early intervention leaves such damages from the task completion perspective.




Fig. 10.5 Organizational performance over time, aggregated by the first factor

Impact to task completion timing

The output from simulations can then be assessed with standard statistical tools or network analytic tools like ORA. Using the JDynet of Construct task completion behavior can be examined. JDynet regenerates the task completion status over the simulation period, and it generates a task completion speed, a mission completion speed and a Gantt chart. Task completion timing was one of the areas of investigation.

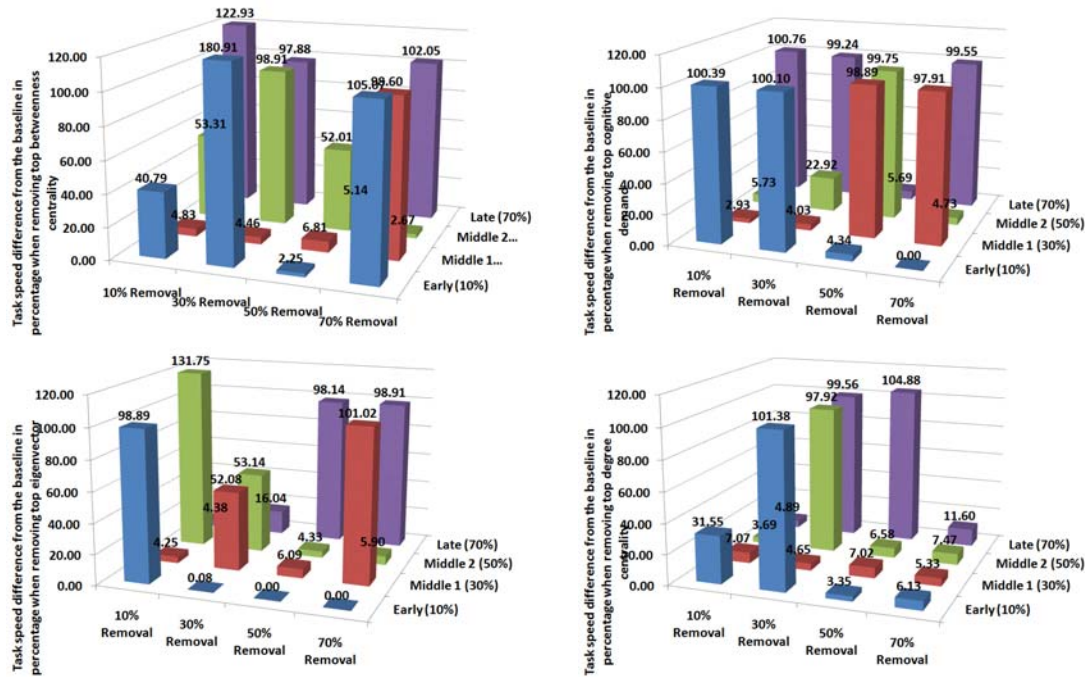


Fig. 10.6 Percentage of Task completion speed to the baseline, 64 virtual experiment cells

Figure 10.6 shows the task completion speeds of the 64 virtual experiment cells. The chart value is the percentage value of a specific virtual experiment cell compared to the baseline case. Therefore, if the value is higher than 100, it means the virtual experiment cell has faster task completion speed. If the intervention timing is late, the task completion speed is higher. On the other hand, if the intervention happens earlier, some tasks are impossible to be executed which makes their task completion speed 0.

Figure 10.6 shows that removing the high degree centrality agents is better in reducing the task completion speed. In most of the cases, the degree centrality based removal shows below 32% of the task completion speed compared to the base line (except four cases that show 101.38%, 97.92%, 99.56%, 104.88%). In general, removing a small number of agents late in the simulation does not cause any impact or damage though it sometimes increased the task execution speed.

Figure. 10.7 shows the mission completion speeds of the 64 virtual experiment cells. Since some removals disabled the organization's ability to execute their entire task dependency network, the cells with successful mission prevention show 0 mission completion speed (infinite execution time). These complete mission disables frequently happen when removing more than 30% of agents at the earlier stage. If the interventions are not successful, some cells show in-

creased mission execution speed (i.e. 103.55% of top cognitive demand agent 10% removal at early stage). Again, large and early removals show better destabilization effect compared to the small and late removals.




Fig. 10.7 Percentage of Mission completion speed to the baseline, 64 virtual experiment cells

Figure 10.8 shows the estimated Gantt chart of the baseline to show the bottleneck tasks and the task durations. This demonstrates the JDynet capability to generate a chart organizations and people frequently use in the real world. This helps identify which task(s) are bottleneck tasks that slow the mission execution speed. The *rent residence* task seems to have the longest execution time and seriously damages the mission execution speed. The *rent residence* is the prerequisite task to performing the *run bomb factory* task which leads the later task chains. Because of the *rent residence* task's delay, other tasks executed in the later phases got held up.




Fig. 10.8 The estimated Gantt chart of the baseline case

Agent interactions during simulations

The simulation allows an analyst or researcher to observe the agent interactions and organizational element transfers over the course of the simulation. Figure 10.8 is the Gantt chart of the baseline case over the course of the mission execution. As the Gantt chart displays, the agents focus on different tasks as the mission progresses. Moreover, the agents are assigned different tasks, which makes their interactions and organizational element transfers change over time. Fig. 10.9 is the collection of the agent interactions and organizational element transfer networks during the course of the simulation.

The agent interaction networks show no significant differences over time. There are minimal changes in the link weights. However, the agent organizational element transfer network changes dramatically, which means that the actual usefulness of the interactions change according to whether the interaction accompanies an organizational element transfer or not. The terrorists are bounded to their cellular network structure, so that the interaction network itself is not an obvious change. We need to see the implied usefulness of the interactions by looking into whether the link was used to actual resources or expertise transfer. In the transfer networks, there are isolated agents who are not used during the particular time period. In this case, a manager may consider reassigning the agents to other tasks which can be executed in parallel. Also, a commander may consider removing heavily used agents at a particular time-step when they can figure out which transfer network is going on at the intervention timing. Empirically, *Fazul Abdul-lah Mohammed*, *Al Owali* and *Wadih el-Hage* are the agents that consistently appear in the transfer network, which means that their removals would be effective in any of time periods.

10.4. Assessing the Impact of Destabilization on Adversarial Movement

One of the key problems in reasoning about adversarial behavior is that they move. Adversaries move from location to location as they plan and take action. ABMs tend to focus on either spatial movement through unrealistic grids or they ignore space entirely. Although the social patterns of adversary groups are useful for identifying them, the transnational movement patterns of adversaries can be an even clearer indicator [232], [171], [233]. A model that predicts important adversarial movement trends would be powerful and useful.

The previous example introduced a model for simulating adversaries' social behavior. This example extends the model by adding adversaries' geospatial behavior. The principal addition is to extend the operations research based social behavior modeled in the previous example with movement behavior. Agents relocate in order to complete tasks, and spatial-proximity influences link maintenance. If an agent is removed, and any neighbors are present, those neighbors are able to recover the agent's links. Finally, locations contain both resources and information, some of which may not be available to the organization at large. Agents may travel to maximize the resources available to their organization. Herein we explore how to adapt Construct so as to enable the adversaries to move from location to location, and in doing so expand the set of interaction logics to account for proximity driven interaction. By including space a number of factors are changed:

- 1) Simulation model iteration management
- 2) Social interaction logics
- 3) Knowledge of space
- 4) Task execution logics

This section describes the implementation of a spatial version of Construct, a virtual experiment exploring aspects of adversarial movement, and the results of that experiment.

We apply this extended model to the 1998 U.S. Embassy bombing incidents in Kenya and Tanzania. Over the course of the incident, the terrorists have to extensively move around regions and cross the borders to align the resources and expertise for the execution.

In this extension, an agent must select an interaction partner agent as well as a relocation destination. Therefore, as the agent chooses an interaction partner in the previous model, the agent chooses a place to move in this model (a hi-level behavior flowchart can be seen in Fig. 10.10). Agents have four different intentions for the relocation, these intentions range from: task performance, resource acquisition, interaction facilitation, and edge recovery (network robustness). The details on how these intentions impact behavior and the way in which we have extended Construct are now described.

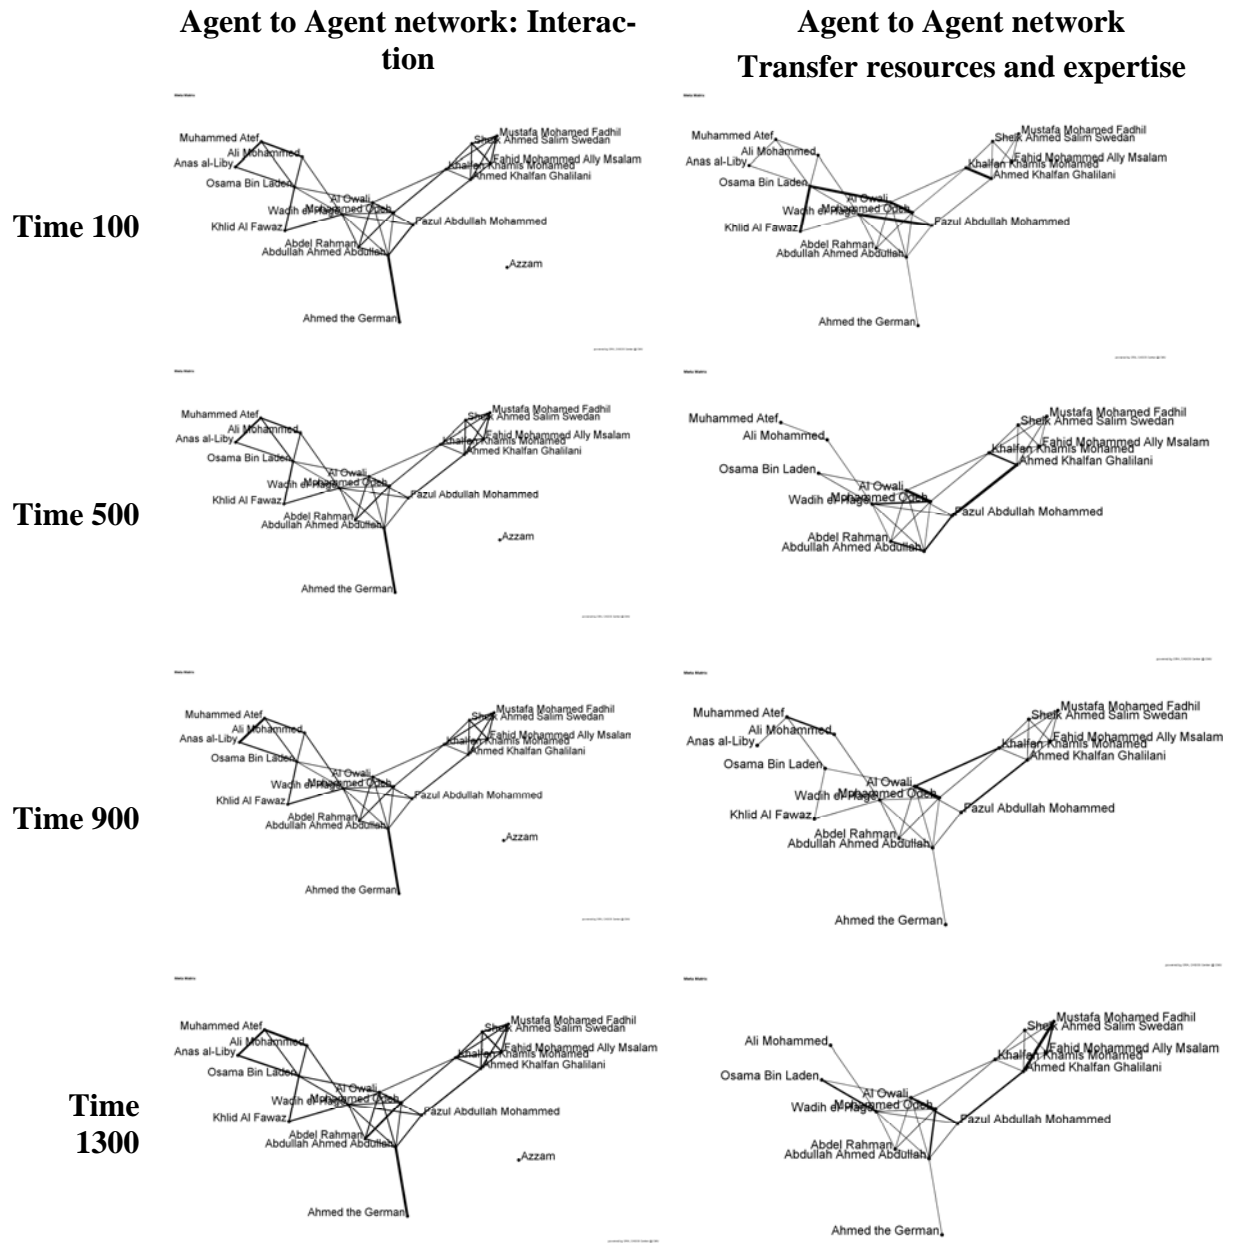


Fig. 10.9 Collection of agent interaction and organizational transfer network over time, link thickness is adjusted to show the frequency of the link usage.

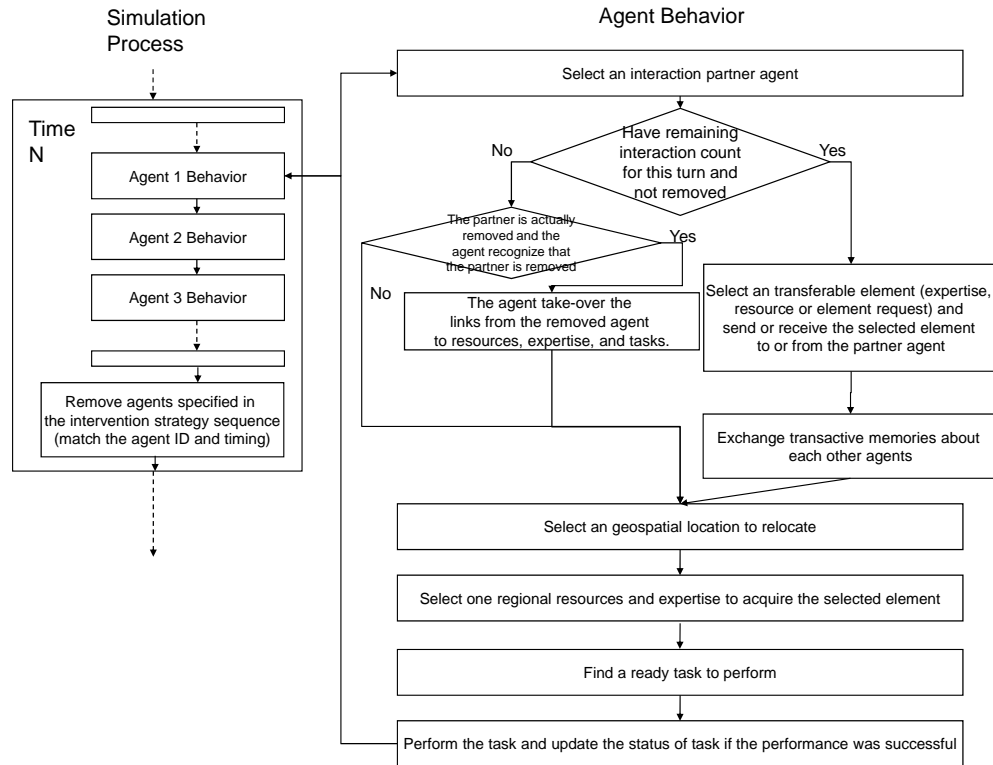


Fig. 10.10 Agent behavior logic. Compared to the previous behavior model, the geospatial relocation and the regional resource/expertise acquisitions are added.

10.4.1 Simulation model iteration management

Simulation iteration is the main loop of the entire model. When a user requests a simulation run, the simulation will execute this loop after loading the simulation model inputs and parameters. The simulation model runs a number of iterations to simulate the passage of time. These iterations are controlled by a loop controlling individual agent behavior. Table 10.6 contains the pseudo code for this outer or main loop.

TABLE 10.6 Geospatial simulation model main loop

<pre> Function main() Load simulation inputs; Setup simulation inputs; Setup random; For i = 0 to num_timestep simulation_iteration(random,i); Generate_and_output_intermediate_outputs //as directed by user configuration End; Calculate_performance_value_for_entire_simulation(); Generate_simulation_outputs; End function; </pre>

As depicted in the above pseudo code, the model runs the number of simulated time-steps with a random factor. This random factor makes this simulation stochastic. The reason behind this randomness is explained in the simulated agent behavior, when the agent must use this random factor to seed their coin-tosses.

After the simulation loop finishes, the program generates the performance outputs. There are two different types of outputs from this simulation model. First, we get the performance numbers, i.e. task completion rate, knowledge diffusion, task accuracy, etc. These numbers are printed out into files right away. The second output type is the estimated network outputs. These outputs are recorded in a DynetML file, so that the file can be loaded in ORA and visualize the over-time changes. Thus, Generate_simulation_output should handle these two types of outputs.

After coding this big loop wrapping the entire model, we code the individual iteration function that will be invoked over time. This simulation_iteration function is shown in Table 10.7.

TABLE 10.7 Geospatial simulation iteration for each time-step

Function simulation_iteration(Random r, int timestep) Agent_behavior_order = Randomized_order(1 to num_agent); While(Agent_behavior_order) i = next(Agent_behavior_order); Execute_agent_behavior(i, r); End; Calculate_performance_value_for_timestep(); End function;

By randomizing the agent behavior order, we can simulate the randomness in the action frequency. Also, by executing every agent's behavior for a single time-step, we can guarantee that the agents will execute their actions for the number of time-steps throughout the simulation.

The simulation's main loop calls the simulation iteration function for the user-specified number of time-steps. The simulation iteration function calls the individual agent behavior functions in a randomized order. The below is the specification of the individual agent behavior function. As discussed earlier, the agents gather knowledge and resources and perform assigned tasks by using gathered elements. Table 10.8 shows the agent behavior at a high level.

TABLE 10.8 High level agent behavior

Function Execute_agent_behavior(int agentID, Random r) Social_interaction(agentID, r); Perform_task(agentID, r); End function;

10.4.2 Social interaction logic description

The following pseudo code is the social interaction behavior pattern in the simulation. There are three social interaction motivations as discussed earlier in this chapter. The three motivations are 1) requested element (knowledge or resource) delivery; 2) other's element delivery request passing; and 3) the agent's element delivery request generation and passing. When an agent has a chance to make a social interaction, the agent makes a weighted random choice to select one motivation out of three. This weighted random choice represents the gap between the agent's intention and action. For instance, having a higher probability for the agent's element request generation is a representation of the agent's intention to get that element. However, in the simulation,

he might have to pass other's element request because of the randomness. Then, his action is different from his intention. Having said this, if he has a far higher probability for a certain motivation, then he is very likely to select the motivation out of the random choice. This reflects the strength of intention and increasing likelihood of his intention realization. Table 10.9 describes this social interaction and the role of neighborhood and proximity is detailed.

TABLE 10.9 Agent's social interaction implementation pseudo code

Function Social_interaction(int agentID, Random r) Neighbor_agents = getSphereOfInfluence(agentID, one social link away, return_only_agent); choice = weightedRandomChoice(r, weight_element_delivery, weight_others_request_passing, weight_my_request_generation); switch(choice) case Element_delivery: Element e = find_requested_and_possessing_element(agentID's elements); Request req = find_request_records_specified_by_element(agentID's received delivery request, e); If (req.sender has done his interaction for this turn) finish this block; If (req.sender and agentID are at the same location) If (transferSuccessProbAtSameLocation < r.nextValue) Unlink(agentID,e); Link(req.sender,e); End; Else If (transferSuccessProbAtDifferentLocation < r.nextValue) Unlink(agentID,e); Link(req.sender,e); End; End; Remove_request_records(req); Case Others_request_passing: Request req = find_request_records(agentID's received delivery request); interactionPartnerID = pick_one_agent_with_the_element _based_on_the_transactive_memory(agentID, Neighbor_agents); If (interactionPartnerID has done his interaction for this turn) finish this block; Request newReq = new Request(req.element, agentID); Put_in_the_request_list(interactionPartnerID,newReq); Case My_request_generation: Element e = find_required_element_not_in_possession(agentID); Request req = new Request(e, null); Put_in_the_request_list(agentID,req); End switch; transactiveMemoryExchangePartnerID = pick_one_agent_randomly(Neighbor_agents); exchangeTransactionMemory(agentID, transactiveMemoryExchangePartnerID); End function;	
--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--

10.4.3 Knowledge of space

At the end of the social interaction, the agent exchanges his transactive memory with a randomly selected neighboring agent. This is a simulation of interactions passing the information about the current simulated situation. The transactive memory element exchanges are done as the following pseudo code. This exchanged transactive memory becomes the basis for agent social behavior: finding required elements, finding interaction partners, etc. The key, as can be seen in Table 10.10 is that agents who are not proximal are not finding each other as interaction partners and interacting. Future work would refine this by allowing for different types of communication some of which is proximity based.

TABLE 10.10 Agent's transactive management pseudo code

```
Function exchangeTransactiveMemory(int agent1ID, int agent2ID)
  Agent1_Neighbor_nodes = getSphereOfInfluence(agent1, one social link away);
  Agent2_Neighbor_nodes = getSphereOfInfluence(agent2, one social link away);

  N = (number of transactive memory elements exchanged);
  For i = 1 to N
    Agent1_neighbor_node = randomly_pick_one_node(Agent1_Neighbor_nodes);
    Put_transactive_memory_tuple( agent2ID,
      new TransactiveMemoryElement(agent1ID, Agent1_neighbor_node));
  End;
  For i = 1 to N
    Agent2_neighbor_node = randomly_pick_one_node(Agent2_Neighbor_nodes);
    Put_transactive_memory_tuple( agent1ID,
      new TransactiveMemoryElement(agent2ID, Agent2_neighbor_node));
  End;
End function;
```

10.4.4 Task execution logics

Finally, the agents perform task execution behavior. A task is not ready to be executed if the task's prerequisite tasks are not done yet. A task is ready to be executed if the task's prerequisite tasks are done. A task is done if the group of assigned agents has all the required resources, knowledge and is placed at required locations. This task execution model is described in Table 10.11.

TABLE 10.11 Agent's task execution implementation pseudo code

```
Function Perform_task(int agentID, Random r)
  task_list = getSphereOfInfluence(agentID, one social link away, only_task_nodes);
  ready_task_list = select_only_ready_task (task_list);
  task_to_execute = randomly_pick_one_task_that_all_required
    _elemets_are_gathered(ready_task);
  If ( taskExecutionSuccessRate < r.nextValue )
    recordTaskIsDone(task_to_execute);
  End;
End function;
```

10.4.5 Link to the previous description

Figure 10.11 shows which simulation flowchart components correspond to which pseudo codes in the previous sections. The simulation process is managed by the simulation model main loop, Table 10.6, and the simulation iteration function, Table 10.7. In the simulation iteration function, each agent is called in the randomized order, and the agent executes three aggregated behavior patterns. The first behavior pattern is the social interaction that is implemented as Agent's social interaction implementation pseudo code in Table 10.9. Then, the second pattern is the task execution implemented as Agent's task execution implementation pseudo code.

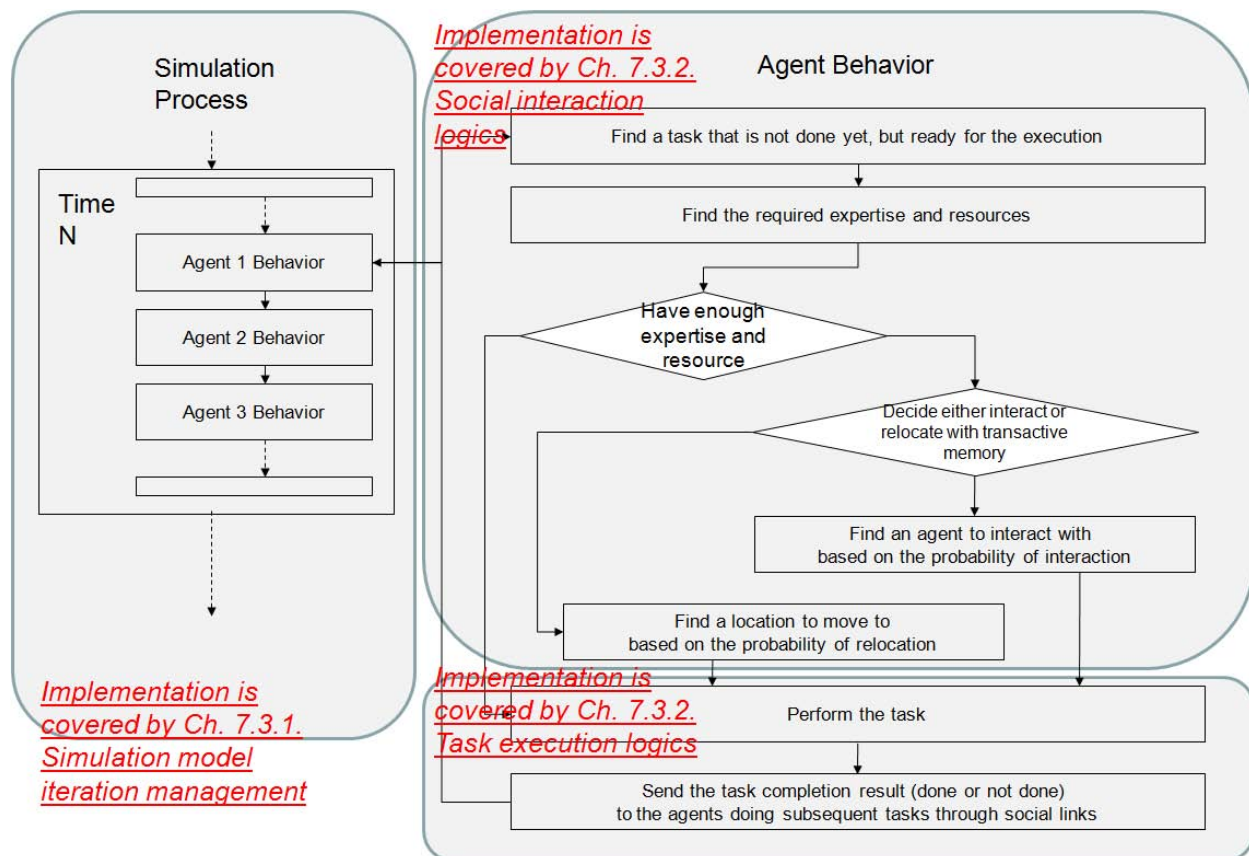


Fig. 10.11 Annotated simulation procedure flow chart. The annotation specifies which items in the flow chart correspond to the pseudo code.

Table 10.12 contains a list of key parameters and how these parameters are utilized by the pseudo-code (Tables 10.6 through 10.11). These parameters were introduced earlier. However, the earlier introductions were about their types and implications. Table 10.13 provides the links between the pseudo code functions and the used parameters. As such, it provides information about where the user parameters are used in which part of the simulation model.

10.4.6 Virtual Experiment Design

The experimental design used here is similar to that of the previous example. The geo-analysis virtual experiment is described in Table 10.13. The control variables are:

- How targets are selected
- The number of targets removed (size of intervention)
- The timing of the removals (intervention timing)

Targets are selected based on one of four factors, each of which is a measure of node centrality. These four factors are: a) degree (number of connections an agent has), b) betweenness (the number of shortest paths that run through this agent), c) eigenvector centrality (how connected is this agent to well-connected agents), and d) cognitive demand (multi-mode measure considering an agent's connection to resources, skills, as well as other agents). The size of the intervention is another factor in the experimental design – top ranked agents up to the number required are removed (so, in the largest case, the top twelve agents for the given selection scheme are re-

moved). When the interventions take place is another factor, the interventions occur from 5 to 40% of the way through the simulation's time-course.

TABLE 10.12 A table describing the key parameters in the simulation and the implication of setting these parameters

Name (Default value in the parenthesis)	Which pseudo code function uses the parameter	Implication
Boost for interaction if two agents are co-located (1.5)	Social_interaction, (Table 10.9)	If two interacting agents are co-located, the agents will have higher chances of transfer success.
Boost for removal recognition if two agents are co-located (1.5)	Social_interaction (Table 10.9)	If an agent tries to recognize one removed agent at the same location, the agent will have higher chance in recognizing the removed agent.
Number of time-step (5000)	Main (Table 10.6)	The number of simulated time-steps
Weights for requested element delivery (0.33), others' request passing (0.33), or the agent's request passing (0.33)	Social_interaction (Table 10.9)	Only used in task performance agent interaction model. Weights for selecting an agent interaction purpose. An agent selects one purpose out of three, requested organizational element (expertise or resource) delivery, his required element request to others, or passing others' request to different others.
Interaction count for time-step (3)	Social_interaction (Table 10.9)	An agent cannot interact with another agent after this maximum interaction count.
Cognitive power for time-step (3)	Social_interaction (Table 10.9)	An agent can only respond to the number of interactions specified by this parameter.
Exchange success rate (0.75)	Social_interaction (Table 10.9)	If an agent diffuses information or passes a resource to another agent, there is a success rate of such trials.
Interaction social distance radius (1)	Social_interaction (Table 10.9)	Interaction candidates are limited to agents who are within N social link radius from the interaction initiating agent.
Task execution success rate (0.5)	Perform_task (Table 10.11)	When an agent performs a task, the agent can accomplish the task with this success rate. If the task is not ready (the ready state is elaborated later), an agent cannot perform the task.

TABLE 10.13 Virtual experiment design for simulation parameters (30 replications, 2500 simulation time-steps)

Name	Value	Implication
Removal target selection scheme	Degree, Betweenness, Eigenvector centralities and Cognitive Demand (4 cases)	Agents with high network values are considered critical, and their removal should impact the organization. This is how we pick target agents to remove.
Intervention size	1, 5, 9, and 12 agent removals (of 19 agents total, 4 cases)	The intervention size specifies how many agents to remove with this intervention.
Intervention timing	125, 250, 500, and 1000 time-step (removing at after 5%, 10%, 20% and 40% timeflow, 4 cases)	The intervention happens at a specific stage of simulation period.
Total virtual experiment cells: 64 cells (4x4x4 cases)		

10.4.7 Results

The question we are trying to examine is how large, how early, and what kinds of agents are the best to remove to significantly impact the performance of an adversarial organization.

Results differ slightly from the social-only model previously discussed. Table 10.14 is the regression analysis, which compares end-run (after all 2500 turns have concluded) performance for each simulation. As could be predicted, the size of the intervention has the most drastic effect on the organization's performance. Larger interventions impact performance more. Earlier interventions tend to have much larger impacts than late interventions. There is no clear trend among the various selection schemes, although there is some indication that different selection criteria should be used for different mission objectives.

Table 10.14 A table of standardized coefficients for regression to the six organizational performance metrics at the end time using the virtual experiment settings (treating removed agent selection scheme with four categorical values) (N=64 cases) (* for P<0.05)


Standardized Coefficient	Mission Speed	Task Speed	BTA	ETA	Diffusion	Task Completion
Intervention Timing	0.521*	0.355*	0.470*	0.129*	-0.142*	0.602*
Intervention Size	-0.663*	-0.337*	0.013	-0.981*	0.975*	-0.630*
Degree Cent.	-0.058	0.142	0.036	0.040	-0.027	0.002
Betweenness Cent.	0.053	0.073	0.092	0.018	-0.023	0.050
Eigenvector Cent.	-0.008	0.090	0.017	0.013	-0.010	-0.012
Cognitive Demand	0.000	0.000	0.000	0.000	0.000	0.000
Adjusted R-Square	0.678	0.171	0.144	0.960	0.952	0.725

Figure 10.12 shows the organizational performance over-time. In terms of Energy Task Accuracy and Diffusion, removing more terrorists increases the inflicted damage (as can be seen in the largest interventions). However, from the task completion perspective, early interventions


are more important and their impacts to the organization are prolonged. This suggests that the intervention tactic should be adjusted according to the objective of the intervention. If a human analyst wants to stop the spread of expertise and aligning resource distributions, then the analyst should focus on removing more agents. If the analyst wants to prevent an event occurring, then the analyst should focus on removing agents earlier.

Additionally, it should be noted that the late interventions (removing agents after 40% of time-steps) and the small interventions (removing only 10% of agents) do not make significant damage in task completion over time. Therefore, such interventions do not result in the disruption of adversarial task performance and should be avoided.


Further, patterns of movement can be described. The training center in Somalia and the resource center in Afghanistan attract agents (with Pakistan serving as a go-between) until necessary skills and resources have been acquired. Operations occurred in both Kenya and Tanzania and as actors gained training and resources, they tended to move to one of these two target zones.




(a)



(b)



(c)



(d)

Fig. 10.12, a, b, c, d Changes in task metric performance due to interventions.




Fig. 10.13 Agents gathered resources and skills and then moved to operational centers.

10.5 Conclusion

Dynamic Network Analysis (DNA) has a wide number of applications vis-a-vis reasoning about adversarial behavior. Three such applications have been described. The key to reasoning about the adversary is taking social networks and embedding them within the spatio-temporal context. Organization theory and task processing analysis facilitate this embedding by providing the constraints and enablers on task-related activity.

Future work should build on this spatio-temporal-network reasoning. Two avenues we expect to be particularly valuable are: 1) hierarchical modeling where the adversaries and blue-force are modeled simultaneously at the agent and the group/organizational level, and 2) multi-model analysis where diverse models are run for the same geo-temporal-network context and the results are used to inform each other.

PART IV: MULTI-MODELING AND META-MODELING

Chapter 11: Introduction to Multi-modeling and Meta-modeling

Chapter 12: Meta-modeling for Multi-modeling Interoperation

Chapter 11

Multi-modeling and Meta-modeling of Adversaries and Coalition Partners

Alexander H. Levis, Lee W. Wagenhals, Abbas K. Zaidi, Tod Levitt

11.1 Introduction

No single model can capture the complexities of human behavior especially when interactions among groups with diverse social and cultural attributes are concerned. Each modeling language offers unique insights and makes specific assumptions about the domain being modeled. For example, social networks [58] describe the interactions (and linkages) among group members but say little about the underlying organization and/or command structure. Similarly, organization models [59] focus on the structure of the organization and the prescribed interactions but say little on the social/behavioral aspects of the members of the organization. Timed Influence net models, [13], [60] a variant of Bayesian models, describe cause-and-effect relationships among groups at a high level.

In order to address the modeling and simulation issues that arise when multiple models are to interoperate, four layers need to be addressed (Fig. 11.1). The first layer, Physical, i.e., Hardware and Software, is a platform that enables the concurrent execution of multiple models expressed in different modeling languages and provides the ability to exchange data and also to schedule the events across the different models. The second layer is the syntactic layer which ascertains that the right data are exchanged among the models. The Physical and Syntactic layers have been addressed through the development of two testbeds: C2 Wind Tunnel (C2WT) [61], [62] by Vanderbilt University in collaboration with UC-Berkeley and George Mason University (Appendix E) and SORASCS developed by CASOS at Carnegie Mellon University (Appendix F). Both have been used and developed further in this project.

Once the testbeds are available, a third problem needs to be addressed at the Semantic layer, where the interoperation of different models is examined to ensure that conflicting assumption in different modeling languages are recognized and form constraints to the exchange of data. In the Workflow layer valid combinations of interoperating models are considered to address specific issues. Different issues require different workflows. The use of multiple interoperating models is referred to as *multi-modeling* while the analysis of the validity of model interoperation is referred to as *meta-modeling*. Such an approach has been used in simulation mode or to explore the possible outcomes of proposed courses of action; it has not been used to predict outcomes.

In this chapter, we focus on issues relating to the syntactic and semantic layers. First, in sections 11.2 and 11.3 the concepts of multi-modeling and meat-modeling are explored. Both are subjects on which some basic research has been done but much more needs to be done. The modeling languages currently implemented in the C2WT are described briefly in section 11.4. In section 11.5, the approach taken for the meta-modeling analysis is presented. Finally, in Section 11.6, the approach is illustrated through a complex scenario that involves Intelligence and Surveillance in order to defeat adversaries from using IEDs and developing weapons of mass destruction (WMD).

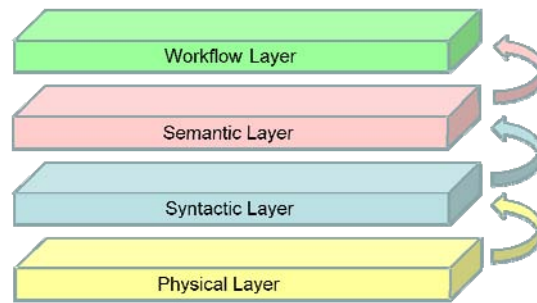


Fig. 11.1 The four layers of multi-modeling

11.2 Representation Issues in Meta-Modeling for Multi-Modeling

When we work in multi-modeling we actually are working, either implicitly or explicitly, over at least five levels of abstraction:

1. Meta-Modeling Language
- 2a. Meta-Model of the Domain Application Area
- 2b. Meta-Model of the Model Type
3. Model Type
4. Model Type Specialization to Domain Application Area

Typically only level 3, the model type, is explicitly specified. For example we may build an influence diagram by using the Pythia tool (Appendix B)**Error! Reference source not found.**, and even though it is being done to analyze military courses of action (CoAs) there is no explicit representation of a CoA as a class, nor are the actions in the CoA called out as being of class “action”. Rather the CoA generation is a method that outputs a sequence of instances of the class “proposition”, because the meta-model of the model type, “influence network” is about links that are of class “influence” between nodes that are of class “proposition” as shown in Fig. 11.2. Furthermore, the meta-model of Fig. 11.2 is typically only kept in mind of the researcher, rather than being explicitly represented.

Meta-modeling for multi-modeling has even greater complexity because of the distinction between knowledge representation of the world for domain-specific reasoning and knowledge representation of software for building and executing domain-specific models. Figure 11.3 pictures the issues.

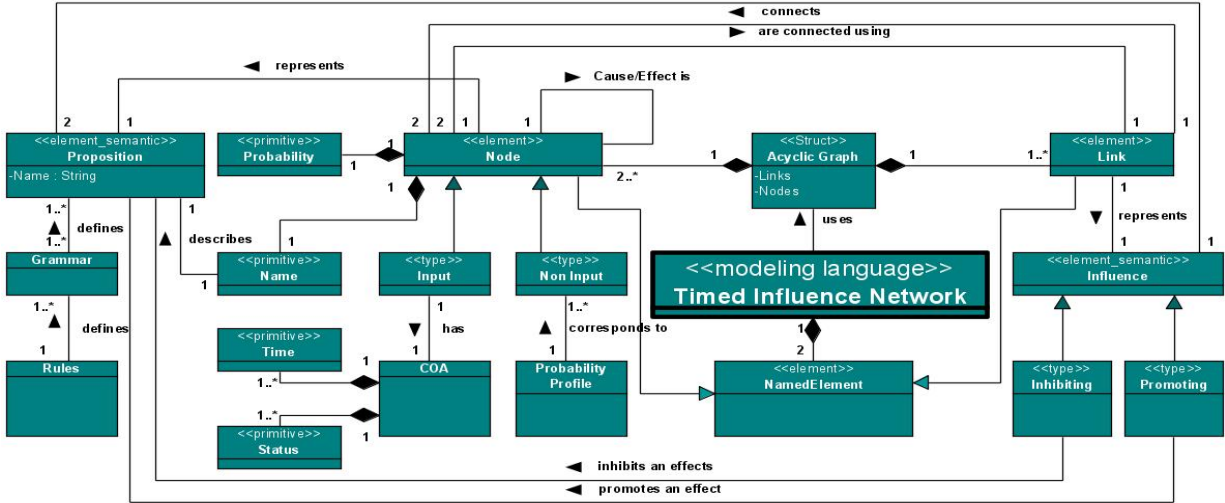


Fig. 11.2 Influence Network meta-model

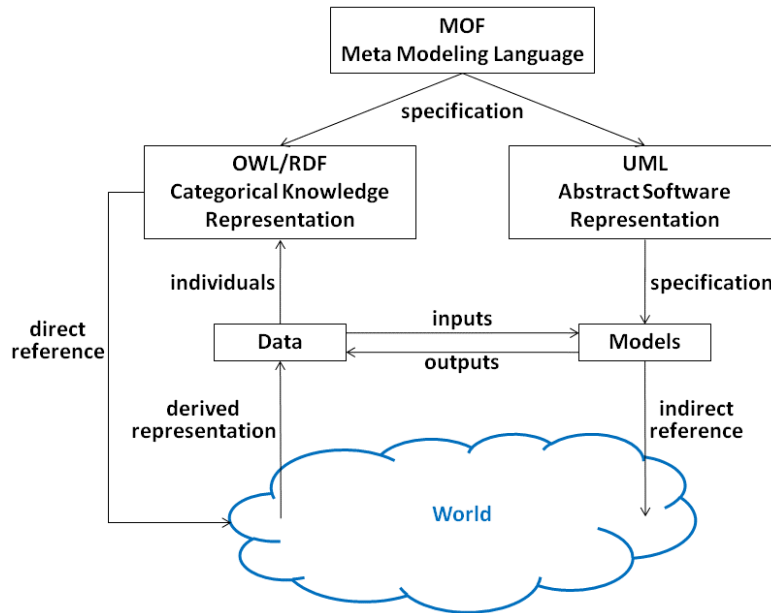


Fig. 11.3 Representation of knowledge and software

The distinctions between knowledge representation for reasoning about things in the world versus that needed to build scalable, maintainable software has been sharpened over the last decade by efforts to develop software standards supporting the World Wide Web (WWW) [63]

Historically the Unified Modeling Language (UML) [64], among others, was developed in response to the need for abstract representations and specifications for building maintainable, scalable software.

However the need to search and reason about information available on the WWW pushed the distinction between data as knowledge for general use versus data as inputs to computer programs and led to the development of formal ontologies, notably the Resource Description Framework (RDF) [65] and the Web Ontology Language (OWL) [66] for reasoning about knowledge independent of the goal of writing a specific software system. In addition knowledge man-

agement frameworks such as Protégé [67] were developed for working with knowledge representations independent of specific software models.

This has led to the necessity to unify and move agilely between OWL-like reasoning about knowledge and UML-like engineering of software, see, for example, [68]. Multi-modeling in particular requires this capability.

When we investigate the world, for example to learn about decision makers and decision making processes in a foreign country, we incrementally discover knowledge distinctions that cause us to augment and revise previous representations we might have developed. For example we might find that a Dr. Smith was involved in a debate on the use of nuclear weapons and then later find out that she is the Director of Nuclear Weapon Development for the country of interest.

From an ontological viewpoint we have either changed Dr. Smith's class representation or, more likely, added properties to her class representation as an agent including that she is a member of the government, and that she holds a specific office in the government.

In the OWL sense of ontological knowledge representation, we have specialized the class(es) to which Dr. Smith is a member, because a class represents the ontological category for Ms. Smith as an individual in the world, and therefore as a "resource" in OWL.

Although we could make such distinctions in the UML representation of Dr. Smith we have probably not done so unless it is necessary for the execution of a software model in which Dr. Smith appears as data. For example her communications on nuclear weapons might appear as propositions in an influence network, but the fact that she is a member of the government organization is not explicitly represented because it is not needed for any distinctions implicitly made by the inference methods that operate in the influence network model.

More generally OWL is about representing distinctions about the world that ultimately are observed as data, while UML is about representing distinctions in software that are used to build tools to process data.

This distinction appears in the fundamental capability in OWL to make logical deductions that can infer previously unrepresented properties about individuals observed in data, whereas UML has no comparable facility. From the software development viewpoint, such inferences would only be done by a tool that was built using a UML specification.

Conversely UML has capabilities to specify distinctions about data encapsulation within objects instantiated from class representations and to set default values whereas OWL does not because "class instantiation" and "data encapsulation" are not inherent properties of world entities. From the OWL perspective such representations would be done by custom-built facilities reasoning about world sub-domains (e.g. software itself) where such distinctions hold.

When we do multi-modeling we need both kinds of reasoning and representation. On the one hand we need to seed models with knowledge derived by an ontologically-driven search of available data sources. Then we need to use that observed and inferred data to do multi-modeling with different models, to better understand or solve a problem (in the world). Inter-operations between models that can realize such insights or solutions must be done in a formal software manner if they are to be effective, maintainable and re-useable.

Meta-languages have been developed to support inter-operation between OWL-like and UML-like reasoning, notably the Meta-Object Facility (MOF) [69]. MOF can be used to specify


representations in both OWL and UML. In fact MOF is used to represent MOF itself closing the implicit loop on meta-representation [70]. Kermeta [71] is a MOF extension language that enables not only specification but also software execution of meta-models.

11.3 Meta-Modeling Operations for Multi-Modeling (M4Ops)

In research performed so far we have identified five types of meta-modeling operations for performing multi-modeling.

- Concatenation: models share representations and so can get instances from each other
- Amplification: model adds or augments class representation from another
- Parameter Discovery: one model provides parameters for algorithms to another model's method
- Model Construction: one model is used to construct models of another type
- Model Merging: meta-model for a new model type is created that merges structure from one model with methods from another

In the case of concatenation because the models share the same class representations in question, no data translation needs to be performed. Concatenation is pictured in Fig. 11.4.




- **Same classes shared between models A and B**
- **New class instances sent from A to B**

Fig. 11.4 Concatenation

Amplification to date is performed manually by implicitly translating a representation in one model to another representation in another model. We use the notation A^M to mean the meta-model representing the structure and methods of model A.¹ In Fig. 11.5 we use the notation “meta-query” to mean the human interaction to query for ontological knowledge that might be resident in model B but relevant to model, A if an appropriate class re-mapping was invoked. For example a “decision-maker” in model B might be re-mapped to an “agent” in model A. In the formal meta-modeling operation we augment the meta-model in model A, i.e. A^M , a specialization of the class “agent” called “decision-maker” would be created. Thereafter data searches for instantiation of models of type A would include looking for “decision-makers”.


¹ It is technically necessary to distinguish the meta-model of a generic model of type A, such as a generic Petri net model, versus the meta-model of an instantiated model of the use of a Petri net to represent, for example, the membership and communication relationships in an organization. At this point in multi-modeling research we have not developed the technology far enough to merit making such distinctions in notation. It will be necessary in future versions of this report.



- **Meta-query requires knowledge not resident in model A**
- **Model A meta-model augmented by refactoring with B's meta-model**
- **A acquires new classes and new instances**

Fig. 11.5 Amplification


Parameter Discovery occurs when one model uses another as a supplier of method parameters. This can be done in a “concatenation mode” where the parameters produced by model B are of the same class as those required by the method of model A. However it can be that the parameters produced by model B require translation, or, alternatively, the parameters of the method of model A need to be augmented to match the classes of those in model B. (Fig. 11.6)



- **Query for parameters needed to execute A's method**
- **B performs a discovery process to find and/or compute parameters**
- **A can build new knowledge using its method that now has necessary parameters**

Fig. 11.6 Parameter Discovery


Model Construction addresses the case when one model serves as a factory for constructing models of a different type. For example a model that learns Bayesian networks from data is a constructor for Bayesian network models. Note that there could be multiple, different Parameter Discovery operations that provided data to such a constructor operation. For example there might be one Parameter Discovery operation that searches the WWW for data to feed a Bayesian network constructor model, and another Parameter Discovery operation that searches relational databases for parameters for the Bayesian network constructor model. (Fig. 11.7) This illustrates the reason for distinguishing the “operation” from the “model”. This distinction needs clarification.



- Model A is a “factory” for models of type B
- A can build multiple different instantiations of models of type B
e.g. learning models from data

Fig. 11.7 Model Construction

In “Model Merging,” a new model type is created that merges (part of) the class structure from one model with one or more methods from another model.



- Model A contributes a structure, such as a relation between agents
- Model B contributes a method, such as an inference algorithm
- A mixture of meta-models from A and B is used to construct a meta-model of a new model type, C, that can execute the method of B over the structure of A

Fig. 11.8 Model Merging

For example we can envision developing an organizational membership model (a type of social network model) with a probabilistic inference method (from a Bayesian network model type) in order to create a model in which one can probabilistically project the growth/decay or general health status of the organization.

11.3 Multi-Modeling

Current military operations need, and future operations will demand, the capability to understand the human terrain and the various dimensions of human behavior within it. Behaviors in the human terrain context extend across the spectrum from adversaries to non-combatant populations, to coalition partners, and to government and non-government organizations. As the type of missions that current and future commanders must address has expanded well beyond those of traditional major theater combat operations, the need to broaden the focus of models that support planning and operations has become critical. Actions taken by all agents together with the beliefs, perceptions, intentions, and actions of the people involved in an area of operations, interact

to affect the outcome of a conflict or coalition operation, a disaster relief plan, and/or a peace-keeping effort. No single set of models and tools can support the operational commander addressing the challenges of conducting non-conventional warfare missions. For example, while there are many models using diverse data bases, none can address the complexities of coordinating kinetic and Information Operations when the adversary is embedded within a complex non-combatant population.

The MURI team at George Mason University and Carnegie Mellon University [72] have developed a suite of models that address different aspects of modeling the attitudes and behaviors of adversaries. While they address man-made threats, they are proof of the concept that model interoperation is feasible. A subset of the currently available suite of tools is shown in Fig. 11.9.



Fig.11.9 Modeling applications using different modeling languages

These models have been described in Parts II, III, and IV of this report. Consequently, only a brief description is given here. **CAESAR III** [73] is a Colored Petri Net tool for designing and analyzing organizational structures. The Timed Influence Net application, **Pythia** [38], [22] is used to develop Courses of Action and compare their outcomes. **Temper** [74] is a temporal logic inference tool that is used to address the temporal aspects of a course of action. **Ruler** [75] is a tool for evaluating whether a proposed course of action is in compliance with the prevailing legal and regulatory environment. **SEAT** is a tool for visualizing and comparing the results of the simulations with measures of performance and measures of effectiveness.

ORA [76] is an application for the construction and analysis of social networks while **DyNet** [77] is a computational model for network destabilization. In addition, **WebTAS** [78], a GFE visualization and timeline analysis tool developed by AFRL/RI that accesses data in data bases and can receive streaming live data form sensors, has been integrated in the C2 Wind Tunnel thus enabling showing data and results on maps. Each of these federates uses a different modeling language and a different simulation engine.

Multi-modeling can serve as a means of reducing data ambiguity and identifying missing data. For example, by comparing two distinct, incomplete models and using inferences from one model to inform the other, a much less ambiguous representation of the system of interest can be obtained. When this approach is coupled with meta-modeling research to determine what inferences from one model can be validly “exported” to another, the opportunity arises for exploring ambiguity, uncertainty, and missing data issues.

Effective multi-modeling requires that the modeling languages used, the models themselves, and the supporting data, do not contain assumptions that invalidate the specific model interoperation. This leads to the need for meta-modeling analysis.

11.4 Meta-Modeling

With the capability provided by the C2WT infrastructure, different models addressing the same problem can be utilized, all drawing upon the same data set. This provides a means for exploring the suitability of using models of different resolution and exploring their range of applicability through computational analysis and evaluation.

Our approach to understand modeling language semantics so that multiple models can be used together, i.e., can interoperate, has been to use concepts maps [79] to describe the characteristics of the set of modeling languages and data that are available to support analysis. A fragment of the concept map for the Timed Influence Net modeling language is shown in Fig. 11.10. A detailed description of the approach is given in the next chapter (Chapter 12).

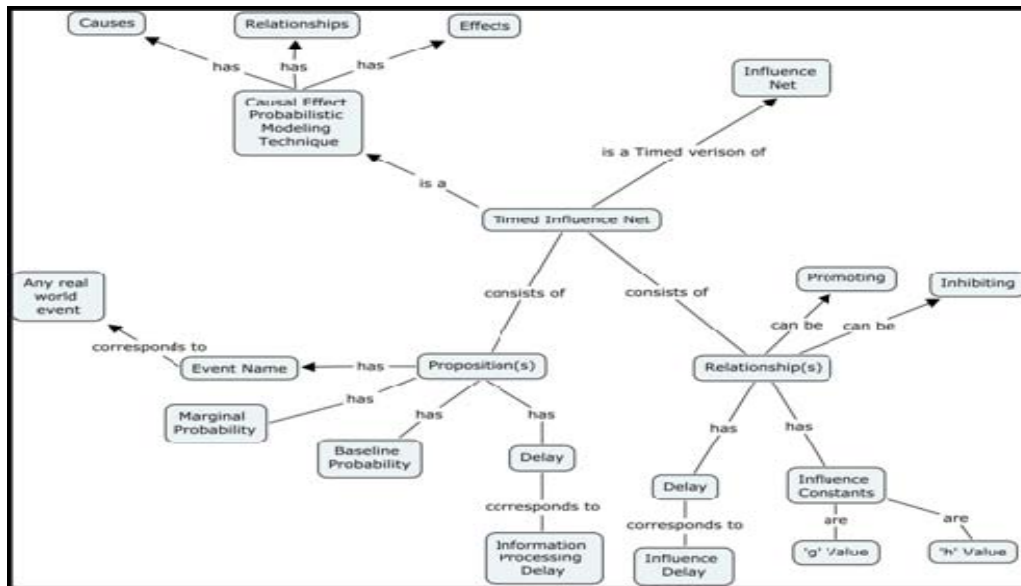


Fig. 11.10 Fragment of the concept map for Timed Influence Nets.

Meta Modeling analysis indicates what types of interoperation are valid between models expressed in different modeling languages. Note that model interactions can take a wide variety of forms: (1) One model runs inside another; (2) Two models run side-by-side and interoperate. The interoperation can be complementary where the two run totally independently of each other supplying parts of the solution required to answer the questions, or supplementary where the two supply (offline and/or online) each other with parameter values and/or functionality not available to either individual model; and (3) One model is run/used to construct another by providing design parameters and constraints or constructs the whole or part of another model (Fig. 11.11). These are all aspects of the need for *semantic interoperability*.

We assume that two models can interoperate (partially) if some concepts appear in both modeling languages. By refining this approach to partition the concepts into modeling language input and output concepts and also defining the concepts that are relevant to the questions being asked by the analysts and decision makers, it becomes possible to determine which sets of models can interoperate to address some or all of the concepts of interest, and which sets of models use different input and output concepts that are relevant to those questions.

In order to support semantic interoperability we must be able to interchange models across tools. This requires model transformations. The transformations are formally specified in terms of the meta-models of the inputs and the outputs of the transformations. From these meta-models and the specification of the semantic mapping we synthesize (generate) a semantic translator that implements the model transformation. [80]

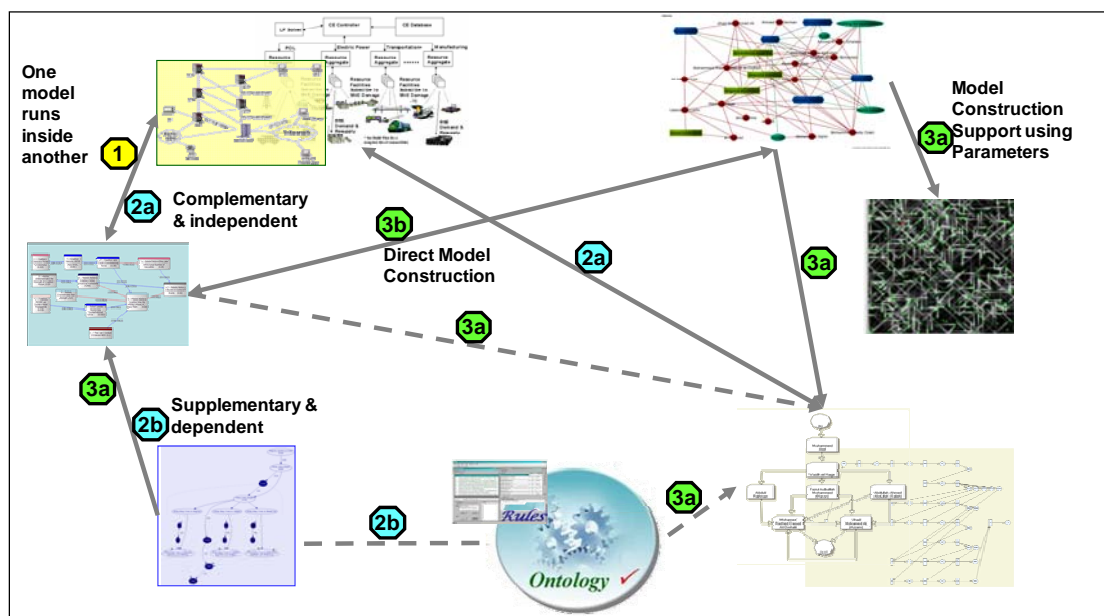


Fig. 11.11 Multiple types of model interoperation

11.5 Human Terrain Example

The C2WT technology has been applied to several different defense-related domains. However, to demonstrate the technology and the use of the C2WT as a platform for computational and mixed initiative experiments, a scenario was developed that required the inter-operation of several disparate models. The scenario is summarized below.

Blue suspects that Red is engaging in the development of WMD, but does not know where the development is being done. The key idea of the scenario is that Red has a covert house in which vehicle borne improvised explosive devices (VBIEDs) are manufactured and Red is also using the same house for the WMD development. Blue does not know that they co-located.

Blue has some limited intelligence about both of these activities and would like to locate the facilities. A shipment by truck of WMD materials provides an opportunity for Blue to discover the WMD factory. However, concurrently, Red decides to deploy a VBIED following his usual Tactics, Techniques and Procedures that have proven effective in the past. Blue has to make de-

cisions as to how to allocate scarce ISR assets (two Unmanned Air Systems, UAS) among the two Red activities while at the same time he wants to make the VBIED ineffective and gather additional intelligence about the VBIED making factory. A cyber cell was added to the Blue organization to provide additional ISR capabilities (SIGINT) for Blue.

The scenario was designed to illustrate the benefits of information sharing over networks. If different Blue operators and decision makers share the information that they collect from different resources in a timely manner, then the tactical and operational level commanders will be provided with the maximum information that, in the case of the scenario, results in Blue locating the combined WMD and VBIED factory and identifying and locating the VBIED operation. If some of the information is not shared for any reason, such as operator error by the Blue operator – the UAS pilots - or Red interference, Blue may be unable to identify the location of the covert bomb factory or track the VBIED to its final location.

GoogleEarth was selected as the scene visualization application because it is freely available and very realistic. Human operators may be used to control the two UASs; alternatively, the UASs may fly autonomously. The operators received directives from the Colored Petri Net representation of the CAOC operators and the information from the UASs was sent to those CAOC operators over the simulated network modeled in OMNeT++.

Figure 11.12 shows some of the screens generated by the execution of this scenario on the C2WT. One screen shows, using GoogleEarth, what the sensors on the UAS sees; the second shows what the human operator of the UAS sees after the sensor data are transmitted through the network. In this way, the effects of network delays or jamming or cyber exploits can be observable directly. Another screen shows the ground truth – the location and tracks of the UASs and of the targets as they drive through an urban environment. Still another screen shows the situation reports and other messages that go through the system. Finally, when the driver of the VBIED is identified, ORA generates the social network to which he belongs and thus helps identify other actors in the exploit.

The entire scenario with all federates was created and installed on multiple computers with several large screen displays. The execution of this particular scenario was conducted in real time because of the presence of the human operators and lasted about 20 minutes. Furthermore, the C2WT was instrumented for the collection of data so that excursions from the scenario could be evaluated. Two extensions were created to the main scenario involving Red's use of jamming to attempt to reduce the effectiveness of Blue. In the first case, the jamming was effective. In the second case, anti-jamming procedures were employed by Blue and the data links from the UASs to the CAOC were re-established. However, it was determined that the length of the time interval during which the signal was jammed was a critical parameter. If the interval was too long, then the UAS was not able to re-acquire the target in the urban environment. Currently, a series of experiments is being conducted with variants of this scenario exploring the effect of cyber exploits on performance.

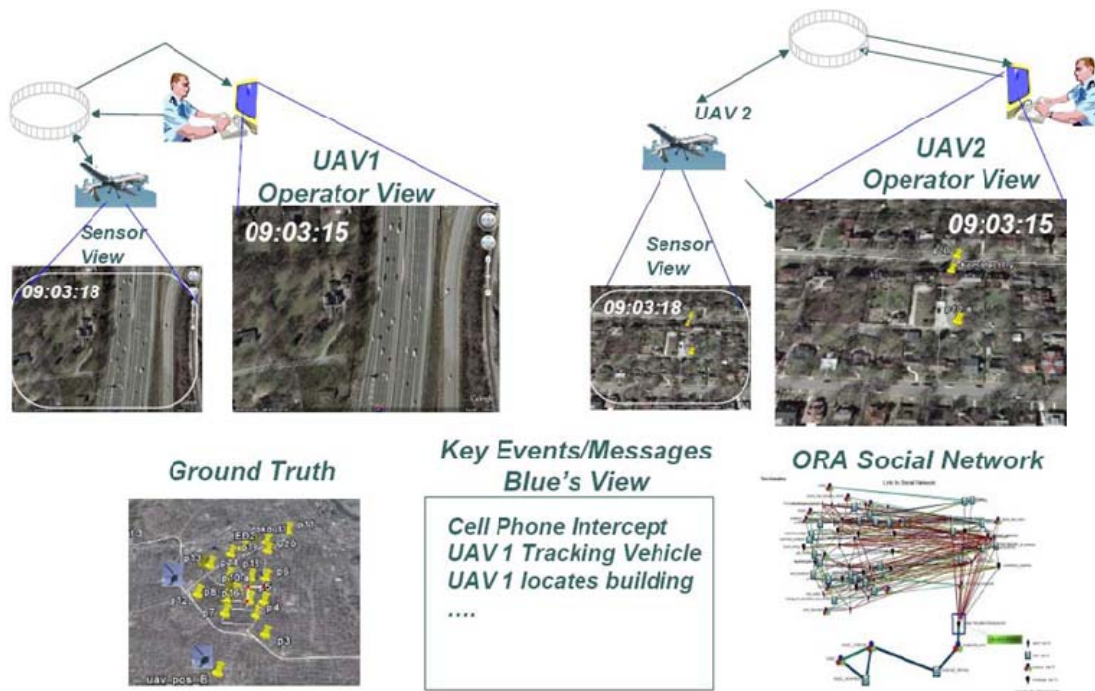


Fig. 11.12 Large Screen Displays for C2WT Demonstration

11.6 Conclusions

The concept of multi-modeling has promise for addressing the complexities of modeling adversary behavior especially when only incomplete data are available. This has led to the need for research on meta-modeling to ascertain that the interoperation of multiple models for a particular problem is valid. Such research is continuing while additional federates are being prepared for installation on the C2WT.

Chapter 12

Meta-Modeling for Multi-Model Interoperation

M. Faraz Rafi, Alexander H. Levis and Abbas K. Zaidi

12.1 Introduction

A model is an abstraction of a real phenomenon capturing some aspects of a domain of interest. The analysis performed using the model is therefore limited to only those aspects that are addressed in the model and the ability of the modeling language (used to construct the model) to allow the analysis (i.e., queries processed, metrics estimated, etc.). A model is generally constructed by a modeler (e.g., a subject matter expert or a knowledge engineer) who uses relevant (i.e., relevant to the problem being studied) data about the domain for which the model is being constructed and organizes it in a form conformant with the syntactical and semantic rules of the modeling language employed to construct it. Figure 12.1 depicts this process with a meta-model that alludes to these underlying notions of modeling language requirements and the relevant data needed for this model building exercise. In our use of this term, a meta-model is a precise definition of the constructs and rules needed for creating problem-specific semantic models. [81] The notion of a meta-model is a well understood term especially in the Software and System Engineering disciplines. [82], [83], [84]



Fig. 12.1 Model building overview

The modeling of a human conglomerate for the analysis of its behavior in response to external (e.g., social, economical, political, etc.) stimuli is a complex problem and requires development of an ensemble of several models. Each model, in this set, offers unique insights and makes specific assumptions about the domain being modeled that the model may or may not share with another model of the same domain. For example, Social Networks [58] describe the interactions (and linkages) among group members but say little about the underlying organization and/or command structure. Similarly, organization models [59] focus on the structure of the organization and the prescribed interactions but say little on the social/behavioral aspects of the members of the organization. Timed Influence Net models, [13], [85] a variant of Bayesian models, describe cause-and-effect relationships among groups at a high level. The idea of using a variety of techniques/models to solve a complex modeling and analysis problem is not a new one: an earli-

er survey and a collection of papers on the use of multiple *strategies* for machine learning problems can be found in Michalski and Tecuci [86]. Some examples of more recent works employing multi-modeling are: the use of multiple *specialists* in modeling human cognition in an approach called Polyscheme [87], [88]; integration of *effects* based (i.e., operational/strategic level) and *attrition* based (i.e., tactical level) models to study effects of the quality of the commodity services provided by an infrastructure on the socio-cultural attitudes and actions of the population of a small geographical region [89], [13]; the integration of First-order Logic with Bayesian probability theory in the form of an approach called MEBN [90], [91]; [92]; Interoperable Technosocial Modeling (ITM) which focuses on the integration and evaluation of human and physical models across diverging modeling platforms, e.g. Bayesian Nets and System Dynamics [93], [94], [95], and the interoperation between organizational decision models and socio-cultural trait models in Kansal et al. [96] and Levis et al. [73].

Figure 12.2 depicts a multi-modeling environment where each individual model is shown contributing to a piece of a larger puzzle (i.e., solution to the problem under consideration) with the analysis results obtained by it. The figure presents an over-simplification of possible inter-operations among the models in the ensemble employed to address a modeling and analysis problem. A variety of different types of interactions among disparate models have been reported in the literature. The following is a list of these types with some example citations:

1. A model-*i* uses another model-*j*, where each model is constructed using a different modeling language, to complete a computational/analysis task. For example, in Waghans et al. [13] a Civil Environment Model (CEM) is shown to provide inputs to a Timed Influence Net model developed for effects based planning.
2. Two or more models run concurrently and supply complementary parts of a solution. This case is also depicted in Fig. 12.2.
3. A model-*i* supplements the computational/analysis task of some other model-*j* by providing it with analysis results and/or parameter values. For example in Kansal et al. [96] a socio-cultural trait index is used to introduce design constraints, incorporating cultural attributes of a group of individuals, on the allowable interactions within an organization. The organization is modeled as a discrete-event system. In Zaidi et al., [27] and Haider et al. [97] a temporal modeling and reasoning approach is employed for a Timed Influence Net (TIN) model to supplement its probabilistic analysis with a temporal analysis capability.
4. A model-*i* is used to construct the whole/part of another model-*j*. For example, Moon et al. [98] propose an approach to automatically construct a Timed Influence Net (TIN) model from a Social Network and a rule model. A similar approach is presented in [99] that attempts to construct a TIN model from an ontological knowledge model of a domain.

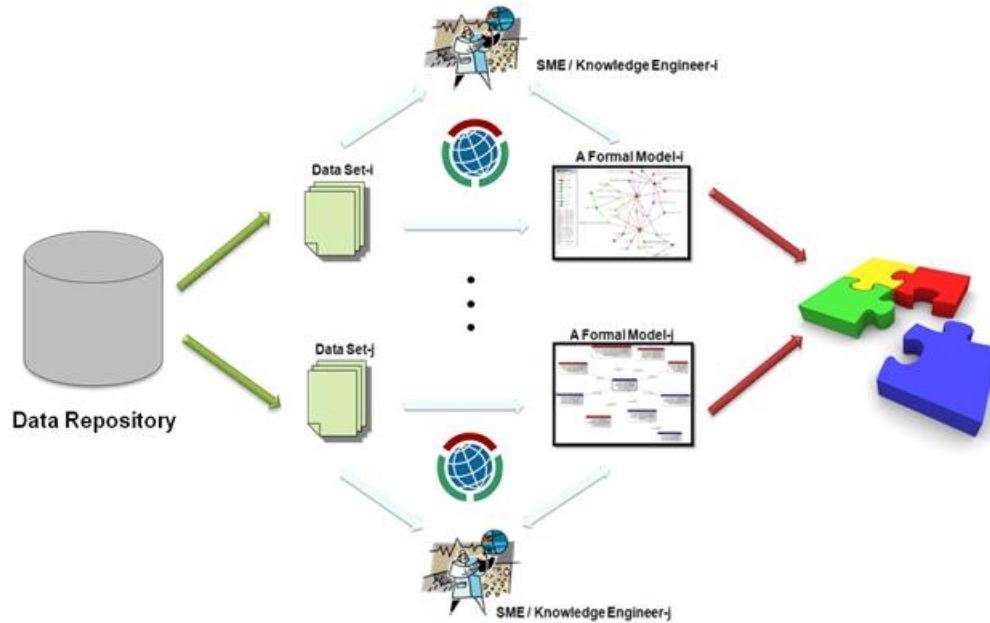


Fig. 12.2 A multi-modeling environment

The examples of interoperations, cited above, involve specific models of some domain that are linked together for some example application. Most of these approaches are ad-hoc in the sense that they do not provide a theoretical underpinning of the interoperations presented in them. They are valid for the example instances that are used to present them without a generalization that can be used to interoperate any two (or more) models expressed in the same modeling languages as the example models.

In this chapter, we present a framework for a formal study of syntactic and semantic interoperability of disparate models and modeling languages employed to address a specific problem of interest. The rest of the paper is organized as follows: Section 12.2 presents the Meta-modeling approach for the study of multi-model interoperation by formalizing the overlaps among models expressed in different modeling languages. The details of an application of the approach are presented in section 12.3. This application uses Social Networks and Influence Nets as the two candidate modeling languages for the development of a meta-model for possible interoperations between the two classes of models. Section 12.4 analyzes the meta-model developed in Section 3. Section 5 provides concluding remarks and directions for future research.

12.2 The Meta-Modeling Approach

A meta-model is an abstraction highlighting properties of a model constructed using a modeling language. A model conforms to its meta-model in the way that a computer program conforms to the grammar of the programming language in which it is written. The presented meta-modeling approach is an analysis of the conceptual foundations of a model ensemble so that individual models, constructed to address a specific problem in a domain of interest, can interoperate as part of a workflow developed to address the problem. This meta-modeling approach extends earlier works by Kappel et al. [100] and Saeki and Kaiya [101] for a class of modeling languages primarily used for behavioral modeling problems.

It is a phased approach that uses concept maps, syntactic-models, and ontologies. It is based on comparing and merging the ontologies (for each modeling technique) to help identify the similarities, overlaps, and/or mappings across the models under consideration. The data set populating each model in this ensemble may have been derived from a single large repository of information and may have overlaps (i.e., concepts, relationships, constraints, etc.) with the data sets required by other models. We assume that two models can interoperate (partially) if some concepts appear in both modeling languages. By refining this approach to partition the concepts into modeling language input and output concepts and also defining the concepts that are relevant to the questions being asked by the analysts and decision makers, it becomes possible to determine which sets of models can interoperate to address some or all of the concepts of interest, and which sets of models use different input and output concepts that are relevant to those questions.

Figure 12.3 provides an overview of the proposed approach. The underlying idea is to first develop separate meta-models for the different modeling languages employed, and then merging (or comparing) these individual models into a unifying meta-model for the ensemble. The merged meta-model identifies the similarities, overlaps, and/or mappings across the models under consideration that can then be used to formalize interoperation among these models in a workflow.

The approach starts by specifying a modeling language by constructing a generalized Concept map [102] that captures the assumptions, definitions, elements and their properties and relationships relevant to the paradigm. This is termed as the *Conceptual Modeling Level* in Fig. 12.3. This concept model is a structured representation, albeit not a formal one, and, therefore, not amenable to machine reasoning. The knowledge items, i.e., concepts and roles, captured at this level represent the *intensional knowledge* [103] about the problem domain. The reason for concentrating only on the general knowledge about the domain, as opposed to *extensional* or more specific knowledge about the instances in the domain, is to relate concepts and roles from one modeling paradigm to another. Such a mapping, if established, can then be applied to the extensional knowledge about a domain to instantiate both the models and the interoperations between them for a specific problem/domain of interest. The Concept map representation with the intensional knowledge is then formalized using syntactic model. The aim of constructing the syntactic model is to reveal the structural aspects of the modeling technique and to lay down the foundation for its ontology. This step is shown as the *Syntactic Modeling Level* in Fig. 12.3. A basic ontology, referred to as a *pseudo ontology*, is constructed which mirrors the syntactic model and serves as the foundation ontology; it does not contain any semantic concepts (related to the modeling technique and to the modeled domain) but acts as the skeleton for the ontology. In the next step, semantic concepts and relationships are added to this foundation ontology to obtain the *refactored ontology*. Once the individual ontologies are completed for each modeling technique, mapping of concepts across the ontologies is started. The resulting ontology which contains these concepts and relationships within and across multiple ontologies is called an *enriched ontology*. In our use of the terms, the pseudo, refactored, and enriched ontologies are all Description Logic (DL) [103], [104] knowledge bases, $O = (Tbox, Abox)$, with only *terminological axioms*. In our approach, terminological axioms are collection of concept axioms and role axioms expressed in OWL. An ontology, as defined here, has semantics similar to a knowledge base expressed in first-order predicate logic with implicit knowledge that can be made explicit through inferences. [103] The enriched ontology so constructed for several modeling languages can therefore be reasoned with using the logical theory supporting the ontological representation (i.e, OWL). The

inferences applied to such an ontology can verify the consistency of concept definitions and classify concepts with respect to subconcept-superconcept relationships. The inference mechanism, when applied to the enriched ontology, therefore, can identify these mappings among concepts across different modeling languages. The mappings suggest possible semantically correct ways to ensure consistency and to exchange information (i.e., parameter values and/or analysis results) between different types of models when they are used in a workflow. The steps of constructing the pseudo, refactored and enriched ontologies are carried out as part of *Ontological Modeling Level* in the proposed approach (Fig. 12.3).

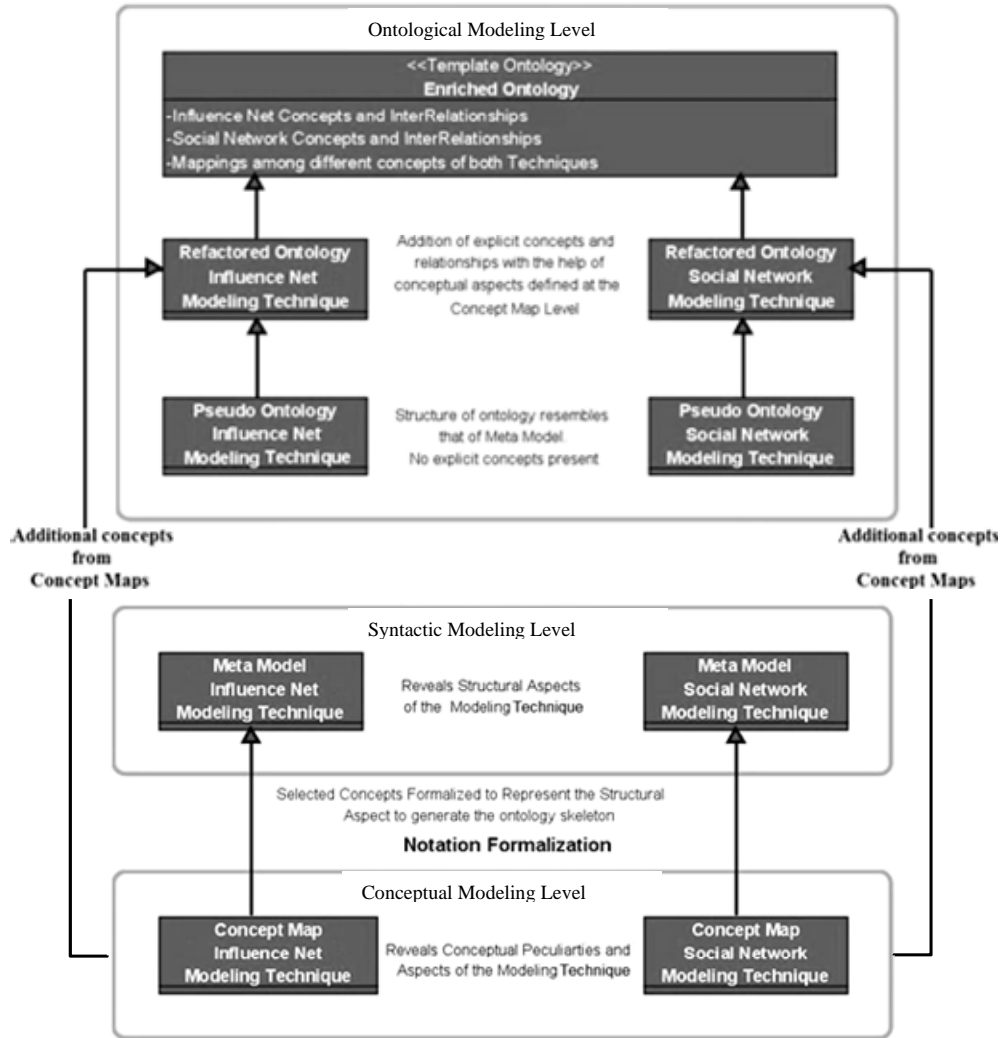


Fig. 12.3 Overview of the meta-modeling approach

12.3 Application

In this section, we present an application of the proposed approach of Fig. 12.3 to explore interoperation between two modeling languages for behavioral modeling problems. The two modeling languages selected for this study are Social Networks [58] and Influence Nets [13]. Social Net-

works describe the interactions (and linkages) among group members. Influence Net models, a variant of Bayesian Networks, describe cause-and-effect relationships among groups at a high level. A considerable number of models for real world scenarios have been developed and analyzed using both techniques. This served as one of the main reasons for selecting these techniques for the study presented in this paper. A brief description about each technique follows:

Influence Nets

Based on two well established techniques, i.e. Bayesian inference net analysis and Influence diagramming technique, Rosen and Smith [5] proposed a formalism called Influence net to perform probabilistic modeling in order to model the rationale of some group or organization. In an Influence Net model, the nodes represent random variables (propositions) such as beliefs, actions, events, etc., whereas an edge represents a causal relationship (influence) between two nodes (propositions). The parent and child nodes are often called *cause* and *effect*, respectively. The causal relationship between a cause and an effect can either be promoting or inhibiting as identified by the edge attributes as shown in Fig. 12.4. In Fig. 12.4, *Event A* has an inhibiting influence on *Event B* and a promoting influence on *Event C*, similarly *Event B* has a promoting influence on *Event C*. A UNIX based application *SIAM* supports the development of Influence net models. Another Window based tool called *Pythia* supports modeling of the timed version of Influence Nets called *Timed Influence Nets (TINs)*.

Social Networks

The Social Network definition used in our study is of Carley [58], where a Social Network is a structure composed of real world entities and associations among them. In this definition, a node, or an entity, can be an agent, organization, action, knowledge, and/or resource. In the sample Social Network shown in Fig. 12.5, the circular nodes represent entities such as human beings and the edges connecting these entities represent associations (e.g., relationship) between them. These associations can be an interaction between *Albert* and *Cynthia* in the form of mutual discussions; it can also be a kinship relationship between *Branden* and *Debby*. The graphical form of the social network can also have a matrix representation in which the entities are represented in the matrix rows and columns and the matrix entries indicate their interaction. ORA, a Social Network analysis tool, supports constructing these matrices and the models. These matrices can either be single-mode or multi-modal. Single-mode matrices represent networks containing only one type of entities (e.g., people or agents only) while multi-modal matrices consider networks with multiple types of entities (e.g., agents, action, organizations, knowledge etc.). These matrices collectively make up a meta-matrix, a framework that integrates multiple and related network matrices into a single interrelated unit as defined by Carley [58].




Fig. 12.4 Example Influence Net




Fig. 12.5 Example Social Network

The following is a detailed, step-by-step, description of the approach (Fig. 12.3) when it is applied to the two modeling techniques, i.e., Social Networks and Influence Nets.

Conceptual Modeling Level

The first step, in our approach, is the construction of Concept maps for each modeling language under consideration. In a Concept map, a concept is represented using some type of geometrical shape (rectangular, circular, elliptical etc.) and is connected with other concepts using a directed link. This link can be tagged with a description of the relationship between the two concepts. Concepts connected together with a relationship referring to a meaningful entity define a proposition, for example, *Influence Net is composed of Nodes and Links*. Concept *Influence Net* is connected with concepts *Node* and *Link* with a relationship *is composed of*. The aim of constructing a Concept maps for both Social Networks and Influence Net is to gain a syntactic and semantic, albeit informal, insight into both modeling techniques to reveal aspects which will ultimately facilitate the ontology construction process later on. The construction of a Concept map is an iterative process that requires brainstorming and frequent revisions until a final and concrete Concept map is developed. The steps of constructing a Concept map include *identification of focus questions*, *construction of parking lot* (pool of concepts) and *establishing cross links* between these concepts.

Figure 12.6 shows a fragment of a Concept map constructed for Influence Nets in an attempt to address the following focus question: *What are the constructs of an Influence Net?*

Syntactic Modeling Level

After the conceptual modeling level, only selected concepts from it are formalized to represent the structural aspects of the modeling techniques in the form of a syntactic model. A syntactic model is an abstraction layer above the actual models and can be considered a meta-model describing the syntactical rules and requirements for the construction of an instance model. The objective at this level is to retrieve a basic ontology skeleton. Syntactic models do not contain any detailed concepts and relationships of the domains they model, but their structure can be used as the basis for the first ontology to be constructed in the next level. Figure 12.7 shows the syntactic model for the Influence Net modeling language.

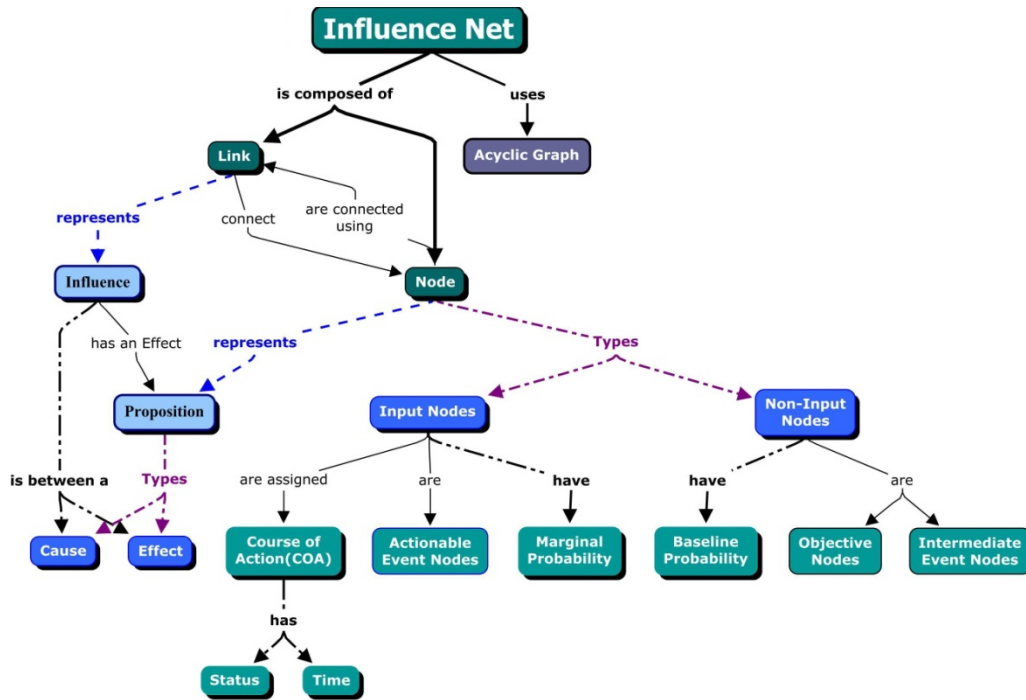


Fig. 12.6 A sample Concept Map for constructs of Influence Net focus question

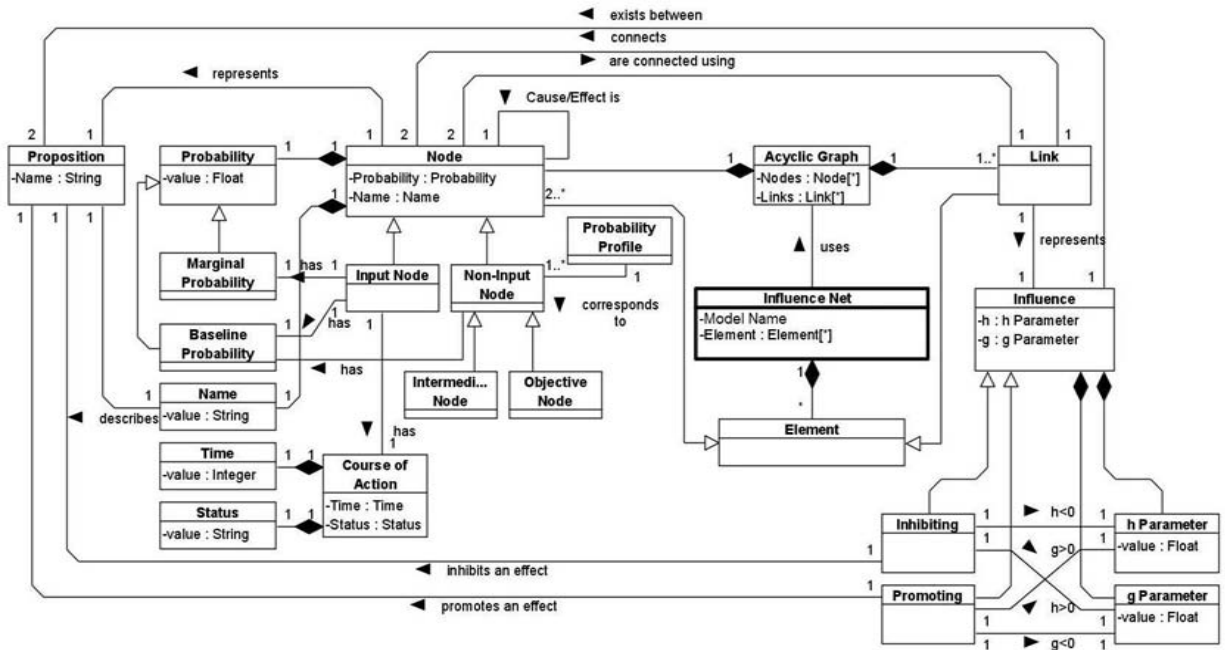


Fig. 12.7 Influence Net Syntactic Model

Ontological Modeling Level

There are three sub-levels in this step (Fig. 12.3) that eventually yield an ontology enriched with concepts and relationships from both modeling languages. Kappel et al. [100] refer to the process of formalism shifting as reducing the gap between the implementation oriented focus of syntactic models and the knowledge representation oriented focus of ontologies. This formalism shift is

led by the syntactic model developed in the previous level. We have adapted the technique suggested by Kappel et al. [100] with an addition of the *Conceptual Modeling Level* as the first attempt in understanding the underlying syntactic and semantic constructs of a modeling language and the generic domain concepts that are used in the models. In our study, we manually asserted some semantic equivalences between refactored ontologies, as opposed to Kappel et al's use of COMA++ tool.

At this level, the first ontology that is constructed is called *pseudo ontology* and resembles its syntactic model equivalent. Explicit domain concepts and relationships used by modeling technique are added into the pseudo ontology to construct a *refactored ontology*. Mapping of concepts between the refactored ontologies of both techniques is partly done manually and partly by invoking the DL reasoner to construct an *enriched ontology*. This enriched ontology contains the individual and mapped concepts and relationships of both modeling techniques. It can be considered as a template ontology which contains the *intra* and *inter-modeling technique* concepts and relationships. An ABox of this ontology can be instantiated for a specific domain which will serve as the knowledge container for that domain.

Pseudo Ontologies

In the first sub-level, the pseudo ontology resembling that of its corresponding syntactic model has no explicit concepts present except the ones related to the structure of the modeling technique similar to what is in the syntactic model. The construction of this ontology utilizes the syntactic model in such a way that each syntactic model element becomes an ontology class and the association between the syntactic model elements becomes either an object or a data property in the pseudo ontology. Figure 9 illustrates the pseudo ontology for the Influence Net modeling technique. For the sake of brevity only a small fragment of the ontology is shown.

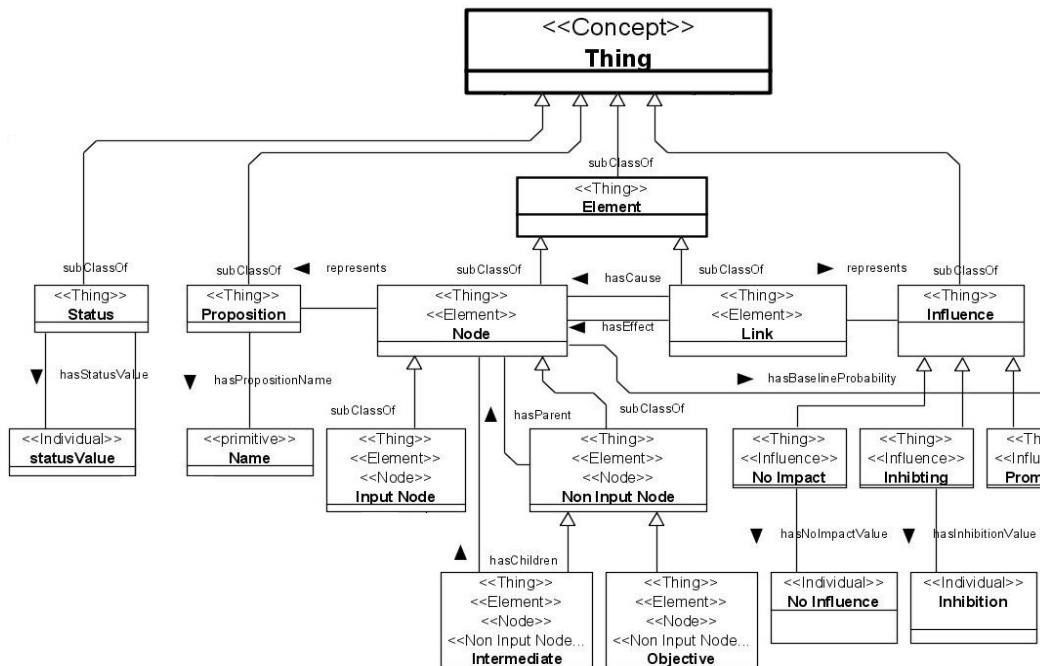


Fig. 12.8 Influence Net pseudo ontology snippet

Refactored Ontologies

Once the pseudo ontology construction completes, the knowledge obtained from the Concept maps, as well as the knowledge of the designer about the modeling technique, is used to define and feed new concepts (classes) and relationships (properties) explicitly into the pseudo ontology to construct a completely specified refactored ontology. Our approach adds additional concepts into the refactored ontology from the concept mapping phase and the domain knowledge of the ontology designer. Since during the concept mapping phase, a large number of concepts were already identified, the ontology designer can make use of that repository of concepts to feed into the pseudo ontology.

The *class hierarchy views* of the refactored ontologies developed for Influence Nets and Social Networks are shown in Figure 12.9 and 12.10. These diagrams were generated using GraphViz plug-in in Protégé 4.0 which graphically describes how the classes are associated (i.e., subclass, superclass, equivalent class) with each other based on the defined object properties (i.e., terminological axioms). Table 12.1 shows the concepts imported into the Influence Net refactored ontology from the Concept maps. Similarly, Table 12.2 shows the explicit concepts added into the refactored ontology of Table 12.1. Classes *subject*, *object*, and *verb* form the constituents of *Proposition* class. The classification of propositions is a critical step. There can be many types of propositions besides the ones shown in the Figure. The actual aim is to come up with an initial set of proposition types that are the most recurrent among the two types of models studied. Since ontology construction is an iterative process, its evolution can continue over time. As new types emerge, they can be incorporated as new classes and properties into the existing ontologies. Table 12.3 shows the concepts imported from the concept maps developed for Social Networks into the Social Network refactored ontology.

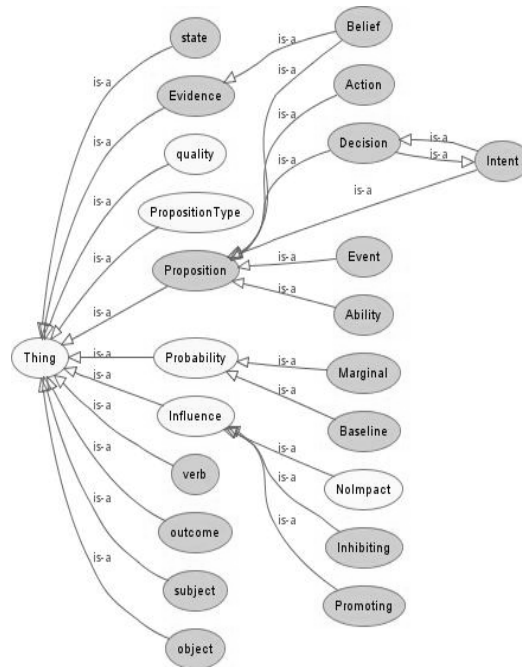


Fig. 12.9 GraphViz Diagram - Influence Net inferred refactored ontology.

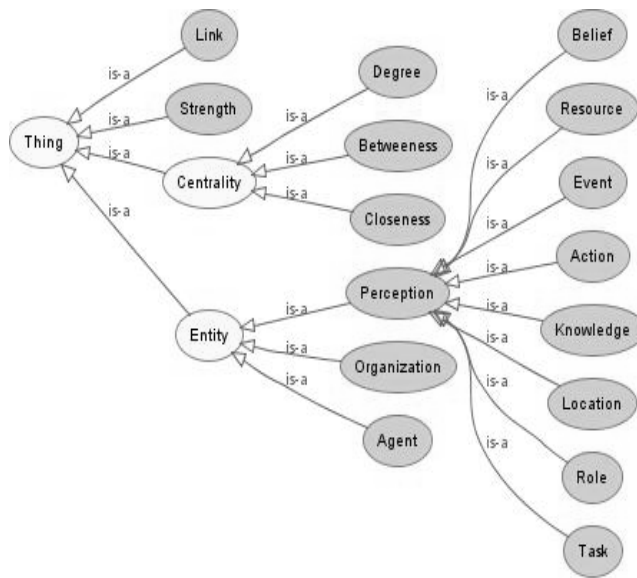


Fig. 12.10 GraphViz Diagram – Social Network inferred refactored ontology

TABLE 12.1 Influence Net Refactored Ontology Elements (Concept Map Imports)

Ontology Domain Class	Ontology Object Property	Ontology Range Class	Concept Map - Propositions		
			Concept 1	Relationship	Concept 2
Proposition	<u>SuperClassOf</u>	Action	Proposition	can define	Action
		Ability			Ability
		Belief			Belief
		Decision			Decision
		Event			Event
		Intent			Intent

TABLE 12.2 Explicit Influence Net Concepts in Refactored Ontology

Category	Ontology Domain Class	Ontology Object Property	Ontology Range Class
Explicit Concepts	subject	hasSubjectValue	subject
	verb	hasVerbValue	verb
	object	hasObjectValue	object
	outcome	hasOutcomeValue	outcome
	Evidence	HasEvidence	Evidence
	quality	hasQualityValue	quality
	state	hasStateValue	state
	PropositionType	hasPropositionType	Affirmative, Negative

TABLE 12.3 Social Network Refactored Ontology Elements (Concept Map Imports)

Category	Ontology Domain Class	Ontology Object Property	Ontology Range Class	Concept Map - Propositions		
Concepts from Concept Maps	Perception	hasPerception Value	Action	Perception	Types	Action
			Belief			Belief
			Event			Event
			Knowledge			Knowledge
			Location			Location
			Resource			Resource
			Role			Role
			Task			Task

Enriched Ontology

The focus of the effort so far has been to construct an enriched ontology filled with the concepts (classes) and relationships (properties) of Influence Net and Social Network modeling techniques within and across them. The motive behind the construction of this enriched ontology is to identify mappings (i.e., subsumption and equivalence) between the concepts of both modeling techniques so that the exchange of information or analysis results between models constructed using both techniques can be formalized. The refactored ontologies serve as the basis for the enriched ontology.

The diagram in Fig. 12.11 illustrates the fusing of concepts from both types of refactored ontologies inside the enriched ontology. This is achieved by defining additional object properties in related classes and asserting them in the new ontology. For instance, the *Agent* and *Organization* classes from the Social Network refactored ontology can be mapped to the *subject* and *object* classes of the Influence Net refactored ontology by adding *hasSubjectValue* and *hasObjectValue* object properties to the existing object property of *Agent* and *Organization* classes. Once, the ontology reasoner is executed, it classifies *Agent*, *Subject*, *Object*, and *Organization* as equivalent classes by inferring that the new object properties added to the *Agent* class map *Subject*, *Object*, and *Organization* to itself as shown in Figure 12.12. The inferred class hierarchy shows these classes with the equivalence sign (Figs. 12.12-12.14). Similarly, *Belief* class in Social Network refactored ontology can be mapped to Influence Net's *Belief* class. The *Event* class of Social Network refactored ontology maps to the *state* class which is the constituent of the *Event* class in Influence net refactored ontology also Social network's Knowledge class maps to quality class which is the constituent of *Ability* class in Influence net refactored ontology. The class *Task* of Social Network can be mapped to the class *verb* of Influence Net refactored ontology. Table 4 summarizes these mapped concepts between the two refactored ontologies. The ultimate result of this mapping is an enriched ontology which is the knowledge container of both Influence Net and Social Network modeling techniques. The class hierarchy view of the obtained enriched ontology is shown in Fig.12.15. Note that the ontology in Fig. 12.15 shows both the asserted and inferred equivalences and relationships among concepts from both Influence Net and Social Network refactored ontologies.

Influence Net Refactored Ontology

Classes

subject, object, verb
Action Proposition, etc.



Social Network Refactored

Ontology Classes

Agent, Organization, Task
Action, etc.



subject (hasSubjectValue some subject or hasAgentValue some Agent or hasOrganizationValue some Organization)
object (hasObjectValue some object or hasAgentValue some Agent or hasOrganizationValue some Organization)
Agent (hasAgentValue some Agent or hasSubjectValue some subject or hasOrganizationValue some Organization)
Organization (hasAgentValue some Agent or hasSubjectValue some subject or hasObjectValue some object)
verb (hasVerbValue some verb or hasTaskValue some Task) etc.
Action (hasElements some subject, verb and object or hasActionValue some Action)

Figure 12.11 Enriched ontology classes



Fig. 12.12 Subject, Object classes mapped to Agent class



Fig. 12.13 Reasoner inferred equivalences



Fig. 12.14 Subject, Object, organization and Agent as equivalent classes

TABLE 12.4 Enriched ontology

Enriched Ontology (Influence Net & Social Network Refactored Ontology Mapped Concepts)					
Influence Net Refactored Ontology Elements			Social Network Refactored Ontology Elements		
Domain Class	Object Property	Range Class	Domain Class	Object Property	Range Class
subject	hasSubjectValue	subject	Agent	hasSubjectValue	Subject
	hasAgentValue	Agent		hasAgentValue	Agent
	hasOrganizationValue	Organization		hasObjectValue	Object
object	hasObjectValue	Object	Organization	hasObjectValue	Object
	hasAgentValue	Agent		hasAgentValue	Agent
	hasOrganizationValue	Organization		hasSubjectValue	Subject
verb	hasVerbValue	verb	Task	hasVerbValue	verb
	hasTaskValue	Task		hasTaskValue	Task
Intent/ Decision	hasElements some Action and hasElements some subject and hasElements some verb	Action, subject, verb		hasElements some Action and hasElements some subject and hasElements some verb	Action, subject, verb
Action	hasElements some subject and hasElements some verb and hasElements some object	Action, subject, verb	Action	hasElements some subject and hasElements some verb and hasElements some object	Action, subject, verb
	hasActionValue	Action		hasActionValue	Action
Belief	hasElements some subject and hasElements some verb and hasElements some (Ability or Decision or Action or Event) and hasEvidence some Evidence	subject, object, verb, Ability or Decision or Action or Event, Evidence	Belief	hasElements some subject and hasElements some verb and hasElements some (Ability or Decision or Action or Event) and hasEvidence some Evidence	subject, object, verb, Ability or Decision or Action or Event, Evidence
	hasBeliefValue some Belief	Belief		hasBeliefValue some Belief	Belief
state	hasStateValue	State	Event	hasStateValue	State
	hasEventValue	Event		hasEventValue	Event
quality	hasQualityValue	quality	Knowledge	hasQualityValue	quality
	hasKnowledgeValue	Knowledge		hasKnowledgeValue	Knowledge

The process described in this section is repeatable for any set of modeling techniques. For instance, to extract semantic knowledge about CPN (Colored Petri Net), similar Concept maps (for the defined focus questions) can be developed, followed by its syntactic model, and then pseudo and refactored ontologies. Ontology construction is an intense brainstorming activity. By the time the refactored ontology is completed, enough insight into the modeling technique should have been achieved that the ontology designer would easily be able to map CPN concepts to the related Social Network or Influence Net concepts (if any). These newly mapped concepts can then be incorporated into this enriched ontology and an updated enriched ontology can be obtained which would serve as the knowledge container for Influence Net, Social Network and CPN modeling techniques altogether.

12.4 Conclusion

In this work, we made a formal attempt to study the syntactic and semantic interoperations among disparate models employed to address a specific problem of interest. In this study, we examined the interoperations among Social Network and Timed Influence Net models that are constructed for an adversarial behavior modeling problem. The approach used for the study is a phased approach, employing a combination of concept maps, syntactic-models, and ontologies, consisting of the following three levels: as part of its Conceptual Modeling Level, it specifies a modeling language by constructing a generalized Concept map that captures the intensional knowledge about the problem that is used in the models done using the language; in the Syntactic Modeling Level, the structural aspects of a modeling language, in the context of the problem domain, are formalized as a UML syntactic model; the final Ontological Modeling Level starts by exporting the syntactic model into an OWL ontology and then refactoring and enriching it semantic concepts and relationships. Once the individual ontologies are completed for each modeling technique, mapping of concepts across the ontologies is started. The resulting merged ontology which contains these concepts and relationships within and across multiple ontologies is called a merged enriched ontology.

The refactored, and enriched ontologies are Description Logic (DL) knowledge-bases with terminological axioms, i.e., collections of concept axioms and role axioms expressed in OWL. In our study of the two modeling languages, i.e., Social Networks and Timed Influence Net, the discovered mappings in the merged enriched ontology suggest possible semantically correct ways to ensure consistency and to exchange information between the two types of models when they are used to solve a problem of interest in a domain. This mapping, in a practical model building exercise, may help a modeler resolve possible ambiguous, uncertain, and missing data issues by identifying what concepts can validly be ‘exported’ from one model to another. For example in our study, the *Agent* and *Organization* classes from the Social Network ontology is shown mapped to the *subject* and *object* classes of the Influence Net refactored ontology after the ontology reasoner is executed, as shown in Figure 12.12. This mapping between the components of the two types of model can help determine the concepts present in one model that may be missing in the other, provided the two types are already developed. In a situation where one of the two models is constructed prior to the other, the merged enriched ontology informs the modeler about the concepts can be validly ‘exported’ from the existing model for the construction of the other.

The presented study of the two modeling languages, Social Networks and Timed Influence Net, shows the promise of the proposed Meta-modeling approach for a Multi-modeling framework in formalizing the possible interoperations between the two types of models. It, however, only identifies interoperations in the form of knowledge items that the two modeling techniques share as part of their design constructs. A semantically more detailed and richer ontology of the two may take this mapping to other levels of interoperations as outlined in Section 12.1.

PART V: COMPUTATIONAL EXPERIMENT

Chapter 13: Cyber Deterrence Policy and Strategy

Chapter 14: Application: The India-Pakistan Crisis Scenario

Chapter 13

Cyber Deterrence Policy and Strategies

Robert J. Elder, Alexander H. Levis

13.1 Introduction

Deterrence as practiced during the Cold War was largely defined in terms of capabilities to impose punishment in response to an attack; however, with growing concern over the proliferation of nuclear and other “mass effect” technologies such as cyber technologies, deterrence has evolved to be understood more generally in terms of cost/benefit calculi, viewed from not only a national perspective, but also recognizing the importance of both friendly and competitor perspectives. With this more holistic approach, the primary instruments used for deterrence are those which encourage restraint on the part of all affected parties. The use of a multiple lever approach to deterrence offers a path to an integrated strategy that not only addresses the cost/benefit calculus of the primary attacker, but also provides opportunities to influence the calculus of mercenary cyber armies for hire, patriotic hackers, or other groups. While traditional denial and punishment approaches to deterrence may not apply, the broader approach to deterrence described in the US Deterrence Operations Joint Operating Concept, which emphasizes the cost/benefit relationship of restraint on the part of potential adversaries, offers a sound foundation for deterring cyber conflict even where multiple actors are involved. This suggests that the strategic concepts of escalation control and extended deterrence which developed from years of experience with nuclear deterrence are applicable and relevant to cyber deterrence.

The letter report “Deterring Cyber Attacks: Informing Strategies and Developing Options for U.S. Policy”² posed a series of questions regarding opportunities to apply the strategy of deterrence developed during the Cold War to future cyberdeterrence strategies: First, is there a model (or modeling approach) that might appropriately describe the strategies of state actors acting in an adversarial manner in cyberspace? And, is there an equilibrium state that does not result in cyber conflict? Second, how will any such deterrence strategy be affected by mercenary cyber armies for hire and/or patriotic hackers? Third, what are the strengths and limitations of applying traditional deterrence theory to cyber conflict? Fourth, what lessons and strategic concepts from nuclear deterrence are applicable and relevant to cyberdeterrence? And finally, how should a U.S. cyberdeterrence strategy relate to broader U.S. national security interests and strategy? This paper addresses these questions and suggests that the use of multi-modeling can serve as a powerful aid to cyberdeterrence analysts and policymakers.

Strategic deterrence during the Cold War was defined in terms of capabilities to impose punishment or deny benefits; however, with growing concern over the proliferation of nuclear and

² Letter Report for the Committee on Deterring Cyberattacks: Informing Strategies and Developing Options for U.S. Policy, Committee on Deterring Cyberattacks; National Research Council, March 25, 2010.

other similar “mass effect” technologies, to include cyber attack, deterrence has evolved to be understood in terms of cost/benefit calculi, viewed from not only a national perspective, but also the perceptions of friends and competitors. The evolving primary deterrence objective is to encourage restraint on the part of all affected parties, and the primary means is to establish mutual understanding among actors designed to prevent one actor from conducting actions or exhibiting behaviors that are unacceptable to one another. In this context, the instruments of deterrence are not only the capabilities to impose punishment or deny the effects of adversary actions, but also the means to identify friendly and competitor vital interests, to communicate with friends and adversaries, to validate mutual understanding of redlines, and to control escalation.

Consider how deterrence was conducted during the Cold War. There were three pillars to the Cold War deterrence concept. The first was the posturing of forces and capabilities to demonstrate readiness. The second was the conduct of visible actions which showcased capabilities. The third was messaging to explain the activities and posture changes to both friends and potential competitors. Messaging was also used to outline United States strategy with regard to nuclear weapon use, particularly with respect to redlines. It is worth noting that ambiguous messaging is a message in itself, even if not intended! The United States postured its forces by putting them on alert both for survivability and to demonstrate their readiness. The US conducted a variety of visible activities to include large-scale, force-wide exercises and small-scale, local exercises to demonstrate its capabilities. And the United States messaged its competitors and friends through use of declaratory strategies (which did not always match the actually implemented strategies), public and private diplomacy, international conferences, and the media. Although it is difficult to assess the effectiveness of individual elements of the Cold War strategy and policies, it is clear that nuclear weapons were not used during the Cold War, and the United States did not fight a major conventional conflict with the Soviet Union that put US vital national interests at risk. This suggests that there might be utility in employing a parallel operational concept applied to cyber deterrence; however, given the complexity of cyberdeterrence, it requires a structured process which can be aided through the use of complex multiple-actor modeling.

US deterrence strategy was effective in preventing a nuclear exchange during the Cold War, but this same strategy had another benefit: fear that one side might employ “tactical” nuclear weapons to avoid catastrophic defeat led both superpowers to avoid a major conventional war in Central Europe.³ Just as one cannot expect cyberdeterrence to prevent cyber network exploitation (CNE) today, Cold War deterrence strategies did not prevent espionage or conflicts involving proxies. This suggests that an effective cyberdeterrence strategy does not stand alone, but serves as an element of a comprehensive defense strategy. Therefore, deterrence in a national security sense is limited to behaviors that threaten the nation’s vital interests such as attacks on indications and warning systems, public safety and health systems, homeland defense capabilities, or national financial systems. Complementary defensive actions are employed to protect against other detrimental, but less onerous, behaviors, much as the US maintained conventional forces during the Cold War to counter attacks against global interests that were not considered vital in nature.

³ Note: Although the effects of all nuclear weapons are strategic in nature, weapons that were employed by tactical forces were called “tactical” nuclear weapons during the Cold War.

Deterrence fails because mutual understanding between actors is lost, or one actor's cost/benefit calculus drives an unacceptable behavior despite the threat of punishment. Therefore, a holistic approach to deterrence requires the US to identify both US and competitor vital interests, to establish a robust, open dialogue with competitors and friends, and to develop and maintain a range of actions to maintain the stability of the relationship. This naturally leads to the development of strategies designed to (1) assure friends and allies, (2) dissuade adversaries from developing capabilities that threaten national well being, (3) deter potential adversaries by encouraging restraint, denying benefits, and threatening to impose unacceptable cost, and (4) maintain capabilities to terminate conflict at the lowest level of destruction consistent with strategic objectives. Regardless of the decision maker, deterrence involves four primary considerations:

1. The perceived cost of restraint (calculus: costs of not taking an action)
2. The perceived benefits of restraint (calculus: benefits of not taking an action)
3. The perceived benefits of taking action (calculus: will action achieve the desired effect?)
4. The perceived costs of taking action (calculus: how will the competitor respond?)

Understanding how these factors are interrelated is critically important to determining how best to influence adversary decision-making.

Multi-modeling, as described in Chapters 11 and 12, is showing great promise as a means to represent the complex interactions among the many actors in cyberspace and thus support analysts and planners as they develop cyber deterrence strategies. The use of a computational infrastructure and systematic workflow provides a rigorous basis to operate various models accessing the same data sets, analyze their use of the data to identify ambiguities and omissions, and conduct coordinated multinational government and non-government (Unified Action⁴) course of action development and evaluation.

To explore the use of deterrence as a strategy for dealing with cyber attack, an understanding of the political and military use of nuclear deterrence during the Cold War will be informative, both from a capabilities and a limitations perspective. And, to get a better perspective of how Cold War deterrence concepts might apply in cyberspace, it is also useful to consider the differences between the Cold War strategy for deterring employment of nuclear weapons and today's strategy for deterring proliferation of weapons of mass destruction.

During the Cold War, the United States faced an overwhelming conventional capability from the Soviet Union and its satellites, and therefore felt the need to have a capability to maintain a non-strategic (tactical) nuclear capability that could be used if NATO's conventional forces were being overrun. An elaborate escalation control regime was established on both sides because there was a fear that the tactical fight in the central plains of Europe could rapidly escalate into a conflict between the two superpowers. Promulgating this fear was intentional on the part of the United States, particularly during the period when the Soviet Union's conventional force capability clearly exceeded NATO's conventional forces in Europe. So, although there were skirmishes between these two great powers and their allies, the potential for use of tactical nuclear weapons

⁴ Unified action synchronizes, coordinates, and/or integrates joint, single-Service, and multinational operations with the operations of other USG agencies, NGOs, and IGOs (e.g., UN), and the private sector to achieve unity of effort. See Joint Publication 1: Doctrine for Armed Forces, pg. II-2.

effectively deterred a conventional conflict. Although there was a clear differentiation between the use of strategic and tactical nuclear weapons, both sides believed that tactical nuclear weapon use would lead to strategic nuclear exchange. Central Europe would bear the cost of tactical nuclear weapon use, but the United States and the Soviet Union would be the targets of a strategic nuclear exchange. As a result, first use of nuclear weapons was unlikely unless one actor became convinced that its vital interests were threatened and believed that only the employment of nuclear weapons would prevent the stronger power from imposing its will on the lesser power. This key element of the Cold War deterrence calculus could be useful in formulating approaches to deter cyber attacks that threaten United States vital interests.

Today's National Security Strategy⁵ places great emphasis on countering the proliferation of technologies which would enable new actors, state or non-state, to develop nuclear weapons. Deterrence strategies have been developed to help counter nuclear proliferation, but these strategies employ significantly different levers to affect the relevant actors' decision calculi. Cold War nuclear deterrence was largely influenced by concerns that use of nuclear weapons by either side would lead to a strategic nuclear exchange against the US and Soviet homelands. Deterring the proliferation of WMD technologies is grounded in the Nuclear Non-proliferation Treaty, with access to peaceful uses of nuclear energy denied to proliferators. But another deterrence lever has been the use of international diplomacy to remind potential proliferators and their supporters that weapons developed from their proliferated technology could also be used against them in the future. It is difficult to prevent a country from proliferating technology for purely economic or political benefit; nevertheless, the non-proliferation regime establishes a framework which enables the effective use of integrated deterrence, defense, and international legal strategies. Similarly, deterring major cyber attacks conducted for economic, military, or political purposes will be difficult, but an international regime could provide a framework to enable effective international cyberdeterrence strategies.

13.2 Deterrence Operations Joint Operating Concept

The Deterrence Operation Joint Operating Concept (DO-JOC)⁶ outlines a basic approach to deterrence and was a first attempt to apply Cold War lessons to post-Cold War challenges. USSTRATCOM has evolved this concept dramatically over the last three years and is in the process of updating the 2006 document. The DO-JOC postulates a series of critical assumptions for effective deterrence of adversarial actions and behaviors that can be applied to cyber deterrence: First, the United States is aware that an adversary (state or non-state) possesses a cyber attack capability that threatens its vital interests. Second, the adversary actions to be deterred result from deliberate and intentional calculations regarding alternative courses of action and their perceptions of the values and probabilities of alternative outcomes associated with those different courses of action. Finally, cyber deterrence must assume that at least some adversary values and perceptions relevant to their decision-making can be identified, assessed, and influenced by others. The DO JOC goes on to note that some actors (both state and non-state) will be extremely difficult to deter; however, truly irrational actors are extremely rare.⁷ Their calculus may be very different from that of the United States but what constitutes rational behavior must

⁵ National Security Strategy of the United States, 2006.

⁶ Deterrence Operation Joint Operating Concept, Version 2.0, Department of Defense, 2006 (Available at www.dtic.mil/futurejointwarfare)

⁷ Ibid, pg.11.

be understood in their terms. The following examination of cyberdeterrence accepts these fundamental assumptions and focuses on deterring rational actors from attacking US vital interests in or through cyberspace.

When most people think of deterrence, the first thought that comes to their minds is the ability to impose significant punishment in retaliation for an attack. However, the Deterrence Operations JOC suggests that adversaries can also be deterred if they feel their actions will not achieve the desired benefits (denial) or that restraint from the action will achieve a better outcome than taking the action the US seeks to deter.⁸ This brings us back to the concept of extended deterrence experienced during the Cold War: The stronger power was deterred from using its overwhelming conventional capability to defeat the weaker power because the weaker power could employ its nuclear weapons to deny the objectives of the stronger power's attacks. In the context of non-proliferation there is another form of "extended deterrence" where the US offers to "extend" its nuclear deterrent force to protect US allies provided they eschew development of their own nuclear weapon capabilities. For most of these countries, the US commitment to protect them from nuclear attack greatly simplifies their cost/benefit calculus regarding the need to develop their own nuclear weapon capability.

One of the lessons we can take from the Cold War experience is the tremendous value of asymmetric weapons to a power that is less capable of conventional warfare, particularly when these powerful weapons can be obtained at relatively little cost. Interestingly, when viewed from the adversary's perspective, the acquisition of these capabilities may be intended to serve as a deterrent to actions and behaviors by the United States. We also see that while the ability to impose great punishment is a key aspect of deterrence, there are also aspects of deterrence which flow from the ability to encourage restraint and deny the benefits of adversary actions. With the end of Cold War, and the emergence of new nuclear powers, the United States recognized that the traditional Cold War approaches to deterrence might not be effective. Now, with the advent of national missile defense, the United States possesses the capability to defend the homeland from a limited missile attack, and since it can use non-strategic weapons to deter less capable potential adversaries of our friends and allies without fear of a retaliatory attack on the US homeland, the credibility of US extended deterrence is enhanced.

13.3 Enabling Deterrence

It will be instructive to assess how the DO JOC applies to cyberdeterrence. With credit to General Larry Welch, cyber deterrence is difficult unless we first understand our critical vulnerabilities and take action to protect them⁹. Another important concept can be found in a 2008 AFSAB report¹⁰ which argued that it is important to protect the United States from the effects of attacks rather than just protect the targets of the attacks. One might think of this protection against effects as "mission assurance" or "cyber resiliency" as contrasted with traditional "information assurance" which focuses on the protection of networks and systems. From a deterrence perspective, the idea is to introduce uncertainty in the adversary's mind that the attacks will achieve the desired effects; if they don't, and there is a possibility that the source of the attack might be determined through forensic analysis or intelligence means, this potential denial of benefit should

⁸ Ibid., pg. 5

⁹ Welch, General Larry, presentation to Cyber Warfare 2008, London, U.K., March 31, 2008.

¹⁰ Defending and Operating in a Contested Cyber Domain, Air Force Scientific Advisory Board, DTIC, 2008.

affect the adversary's decision calculus. The potential for attribution can be improved by "reducing the noise level" through improved security and defense of critical information, systems, networks, and infrastructure, making it easier to detect behaviors that might pose a threat to the United States. This could include establishment of protocols and standards that govern both the public and private sectors in areas that could affect United States vital interests. Daniel Geer recently addressed the need for policy choices that support risk management versus risk avoidance, clearly recognizing that our current risk avoidance approach to cyberspace is attractive, but impractical. He postulates that Americans want freedom, security, and convenience, but they can only have two.¹¹ The Nation must make choices to implement a cyberdeterrence strategy; anything that sacrifices vitally important aspects of national and economic security in cyberspace for purposes of convenience simplifies the attack problem for a cyber adversary.

One of the benefits of multi-modeling is the requirement to capture subject matter expert (SME) knowledge in a systematic manner for expression in an appropriate modeling language, especially in situations that involve multiple actors. The discipline of workflow use has been found useful for information sharing among subject matter experts, and it is critical for the success of complex analyses involving multiple actors, perspectives, behaviors, and capabilities. One of the current experiments to evaluate the basic research conducted thus far involves analysis of alternatives to control escalation that could lead to WMD use involving two countries with which the United States maintains friendly relationships, India and Pakistan (Chapter 14). One of the challenges for the United States is that one country is assigned to the US Pacific Command, and the other is assigned to US Central Command, and so the analysts and commanders operate with different models and perceptions, and although centrally sourced, with different intelligence. The use of multi-modeling allows central decision makers to better understand the interrelationships and differences in perspectives that exist. The models do not make decisions, but provide information that can aid the analyst to assess different courses of action for the consideration of the decision maker. The Pythia timed influence network application uses conditional probabilities established through subject matter expert analysis to allow course of action assessments. Since the models are integrated through the C2WT, potential actions propagate through all actor models to highlight potential second and third order (and beyond) consequences.

13.4 Conclusion

Research into the use of structured analytical approaches to modeling human adversary behaviors suggests that it is possible to gain insights into the strategies of state actors acting in an adversarial manner. Although additional research is required, it appears that the use of Bayesian influence models, informed by social network models, can be used to evaluate escalation control measures that maintain a state of equilibrium that does not lead to major conflict. Furthermore, the use of a multi-lever approach to deterrence offers a path to an integrated strategy that not only addresses the cost/benefit calculus of the primary attacker, but also provides opportunities to influence the calculus of mercenary cyber armies for hire, patriotic hackers, or other such groups. While the denial and punishment approaches, which underpinned Cold War deterrence, may not apply, the broader approach to deterrence described in the US Deterrence Operations Joint Operating Concept, which emphasizes the cost/benefit relationship of restraint on the part of potential

¹¹ Daniel E. Geer, Jr., Sc.D., "Cybersecurity and National Policy," Harvard National Security Journal, vol. 1, April 7, 2010.

adversaries, offers a sound foundation for deterring cyber conflict even where multiple actors are involved. In particular, the strategic concepts of escalation control and extended deterrence which developed from years of experience with nuclear deterrence are applicable and relevant to the broader concept of deterrence.

Chapter 14

The India-Pakistan Crisis Scenario

The GMU/CMU MURI Team

14.1 Introduction

In this chapter, we present the case study that was carried out to provide context for the adversarial and multi-modeling effort.

The scenario is fictitious and has been designed to serve the purposes of the proposed modeling research. It is based loosely on events that occurred between India and Pakistan in June of 2002. The animosity between these two nations has its roots in history and religion and is epitomized by the long-running conflict over the state of Jammu and Kashmir (J&K). China is also administering Aksai Chin at the northeastern corner of Jammu and Kashmir - a situation contested by India. In this scenario, a sequence of terrorist incidents occurs in Sri Nagar, the capital of Indian administered state of J&K, and along the Line of Control (LOC) separating the Pakistani controlled Northern Areas and India controlled Jammu and Kashmir. At the same time, there is growing instability and disaffection with the government in Pakistan, while in India the opposition parties are becoming stronger. The two countries start making a series of escalating moves (e.g. recall diplomatic staff, move troops toward the LOC, reposition mobile missile batteries, and there is activity in both countries' nuclear weapon facilities) while events (e.g. bombings) continue to take place on both sides of the LOC. As is common in such situations, part of the escalation is due to lack of understanding of the adversary's intent (partial knowledge of the state - only some of the kinetic components are known, such as troop movements, but not intent and strategy).

The situation alerts the US Government as well as its armed forces. Specifically, US Central Command (CENTCOM) has Pakistan within its geographic area of responsibility (AOR), while US Pacific Command (PACOM) has India within its geographic AOR. China is also concerned because of the danger of nuclear exchanges between India and Pakistan. The United Nations and Russia may also get involved. The USA's and international community's objective is to dissuade/deter the two adversaries from escalating the situation into a nuclear exchange and affect a rapid de-escalation of the crisis. It has been observed that in such crises, misinformation about the intent of the adversary and misinterpretation of the moves an adversary makes (quite often in response to domestic pressures) tend to escalate the crisis. Consequently, the US develops an Intelligence, Surveillance, and Reconnaissance (ISR) plan to keep its leaders informed. This is complemented with a coordinated information operations campaign in which the US provides improved information about the state of the conflict, thus reducing the ambiguities about the system state, to the two adversaries: CENTCOM to Pakistan and PACOM to India.

The purpose of this scenario is to provide a basis for demonstrating the usefulness and effectiveness in using multiple models together to assess an evolving situation and develop appropriate courses of action (COA). As mentioned, the scenario is based on fictitious events that occurred in the summer of 2002 between India and Pakistan. These events are inspired by similar

real incidents that have taken place in or between the two countries as part of the long-running dispute between them. While the events are fictitious, the actual government structures and the real names of the government officials of the two countries at that time have been used. The actions and/or statements attributed to the government officials, however, are fictitious. The same considerations applied in modeling the US involvement in its attempt first to stabilize the situation and then to de-escalate it. Two US regional commands are involved: CENTCOM and PACOM. As the situation escalates, all instruments of national power become involved. The instruments of US national power are known as DIME: Diplomatic; Information; Military; Economic.

In this scenario, we consider a period from June 1, 2002 to July 20, 2002 and we consider within it two short duration vignettes for the CENTCOM and PACOM analysis and assessment of the situation (see Fig. 14.1).

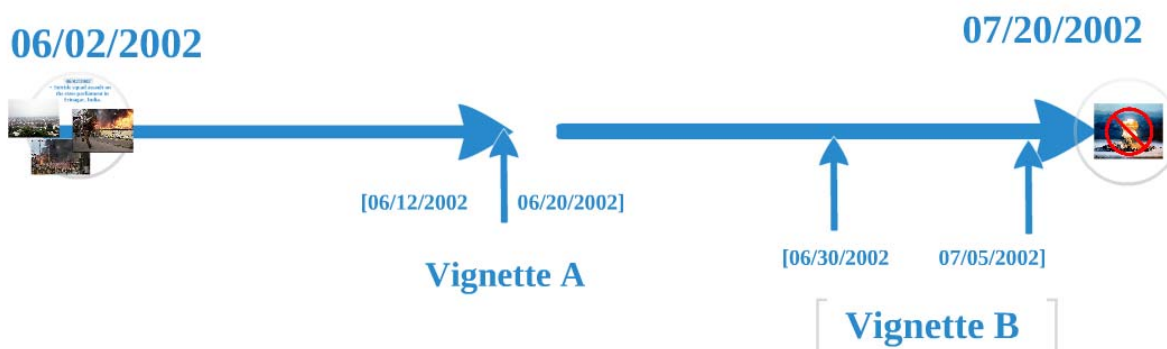


Fig. 14.1 Scenario timeline

The first vignette (A) addresses CENTCOM's and PACOM's observation of the situation and the diplomatic and military maneuvers by India and Pakistan. Each command has a collection of contingency plans (CONPLAN) for their geographic AOR. This situation would cause both commands to more closely observe as well as update their respective CONPLANS and supporting models. Each command would also develop multiple courses of action (COAs) that are focused on broadly defined diplomatic actions and on surveillance and other intelligence activities.

The second vignette (B) assumes that the COAs of vignette A have not produced the desired results and the crisis is escalating. CENTCOM and PACOM assess the situation and develop courses of action to de-escalate the crisis.

Table 14.1, provides details of the scenario events and the two vignettes undertaken by the two command centers.

14.2 Vignette A: Strategic Deterrence/Regional Stability Ops

Vignette A addresses CENTCOM and PACOM observation of the situation and the diplomatic and military maneuvers by India and Pakistan. In this section we present a workflow, comprised of models and their mutual interoperations, used by the two command centers in assessing the situation and in developing contingency plans, if the situation were to escalate. The two commands develop their individual courses of action that are focused on broadly defined diplomatic actions and on surveillance and other intelligence activities. The following is a description of the workflow used by the two commands and the models used by the two.

TABLE 14.1 Scenario and Vignette timeline

<ul style="list-style-type: none"> • 06/02/2002 – 06/07/2002: Contention rises on the Indian part of Kashmir. <ul style="list-style-type: none"> ○ Suicide squad assault on the state parliament in Srinagar, India. ○ Indian government (<i>National Security Advisor</i>) blames Pakistan supported militant organizations (<i>Lashkar-e-Taiba</i>) for the attack. ○ Indian government hands (<i>Director General, Defense Intel Agency</i>) out a list of twenty suspects to Pakistan demanding that Pakistan hand them over to India. ○ Pakistani government (<i>Minister of Foreign Affairs</i>) claims no responsibility and refuses knowledge of the suspects. ○ Pakistani government (<i>Minister of Foreign Affairs</i>) asks India to provide evidence of involvement of Pakistani nationals in the incident citing an agreement reached by the South Asian nations for the extradition of criminals. ○ Indian border security forces (BSF) exchange fire with gunmen trying to cross into Kashmir at LOC. ○ Exchange of fire between Pakistani and Indian forces at LOC. ○ Gunmen killed 11 people in Hindu dominated village of Pogal. ○ Villages along LOC are being evacuated. ○ Pakistan accuses Indian security forces of atrocities against Kashmiri civilian population • 06/08/2002: <ul style="list-style-type: none"> ○ India (<i>Minister of External Affairs</i>) threatens to break diplomatic ties with Pakistan. ○ Pakistani religious (<i>religious party- Jamaat-e-Islami</i>) and nationalist parties (<i>Nationalist Party – Pakistan Peoples Party</i>) stage a demonstration outside Indian embassy in Islamabad. ○ India issues a warning to all its citizens travelling in Pakistan to plan to return. • 06/11/2002: <ul style="list-style-type: none"> ○ Pakistan recalls non-essential diplomatic staff from India. ○ Hindu nationalist parties (<i>Nationalist Party - Shiv Sena</i>) ask Indian government to take punitive actions against terrorist training camps operating inside Pakistan. • 06/12/2002: India recalls non-essential diplomatic staff from Pakistan. • 06/02/2002 -06/12/2002: CENTCOM and PACOM observe events and collect intelligence. • • Vignette A: Strategic Deterrence/Regional Stability Ops: (Monitor the Situation; diplomatic actions) • • 06/11/2002: <ul style="list-style-type: none"> ○ PACOM brings up models (WebTAS, CAESAR III, Pythia, ORA). ○ WebTAS to display events ○ CAESAR III model of Indian Security Council to identify key roles and relationships. ○ ORA model of Indian Security Council and diplomatic contacts; social network analysis for the relevant elements of the Indian Government ○ Pythia to model possible COAs for surveillance and diplomatic actions for de-escalating • 06/11/2002: <ul style="list-style-type: none"> ○ CENTCOM brings up models (WebTAS, CAESAR III, Pythia, ORA). ○ WebTAS to display events ○ CAESAR III model of Pakistani Security Council to identify key roles and relationships. ○ ORA model of Pakistani Security Council and diplomatic contacts; social network analysis of the relevant elements of the Pakistani Government ○ Pythia to model possible COAs for surveillance and diplomatic actions for de-escalating • 06/12/2002: <ul style="list-style-type: none"> ○ PACOM and CENTCOM Coordinate surveillance and military/diplomatic contact COAs ○ Common tasking of Space assets ○ Coordinated tasking of RC-135s ○ Coordinated message to diplomatic and military contacts • • End of Vignette A

- **06/13/2002:** Use space assets begin tracking the designated areas
- **06/15/2002:** PACOM and CENTCOM RC-135s and other assets are collecting data
- **06/20/2002:**
 - Indian heavy troop movement along the border (observed by Pakistan)
 - India increases patrols in the Batalik sector, deploys heavy artillery regiment supported by two divisions of the Indian army and paramilitary forces in the Kargil-Drass sector.
 - India deploys Mirage-2000 Hs, radars and MiG-29 fighter jets along the western border.
 - More specific actions observed by satellites and RC-135s (visuals here; RC-135s and Google Earth)
- **06/23/2002:**
 - Indian mobile missile batteries movement (US data)
 - Unusual movement observed at Kalpakkam (South), Jaipur (West) and Patna (East).
- **06/28/2002:**
 - Pakistani heavy troop movement along the border (observed by India)
 - Movement of multiple mechanized artillery divisions observed at Sialkot and Rann of Kutch sector.
- **06/29/2002:** Diplomatic movements between India and Russia; US aware
 -
 - **Vignette B: Major escalation – COA development for regional deterrence/crisis de-escalation ops.**
 -
- **06/30/2002:**
 - Pakistan equipment movement in and around nuclear facilities (observed by US assets)
 - Movement observed at Kahuta, Khushab, Chaghai Hills and Wah.
- **06/30/2002:**
 - CENTCOM updates models (CAESAR III, Pythia, ORA) and performs new course of action analysis.
- **06/30/2002:**
 - PACOM updates models (CAESAR III, Pythia, ORA) and performs new course of action analysis.
- **07/03/2002:**
 - India equipment movement in and around nuclear facilities (observed by US assets).
 - Movement observed at Manuguru (South) and Thal Vaishet (West)..
- **07/03/2002:**
 - Indian rhetoric on situation heats up
 - Demonstration by Hindu nationalist parties (*Shiv Sena, Bhartiya Janta Party*) in Delhi and other major cities burning Pakistani flags.
 - Media coverage (*News Channels of India- Doordarshan, Star News*) of Indian Pakistani conflict intensifies.
- **07/03/2002:**
 - Pakistani rhetoric on situation heats up
 - Pakistani religious leadership (*leader of Jamaat-e-Islami party*) calls for official declaration of Jihad against India.
 - Media coverage (*News channel of Pakistan – GEO*) of Indian Pakistani conflict intensifies.
- **07/04/2002:** PACOM updates COA
- **07/04/2002:** CENTCOM updates COA
- **07/05/2002:**
 - PACOM and CENTCOM update COA analysis and develop fused COA and submit to JCS and SecDef. (PMEESI)

• **End of Vignette B**

14.3 Vignette A Workflow

Figure 14.2 shows the workflow used in the development of courses of action by the two commands. The activities in the workflow can be divided into the following three broad steps: (1) Identification of data sources; (2) Development of models using the extracted information from the data sources, and (3) Analysis using the models to identify COAs. The development of models, in step 2, requires the use of a number of software tools. The software tools used in the vignette are listed on the left-hand side of Fig. 14.2. Each tool is aligned horizontally with the corresponding modeling and analysis activities in the workflow. The workflow in the figure also lists activities that require interoperation and/or comparison of two or more types of models in order to carry out the refinement and/or analyses on the models. The following is a detailed description of the data sources used, models constructed, and the analyses done on them to assess the situation between India and Pakistan in the context of the scenario of Table 14.1 and to develop some courses of actions to diffuse the tensions between the two nations.

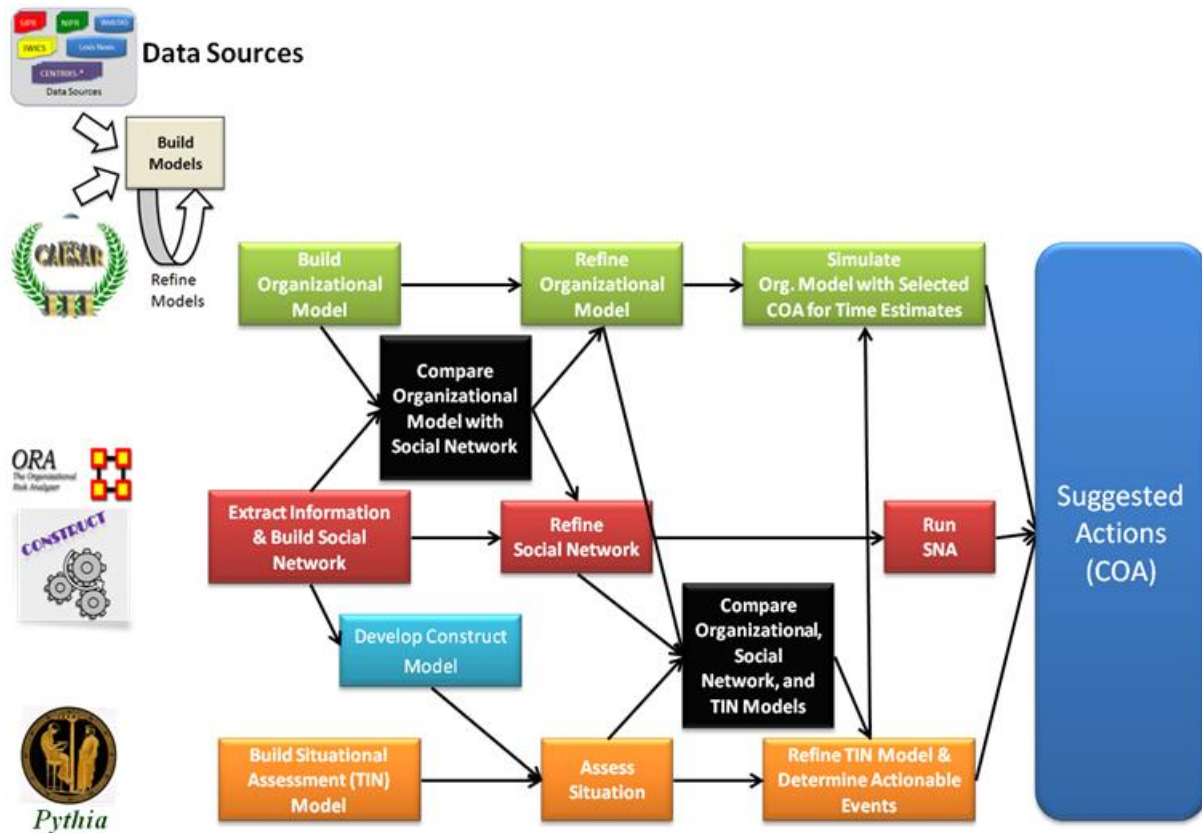


Fig. 14.2: Vignette A workflow

14.3.1 Data Sources

There were, in general, three sets of data the researchers used to support this effort: scenario based inputs (listed in Table 14.1), LexisNexis®-provided data, and web-scraped data from various governments' websites.

The Carnegie Mellon University (CMU) researchers used LexisNexis®-provided data (approximately 3,000 text files representing stored newspaper articles meeting the search criteria) as input into their AutoMap™ tool. The selection criteria within LexisNexis® were the inclusive dates of the scenario's vignettes (20 Jun – 5 Jul, 5 Jul - 22 Jul) and the words “India” and “Pakistan.” AutoMap™ is a semantic extraction tool as well as a social network construction tool. AutoMap™ takes as input generally unstructured sources (e.g. newspaper articles, web scrapings), and through a series of repeatable transformations, generates concept maps/networks, named entities reports, as well as DyNetML™ files for use by Organizational Risk Assessment (ORA)™. The concept maps represent the semantic distance and links between words in the input corpus [105] and helps researches identify which node set(s) individual concepts may belong to. The DyNetML™ files are XML files that represent the nodes (i.e. agents, organizations, locations, knowledge, beliefs, resources, roles, and tasks) and links in the dynamic social networks [106], [107].

The researchers decided a good source for organizational structure data for each nation's national security apparatus (their equivalents to the US National Security Council, Department of Defense, and Department of State) would be the official government web sites of Pakistan and India, including their separately hosted Ministry of Defense web sites. Additionally, to gain a better understanding of how those countries interact with the US, researchers web scraped the official web sites of the US Central Command (USCENTCOM) and US Pacific Command (USPACOM). The US uses these two geographic combatant commands to execute the military and some portions of diplomatic instruments of national power. After these web scraps, CMU had approximately 27,000 text files to input into AutoMap™.

14.3.2 Social Networks

The construction of the social network was an iterative process not completely depicted in Fig. 14.2. There was very tight, though manual, coupling between CAESAR III AutoMap™, and ORA™ to ensure adequate overlap in the real-world and scenario-based population. The manual coupling also ensured correct linkages of the command and control structures and agents in Pakistan, India, and the United States. There was also tight, though still manual, coupling between Pythia™ and Construct™, with parallel coupling between Pythia™, AutoMap™ and ORA™ to ensure adequate overlap in real-world and scenario-based locations and events.

The CMU and George Mason University (GMU) researchers used the AutoMap™ generated networks as direct inputs into ORA. Through a series of human-to-human inter-changes, we identified that the LexisNexis® data had insufficient overlap with the Subject Matter Expert (SME)-generated CAESAR III models built by GMU—it was missing Pakistan's Inter-Service Intelligence (ISI) agency, the Director General of the ISI, India's Defense Secretary, US Combatant Commands, and the US Department of State. Correcting the underlap was primarily through modification of AutoMap's™ processing rules and various thesauri—ensuring key agents did not get deleted, that sensitivity thresholds were raised or lowered, adding entries to thesauri to consolidate like-concepts.

Two separate social networks for each country were built from the CENTCOM and PACOM perspectives. The two networks were a quasi-hierarchy with the national security council (NSC) agents in one network, the NSC-agents plus diplomats relevant to that COCOM's perspective. A third network was built that merged the network models from each COCOM, keeping the agents from the NSC and diplomat networks and all remaining agents relevant to the scenario at hand

(we deleted a great number of cricket and soccer stars, Bollywood stars, and other notables found in the web scrapings that were functionally irrelevant to the scenario).

Tables 14.2, 14.3, and 14.4 summarize the node types, counts, and links between nodes from each US geographic combatant command's perspective.

TABLE 14.2 Vignette A, National Security Council only, CENTCOM & PACOM

Node Sets	CENTCOM (Pakistan)	PACOM (India)
Agents	27	42
Belief ¹²	21	21
Event	40	40
Knowledge	145	145
Location	3250	3250
Organization	321	321
Resource	116	116
Role	247	247
Task	418	418

TABLE 14.3 Vignette A, NSC and diplomats only, CENTCOM & PACOM

Node Sets	CENTCOM (Pakistan)	PACOM (India)
Agents	47	93
Belief through Task ¹⁹	See Table 14.1	See Table 14.1

TABLE 14.4 Vignette A, all agents, CENTCOM & PACOM

Node Sets	CENTCOM (Pakistan) & PACOM (India)
Agents	93
Belief through Task ¹⁹	See Table 14.1 Error! Reference source not found.

¹² Researchers spent the preponderance of their time scrubbing the agent node set. As a function of how the data was collected and scrubbed, the node counts in each of the other node sets are usually identical. This would be unlikely in a set of models maintained by groups of staff members separated by thousands of miles.

CMU then used ORA™ to visualize and analyze the dynamic social networks built through AutoMap's™ text and semantic analysis. More precisely, the focus was on a small set of measures to help inform the development of the other models as well as help refine the suggested courses of action. Initial analysis was through the use of Sphere of Influence reports (see Figs. 14.3 – 14.7), run manually on each actor. With this information, as well as betweenness-centrality, and other 'Key Entities' reports, researchers further 'cleaned' the model of links and nodes not immediately relevant to the purposes of the research (e.g. numerous concepts and agents related to the sport cricket, persons of interest in the US Global War on Terrorism not related to the scenario, historical personages). Researchers maintained records of this additional cleaning to allow better automation through AutoMap™ in follow-on and future efforts. These reporting metrics revealed a consistent and previously unidentified agent with prominent roles in Indian national decision making—the Indian Deputy Prime Minister. The researchers then updated the CAESAR III model to reflect this previously missing agent.






Fig. 14.3 Sphere of Influence Graphic for Indian Foreign Minister during Vignette A's time period. Note the presence of Deputy Prime Minister Advani, who was not in the first iteration of Pythia and CAESAR III models.




powered by ORA, CASOS Center @ CMU

Fig. 14.4 Sphere of Influence Graphic for Pakistani National Security Advisor, for all time periods. There was complete overlap between CAESAR III and Pythia models with this model built through AutoMap and ORA.



powered by ORA, CASOS Center @ CMU

Fig. 14.5 Sphere of Influence Graphic for Indian Prime Minister during Vignette A



powered by ORA, CASOS Center @ CMU

Fig. 14.6 Sphere of Influence Graphic for Indian Prime Minister during Vignette B




Fig. 14.7 Sphere of Influence Graphic for Indian Prime Minister during Vignette C

14.3.3 Construct Model

Construct, based on Constructuralism theory [108], [109], is a turn and agent-based simulation that allows agents to possess and transmit knowledge, forget knowledge, perform tasks (and learn from those tasks), and acquire (potentially incorrect) meta-knowledge about the capabilities and information of other agents in the simulation. A turn-based simulation is one in which every actor is given the opportunity to take a fixed number of actions, when every agent has acted, a new ‘turn’ begins. Actor order is usually fixed. Board-games such as Monopoly, Risk, and Chess can be thought of as turn-based ‘simulations’. Agents in Construct tend to interact on the basis of homophily (a preference for interacting with other agents that share similar perceived traits), but may also seek out expertise (a preference to interact with agents who are perceived to

have a unique or rare knowledge). Construct agents are not required to navigate a virtual-world, although an abstraction of physical proximity can be used to influence the probability of agents interacting with others.

For Vignette A, CMU modeled Indian and Pakistani decision makers as Construct agents. Each of the provocations (as shown in **Error! Reference source not found.**) represented a source of knowledge that contributed to a pro-war belief. Agents started with random distributions of both pro-war and pro-peace knowledge, and were allowed to interact (their interactions were constrained by the found social network built and later refined as part of the workflow). Additional knowledge bits that did not influence pro-war belief were used to provide points of contrast for similarity measures. At various points through the time-course simulation, the number of agents who had the pro-war belief was measured.

In Vignette A, the objective was to assess the impact of all of the provocations that had occurred on networks of Indian and Pakistani decision-makers, assuming both that no further provocations were made, nor any effort to defuse the crisis. As such, a very simple experimental design, as shown in Table 14.5 was used.

TABLE 14.5 Construct experimental design, Vignette A

Parameter	Number of Values	Values
Agent Networks	2	PACOM, CENTCOM
Responses	1	None
Total Knowledge Facts	1	450
Total "Similarity" Facts	1	300
Total "Pro War" Facts	1	50
Total "Pro Peace" Facts	1	100

Using this experimental design, the following two experimental conditions were used for one hundred trials. Figure 14.8 shows the results. The conclusion from this chart is that though tensions that occur throughout the first ten days raise tensions between the two countries, they do not present a clear demand for immediate action to defuse the situation.




Fig. 14.8 The base case presented from Vignette A. We model only those tensions that occur in the first ten days.

14.3.4 Organizational Models

CAESAR III is a software application used for designing information processing and decision making organizations. CAESAR III is developed at GMU. Its design view lets a user design an organization by creating decision makers and defining interactions among them and with the external environment in terms of soft and hard constraints. It uses organization theory, implemented in the Lattice algorithm [52], [110], to generate a solution space that contains all feasible organization structures satisfying the constraints. The organization solution space is bounded by maximally connected and minimally connected organizations. Each organization structure in the solution space is represented in the form of a Colored Petri net. CAESAR III also takes into account the cultural constraints [73], if provided by the designer, within an organization.

Two separate organization models for each country were built from the CENTCOM and PACOM perspectives. The two organizations represent the National Security Council (NSC) structure of the two countries' governments. In a further refinement, diplomats relevant to the two commands' perspective were also added to two NSC structures of the two countries.

The construction of the organization modes was also an iterative process not completely depicted in Fig. 14.2. The initial models of the two organizations were SMPE-generated CAESAR III modes built by GMU. The ORA™ generated social network and analysis results were compared with the CAESAR models to ensure adequate overlap in the real-world and scenario-based population. The manual coupling also ensured correct linkages of the command and control structures and agents in Pakistan, India, and the United States. One important refinement done to the PACOM's model of Indian NSC was due to the discovery of Indian Deputy PM in the ORA sphere of influence reports. The agent was overlooked in the CAESAR model for India.

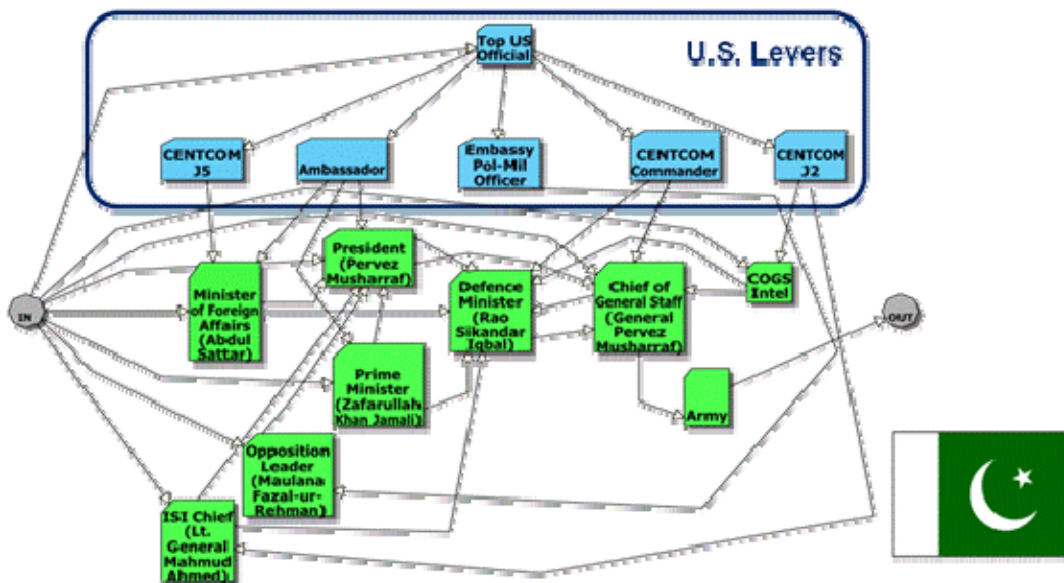


Fig. 14.9 Pakistani Government organization model

Figure 14.9 shows the design of Pakistan's NSC structure in CAESAR III developed from the CENTCOM's perspective. Some of the sources that were manually consulted for the development of this structure are listed below:

- <http://countrystudies.us/pakistan/65.htm>
- http://www.pak.gov.pk/structure_government.aspx

The structure in Fig. 14.9 depicts the Pakistan's NSC as a graph where agents (both government officials, organizations, foreign diplomats), considered relevant to the decision making for security related issues, are modeled as nodes. In this figure the key Pakistani government official and their interactions are shown in the lower half and the primary US officials that can influence those Pakistani government officials are shown at the top. The Ambassador is not a CENTCOM individual but he works closely with CENTCOM. These individuals have been called U.S. Levers to indicate that they could potentially influence key individuals within the Pakistani government. The interactions among them are shown as links between them. The CAESAR III design approach also allows a designer to identify the type of interactions that agents use to communicate with each other in an organization. A class of interactions, e.g., command, information-sharing, results-sharing, etc., is defined to express the types of interactions that can exist among organization members. A narrative description of the organization structure in Fig. 14.9 is given in the following paragraph:

The President is the Head of the government. He gets the information regarding foreign issues related to the country from Minister of Foreign Affairs and situation assessment from Prime Minister. ISI chief passes his intelligence information to the Defense Minister of Pakistan. The President may also get the intelligence information from ISI chief. Opposition leader advises the president upon internal and external security matters of the country. The President discusses his overall assessment of information to the Defense Minister and Chief of General Staff. The Chief of General Staff has an Intel cell that independently assesses the situation inside and outside of the country and provides him with the relevant information.

Figure 14.10 shows the design of India's NSC structure in CAESAR III developed from the PACOM's perspective. Again, the top row of nodes in the model represents the US levers that have contacts (i.e., interaction) with one or more actors (DMs) in the India's NSC. Some of the sources that were manually consulted for the development of this structure are listed below:

- <http://countrystudies.us/india/109.htm>
- http://www.nti.org/e_research/profiles_pdfs/India

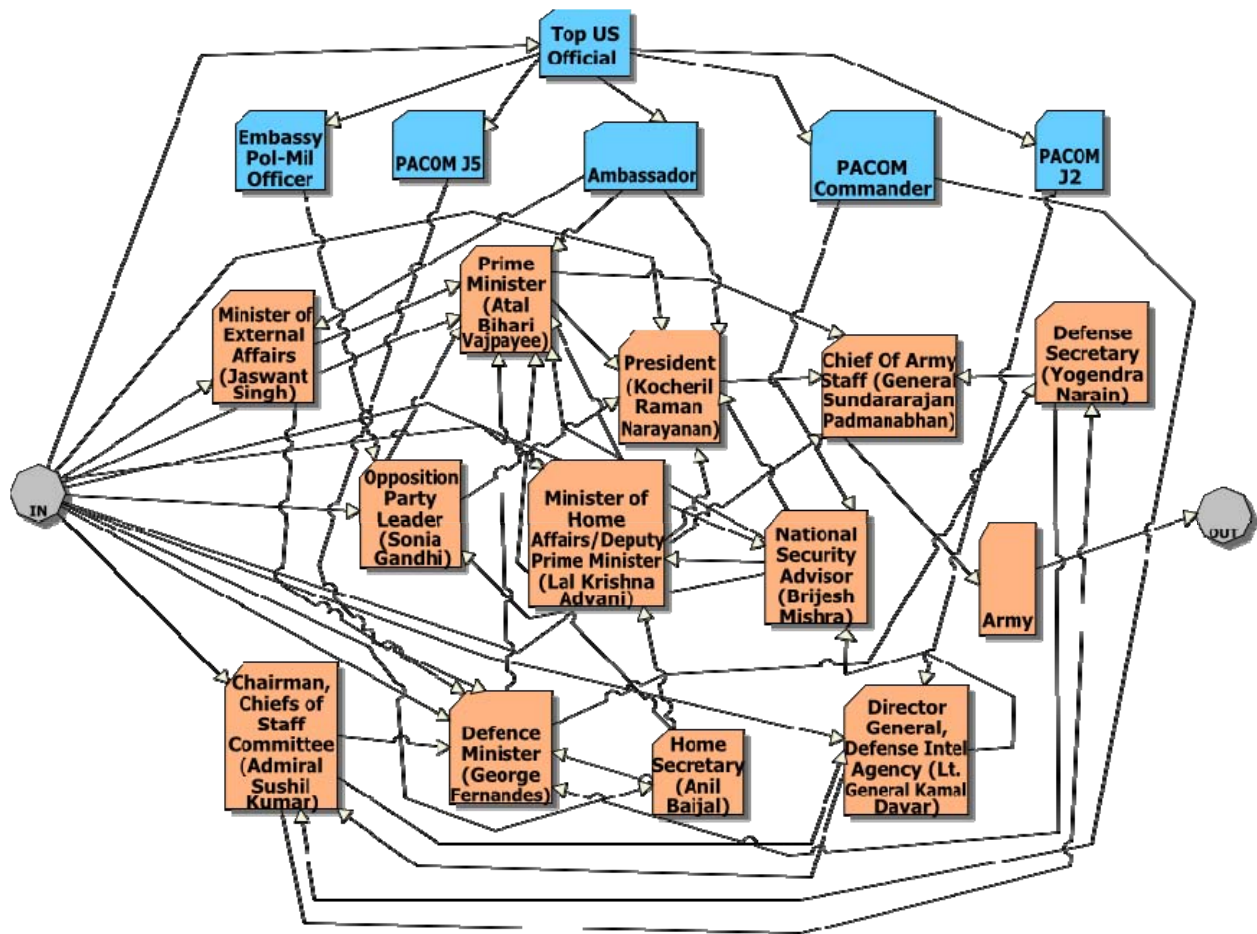


Fig. 14.10 Indian Government organization model

A narrative description of the organization structure in Fig. 14.10 is given in the following paragraph:

The Prime Minister is the Head of the government. President is the Head of the State and acts on the advises of the Prime Minister. Chairman Chief of Staff Committee, Home Secretary, Director General of Defense Intel Agency and National Security Advisor are members of the National Security Council (NSC). National Security Advisor is tasked with regularly advising the Prime Minister on all matters relating to internal and external threats to the country. The Opposition Leader is supposed to give suggestions or advises to Prime Minister on his/her acts. Defence Secretary and Defence Minister are part of Indian Defence Ministry which is charged with the responsibility of internal and external security of the country. Minister of External Affairs is most concerned with foreign affairs of the country. The minister's duties include providing timely information and analysis to the prime minister, recommending specific measures when necessary and planning policy for the future. Deputy Prime Minister seldom carries any powers. Generally, a person with this role also holds key positions like Minister of Home Affairs which is responsible for matters relating to internal security and maintenance of law and order within the country.

14.3.5 Analysis on Organization Models

A sphere of influence analysis was performed on the organization models of the two countries' NSC structures to identify the actors (DMs) in the two governments that can be directly or indirectly (determined by the organization structure and the type of interactions involved) influenced by a particular US lever. For example, Fig. 14.611 represents the sphere of influence of CENTCOM J5 over Pakistani government officials/units. CENTCOM J5 can directly interact with "Minister of Foreign Affairs" of Pakistan. Furthermore, his influence can propagate to the President, defense Minister, chief of general Staff and ultimately to the Army because of the interaction links among these entities. The levers and their spheres of influences for the CENTCOM and PACOM models are summarized in Table 14.6 and Table 14.7, respectively.




Fig. 14.11 Sphere of influence of CENTCOM-J5

TABLE 14.6 CENTCOM sphere of influence report

Lever	Influenced DMs
CENTCOM J5	Minister of Foreign Affairs → President → {Defense Minister, Chief of General Staff} → Army
Ambassador	{Minister of Foreign Affairs, President, Prime Minister} → {Defense Minister, Chief of General Staff} → Army
Embassy Pol-Mil Officer	Opposition Leader → President → {Defense Minister, Chief of General Staff} → Army
CENTCOM Commander	{Defense Minister, Chief of General Staff} → Army
CENTCOM J2	ISI Chief → COGS Intel → {Defense Minister, Chief of General Staff} → Army

TABLE 14.7 PACOM sphere of influence report

Lever	Influenced DMs
Embassy Pol-Mil Officer	Opposition Party Leader → Prime Minister → {President, Chief of Army Staff} → Army
PACOM Commander	{Chairman Chiefs of Staff Committee, National Security Advisor} → {Defense Minister, Director General DIA} → Defense Secretary → Minister of Home Affairs/Deputy Prime Minister → Opposition Party Leader → Prime Minister → {President, Chief of Army Staff} → Army
Ambassador	{Minister of External Affairs, Prime Minister, President} → {Chief of Army Staff, Defense Minister, Defense Secretary} → Army
PACOM J5	Defense Minister → {Prime Minister, President, Chief of Army Staff, Defense Secretary} → Army
PACOM J2	Director General DIA → National Security Advisor → {Opposition Party Leader, Minister of Home Affairs/Deputy Prime Minister, Prime Minister} → {President, Chief of Army Staff} → Army

14.3.6 Situational Assessment

GMU used the influence net modeling tool, Pythia, in concert with two other tools (CAESAR III and Construct) to support an assessment of the situation and course of action analysis in support of Vignette A. It was assumed that CENTCOM and PACOM each had command centers with situation analysis cells. This section discusses Pythia and how it was used during the vignette.

Pythia is a Timed Influence Net (TIN) modeling and analysis tool [13], [25], [111], [112]. With it one can create models of situations that relate Actions or Events to Effects using a sequence of causal relationships. Influence nets (IN) are static (no time) and Timed Influence nets incorporate timing information that enables a modeler to examine the effects over time of a timed sequence of actions.

The models consist of boxes and directed relationships (arrows). The boxes represent random variables that are associated with statements that can be true or false. Typically they represent actions, events, and effects such as beliefs, decisions, and other actions. An arrow that goes from a parent “box” to a child “box” represents a causal relationship. In Influence Nets the relationship can be either promoting (if the parent is true the probability child will increase) or inhibiting (a true parent will reduce the probability of the child).

In creating an influence net a modeler creates the structure of the net and then incorporates the “strengths” of the various influences. The modeler can also specify time delays associated with the arrows (arcs) and the boxes that have parents. Once the model is created, various analyses can be performed with it. If timing information is not included then the Influence net can be used to show the probability that the effects modeled will occur given a set of actions or events. If timing information is included in the model, then by providing the probability of the nodes with no parents as a function of time, probability profiles (probability as a function of time) will be generated for each node that has parents.

To construct and use a TIN to support course of action (COA) analysis, the following process has been defined.

1. Determine the set of desired and undesired effects expressing each as declarative statement that can be either true or false. For selected effects, one may define one or more observable indicators that the effect has or has not occurred.
2. Build an IN that links, through cause and effect relationships, potential actions to the desired and undesired effects. Note that this may require defining additional intermediate effects and their indicators.
3. Use the IN to compare different sets of actions in terms of the probability of achieving the desired effects and not causing the undesired effects.
4. Transform the IN to a TIN by incorporating temporal information about the time the potential actions will occur and the delays associated with each of the arcs and nodes.
5. Use the TIN to experiment with different timings for the actions to identify the “best” COA based on the probability profiles that each candidate generates. Determine the time windows when observation assets may be able to observe key indicators so that assessment of progress can be made during COA execution.

As mentioned earlier, the demonstration was based on two vignettes that were part of a large scenario in which tensions were escalating between India and Pakistan. The scenario starts with a suicide squad assault on the state parliament in Srinagar, India. India blames Pakistan and Pakistan claims that they were not involved in the attack. A series of demonstrations, riots, demands, and diplomatic posturing occurs in both countries along with a increase in military activity. In Vignette A PACOM and CENTCOM each use a set of models created on June 12, 2002 including Pythia to assess the situation.

Figure 14.12 shows the Pythia model that was created for CENTCOM.

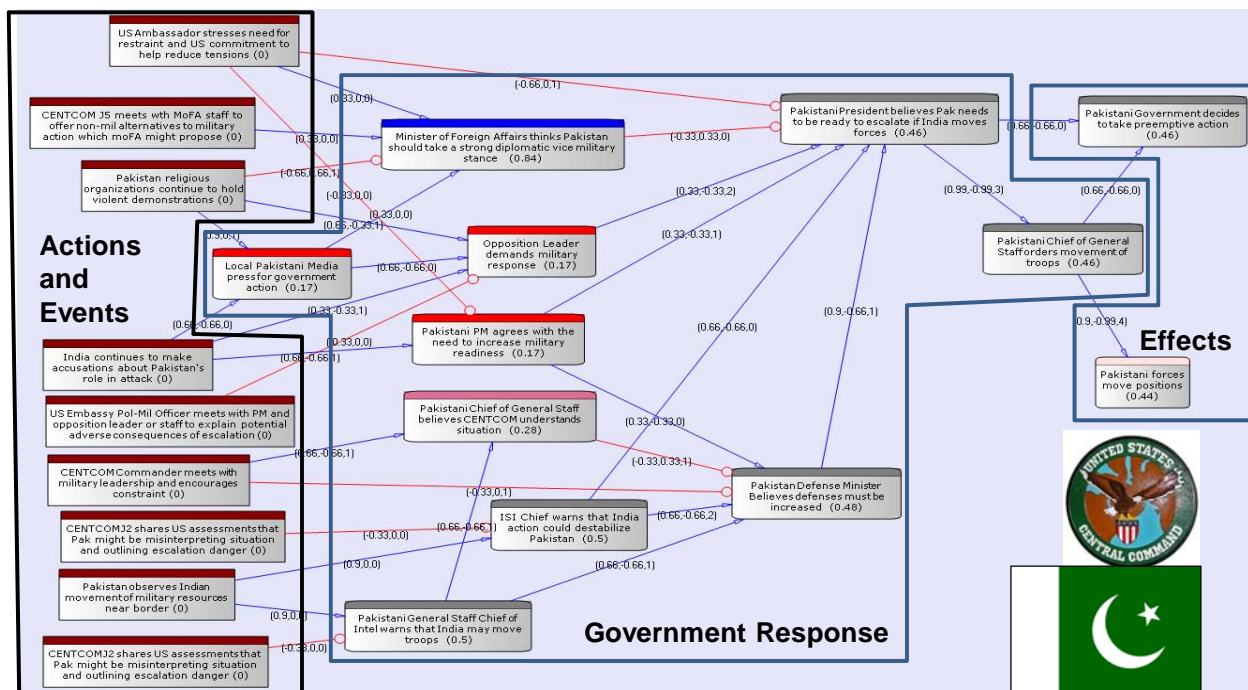


Fig. 14.12 CENTOM Pythia model for Vignette A

The structure of the model can be described as follows. Two main overall affects of concern are shown as boxes on the right side of the model. These effects are “Pakistani Government decides to take pre-emptive action” and “Pakistani forces move positions.” The boxes on the left hand side show the set of actions that CENTCOM may take and a set of possible Events that could happen in the near term. The possible events include (a) “Pakistan religious organizations continue to hold violent demos”, (b) “India continues to make accusations about Pakistan's role in attack” and (c) “Pakistan observes Indian movement of military resources near the border”. The Action nodes include a set of lower level military-to-military (US CENTCOM to Pakistani counter parts), and lower level diplomatic discussions. The boxes in the middle represent the potential reactions of various Pakistani Government Officials to the events and actions and to the reactions of each other. The structure of this model was created using the CAESAR organization model as shown in Fig. 14.10.

A key notion in creating the Pythia model is deciding what statements should be included in the boxes. Subject matter expertise is needed to determine this. The statements include a summary of what each of the levers will say or do and what the expected reaction of the Pakistani government official might be given the parents of the node that represents their reactions. Examples of the types of action are “US Ambassador stresses need for restraint and US commitment to help reduce tensions” and “CENTCOM J5 meets with MoFA staff to offer non-military alternatives to military action which Minister of Foreign Affairs might propose”. An example of the statements indicating possible reactions of the Pakistani officials are “Minister of Foreign Affairs thinks Pakistan should take a strong diplomatic vice military stance”, “Pakistan Defense Minister believes defenses must be increased”, and “Pakistani President believes Pakistan needs to be ready to escalate if India moves forces”.

14.3.7 Assess Situation and COA Analysis

There are many combinations of actions and events that could be analyzed. In the scenario Vignette two cases were examined. The first called worse case was based on all provocative events occurring (e.g., violent demonstrations, Indian government accusations, movement of Indian forces, etc.). In a second analysis Pakistan does not observe the movement of Indian forces. The first analysis done for the worse case was a COA where CENTCOM and the US take no action. **Error! Reference source not found.**Figure 14.13 shows a summary of the results of this analysis. In the probability profile the zero point on the horizontal axis represents the day the analysis was done in the scenario (June 12, 2002).

This analysis was followed by a series of analyses to evaluate the effects of COAs that include the use of the levers and their timing. This was done in two steps. First, the evolutionary search algorithm included the Pythia suite of analysis tools was used to provide and initial set of actions and timing (a COA). The results are shown in Fig. 14.14.

Next, the timing of the actions provided by the search algorithm was adjusted to cause the actions to be taken at the earliest possible time. One constraint was that the CENTCOM Commander should meet after the J2 and J5 have their meetings. This resulted in a slight improvement in the probability profiles of the key effects as shown in Fig. 14.15.

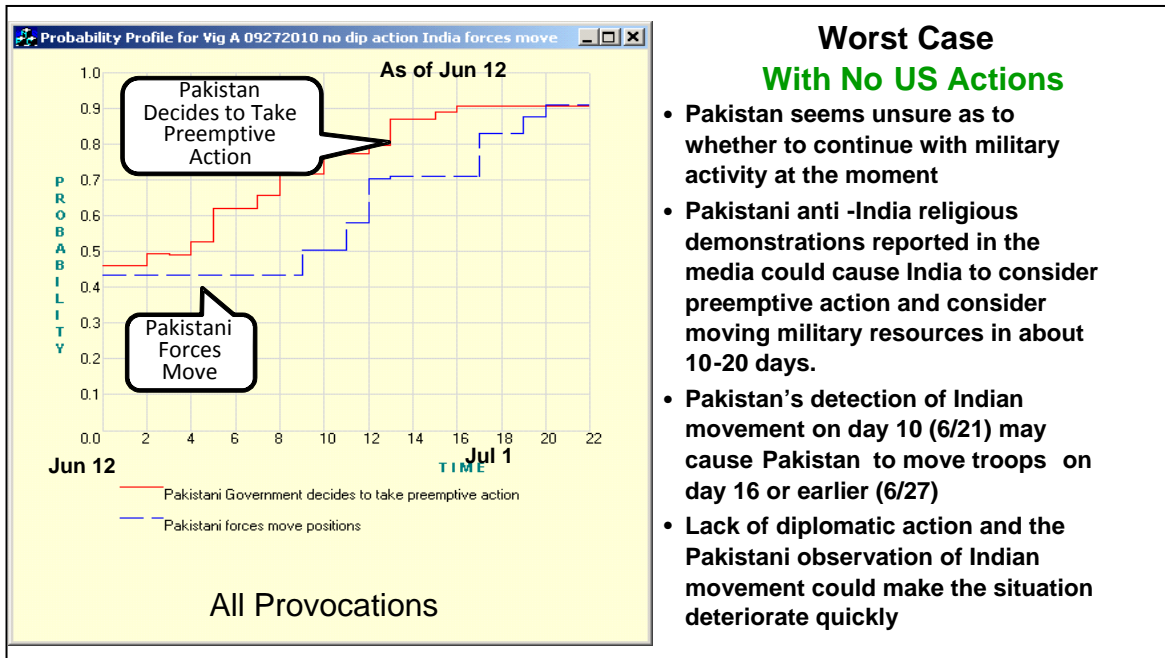


Fig. 14.13 Assessment of worse case situation

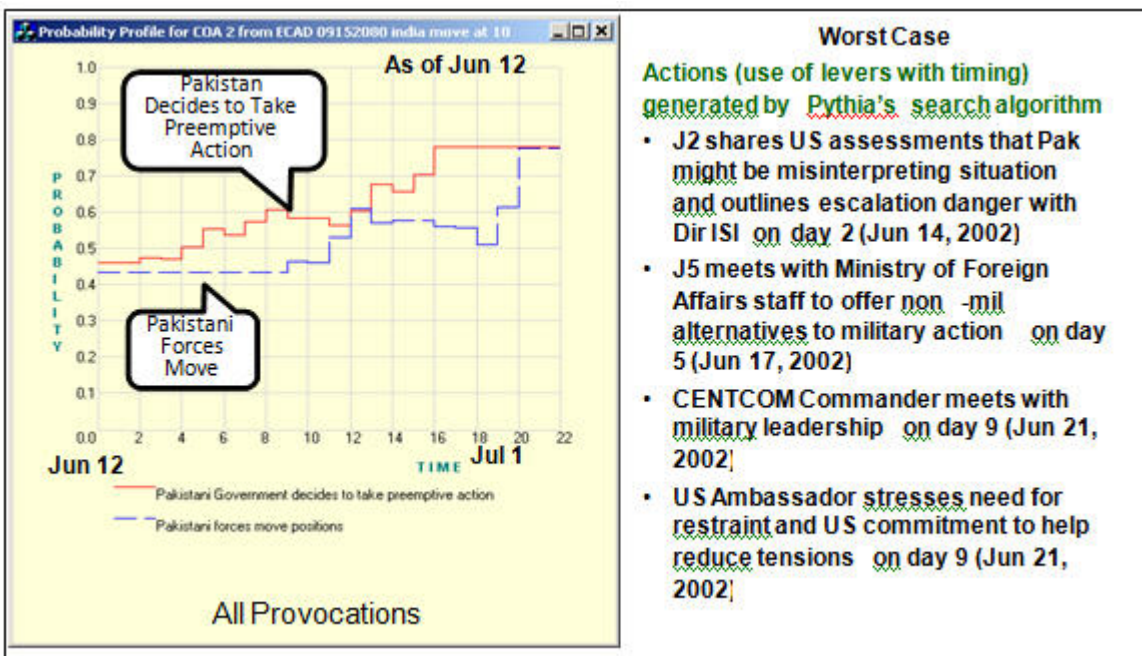


Fig. 14.14 Assessment using Evolutionary Search algorithm




Fig. 14.15 Improved probability profile by taking actions early

An analysis was done using the same actions and timing for the analysis shown in Fig. 14.15, but without the Ambassador actions. This was done to examine what the effect would be if the Ambassador was not able to hold his meeting. Note that it is the call of the US Department of State to direct the Ambassador to hold such a meeting, not CENTCOM. The results are shown in Fig. 14.16.

Finally an analysis was done to see the impact on Pakistan of India moving its forces. Figure 14.17 shows the results when Pakistan does not observe India troop movement. The results of this analysis clearly indicate the importance of trying to convince India to use restraint in its actions. The results of this analysis would likely be discussed with CENTCOM's counterparts in PACOM. CENTCOM should lobby for the Ambassador to participate in the course of action in a coordinated way.

Overall the assessment is that the situation bears considerable watching because there is a reasonable chance the situation could escalate if India moves forces. Even if they do not do that, there is still a possibility of escalation.

The analysis indicates that if there are few further provocations (particularly provocative movements of forces by either side), the situation will remain tense but under control. However if provocations continue, particularly movement of forces, then tensions will escalate. This escalation could occur within 12-20 days. The Pythia analysis suggests using all lower level levers as soon as possible to convince the Pakistani principals that they should follow non-military courses of action. Both military-to-military and the US Ambassador to his counterparts actions should occur. US should focus Intelligence Surveillance, and Reconnaissance (ISR) assets on possible movements of forces (of both Pakistan and India) as well as incidents such as riots, demonstrations, of acts of violence to keep the levers informed of true situation.




Fig. 14.16: No Ambassador involvement




Fig. 14.17: Effect of India not moving forces

A similar analysis was done for the PACOM perspective on India.

Figure 14.18 shows the PACOM Pythia model, and Figs. 14.19, 14.20, and 14.21, show the Worse Case, no action analysis and the best case where all actions are taken aggressively and Pakistan does not move its forces. The summary for India, from the PACOM perspective, was similar to that of CENTCOM for Pakistan: if no major provocations occur, then the situation will remain stable. However, if Pakistan decides to move her forces into the region near the Line of Control this could cause India to consider strong escalation. Messages by all levers should be coordinated and informed by ISR to ensure consistent messaging.

Actions and Event

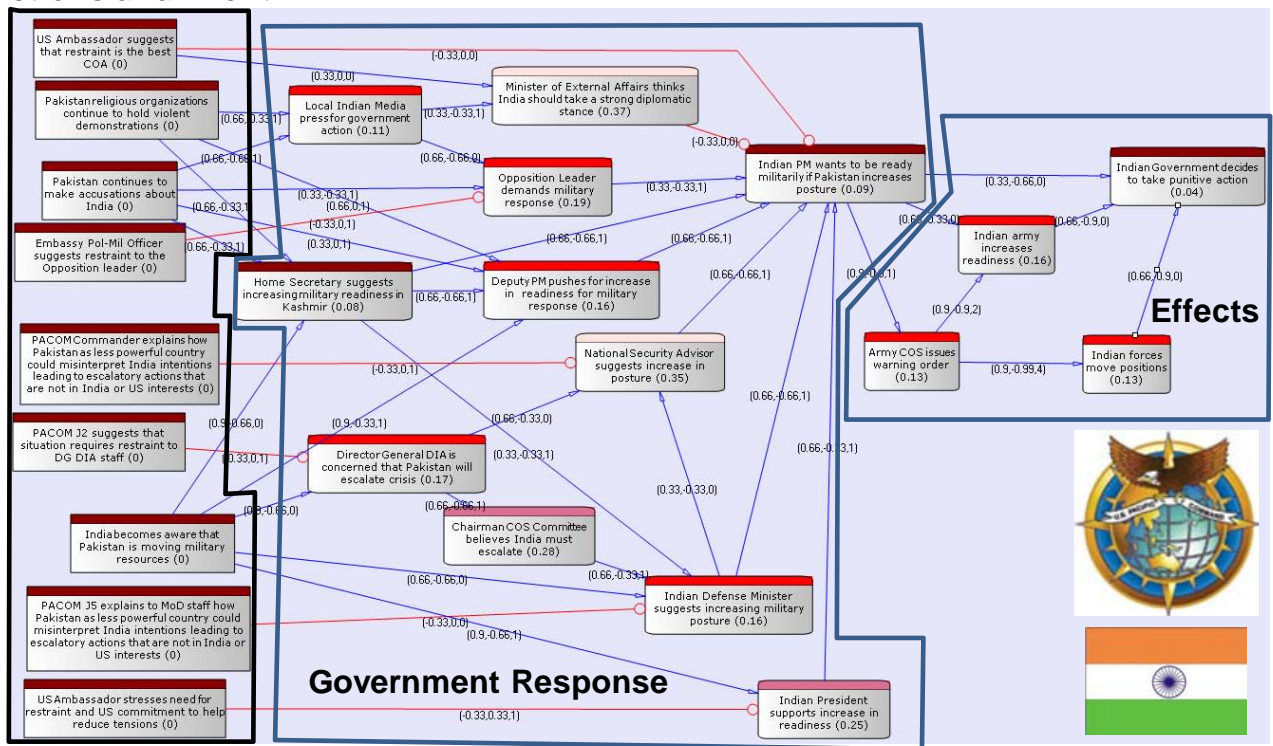




Fig. 14.18 PACOM Pythia model situation



Worst Case No US Action

- Pakistani accusations and demonstrations start to cause India to consider taking punitive action and may cause her to consider moving military resources in about 10 days.
- India detection of Pakistan movement on day 10 (6/21) causes India to potentially move troops on day 16 or earlier (6 days later (6/27))
- Lack of diplomatic action and India's observation of Pakistani movement makes the situation deteriorate quickly

Fig. 14.19 Probability profile for India



Worst case

Adjusted Possible Actions and Timing - accelerated consultations

- PACOM J2 suggests that situation requires restraint to DG DIA staff @1 vs. @4
- PACOM J5 interacts with MoD staff @2
- PACOM Commander expresses views @3 vs. @5
- US ambassador meets with Indian President @4 vs. @6
- US Ambassador suggests that restraint is the best COA (to Minister of External Affairs and the PM) @4 vs. @6

Fig. 14.20 PACOM analysis of situation with all actions

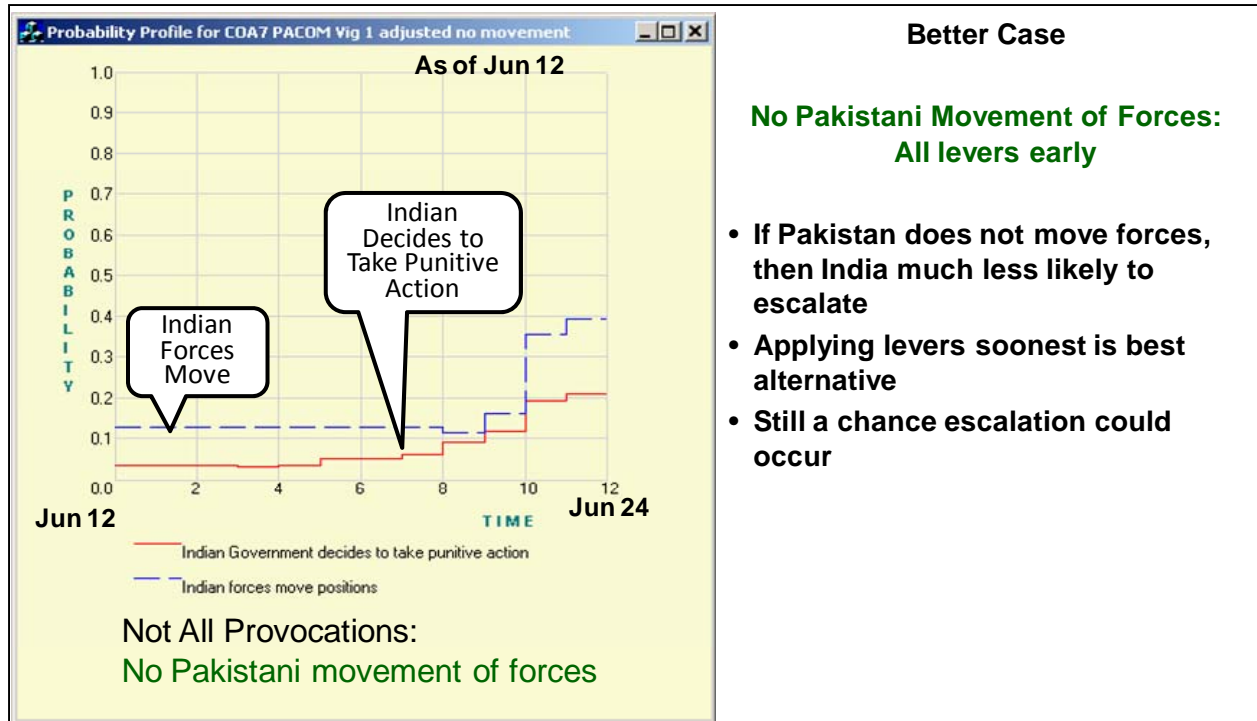


Fig. 14.21 PACOM analysis with no movement of Pakistani forces

14.4 Vignette B: Major escalation: COA development for regional deterrence/crisis de-escalation ops

Vignette B assumes that the COAs of Vignette A have not produced the desired results and the crisis is escalating. As part of this vignette, CENTCOM and PACOM observe the new developments in their areas of responsibility, assess the situation, and develop individual courses of action to de-escalate the crisis. The two commands use the same workflow as employed in Vignette A (given in Fig. 14.2) to develop their individual COAs; however, they develop new models and carry out new analyses in view of the events taking place as part of the scenario. Since the situation has escalated to a point where an armed conflict between the two nations seems eminent, the two commands merge their models for a combined situation assessment and develop a unified COA, as depicted in Fig. 14.22.

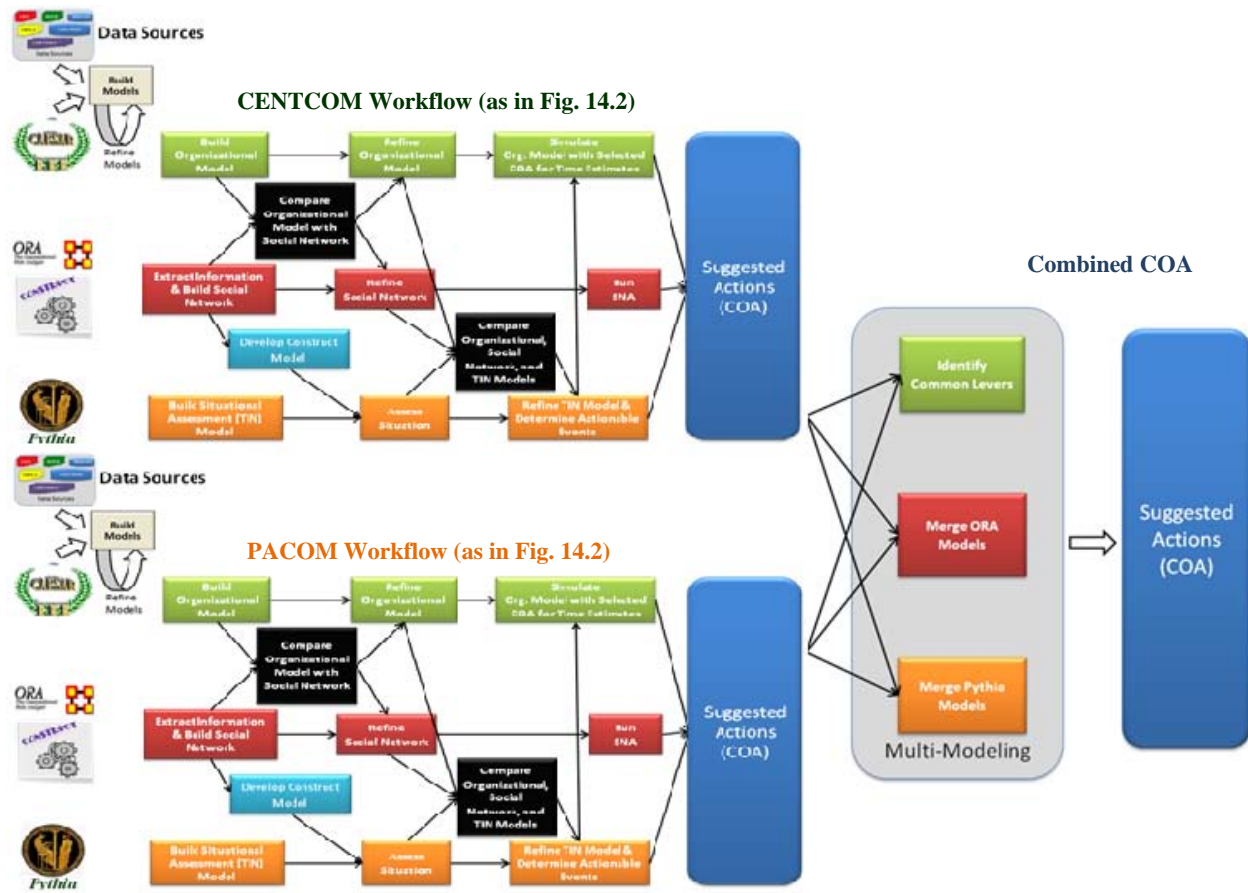


Fig. 14.22: Vignette-B workflow

The following is a detailed description of the data sources used, models constructed, and the analyses to assess the situation between India and Pakistan as part of Vignette B activities by the two commands for both developing their individual and merged COAs.

14.4.1 Data Sources

The data sources for Vignette B were identical to the methods listed for Vignette A. The differences between the data sets were in two of the three data sources: the scenario event list and the LexisNexis data pull. The LexisNexis data pull was for the inclusive dates of the vignettes, driving few minor changes in the network structure. The scenario event list was also different, in both the events listed as well as the locations of those events.

The summary of node types and their respective quantities are provided in Tables 14.8, 14.9, and 14.10.

TABLE 14.8 Vignette A, National Security Council only, CENTCOM & PACOM

Node Sets	CENTCOM (Pakistan)	PACOM (India)
Agents	23	42
BeliefError! Bookmark not defined.	32	32
Event	22	22
Knowledge	148	148
Location	3303	3303
Organization	325	325
Resource	117	117
Role	244	244
Task	424	424

TABLE 14.9 Vignette B, NSC and Diplomats only, CENTCOM & PACOM

Node Sets	CENTCOM (Pakistan)	PACOM (India)
Agents	47 India	111
Belief through Task	See Table 14.1	See Table 14.1

TABLE 14.10 Vignette B, all agents, CENTCOM & PACOM

Node Sets	CENTCOM (Pakistan) & PACOM (India)
Agents	111
Belief through Task	See Table 14.1

The analysis of the network data for Vignette B did not reveal any missing and relevant actors from the network created for Vignette A. The alignment of key agents is likely much more a function of how the researchers collected and pruned the collected data than may be realistic to expect in a real-world environment. As in Vignette A, Sphere of Influence reports, Betweenness Centrality, as well as Key Entities reports were used by researchers to ensure tight coupling with the CMU & GMU tool sets.

14.4.1.1 Sample Key Entity Reports for Vignette B

From the key entity reports, as well as reviews of key entities over time, researchers built a set of plausible underlying analysis. The overall theme of these reports, using the prominence and frequency of presence in the harvested data, is that US diplomatic circles start very high in centrality and influence. There is a downward trend reflected in the concurrent rise in influence and centrality of US military personages (e.g. Secretary of Defense, Chairman Joint Chiefs of Staff).

Figures 14.23 and 14.24 are the Top Ranked leaders histogram from the Vignette A key entity report. Figure 14.23 is from a CENTCOM (Pakistan) perspective, while Fig. 14.24 is from the PACOM (India) perspective. They show the Agents that are repeatedly top-ranked in the measures listed in Table 14.11. The value shown is the percentage of measures for which the Agent was ranked in the top three.




Fig. 14.23 Top ranked leaders, CENTCOM perspective




Fig. 14.24 Top ranked leaders, PACOM perspective

TABLE 14.11 Measures reflected in Key Entity tables

Acts as a Hub (hub centrality)	Leader of Strong Clique (eigenvector centrality)
Acts as an Authority (authority centrality)	Most Knowledge (row degree centrality)
Complete Exclusivity - event (complete exclusivity)	Most Resources (row degree centrality)
Complete Exclusivity - knowledge (complete exclusivity)	Number of Cliques (clique count)
Complete Exclusivity - location (complete exclusivity)	Potentially Influential (betweenness centrality)
Complete Exclusivity - resource (complete exclusivity)	Specialization - event (relatively unique)
Complete Exclusivity - task (complete exclusivity)	Specialization - knowledge (relatively unique)
Connects Groups (high betweenness and low degree)	Specialization - location (relatively unique)
Emergent Leader (cognitive demand)	Specialization - resource (relatively unique)
Group Awareness (shared situation awareness)	Specialization - task (relatively unique)
In-the-Know (total degree centrality)	Workload (actual based on knowledge and resource)

Figure 14.25 and the associated table reflect the agent x agent network of Pakistani and US agents. The table shows the top ten agents in potentially influence as reflected in their betweenness centrality score. Betweenness centrality is a measure that reflects how often a particular agent is in the shortest paths between all pairs of agent nodes in the network. These agents are therefore good to use to quickly pass messages to and from other agents in the network. The relative importance of the US Chairman of the Joint Chiefs of Staff is likely a reflection of the press' coverage of the US high-level military interest in Pakistan just a year after the attacks of September 11, 2001. The high ranking of the US Ambassador to Pakistan and Secretary of State are both reflective of frequent, and subjectively normal, appearance in the data sources given a rise in international tensions.

Figure 4.26 and the associated table reflect the agent x agent network of Indian and US agents. Like the table and figure for CENTCOM (Pakistan), the table shows the top ten agents in potentially influence as reflected in their betweenness centrality score. The relative importance of the President of the US and the Secretary of State reflect their frequent mention in the data sources. Again, the subjective evaluation of this network is of an expected network given the beginnings of an international crisis.

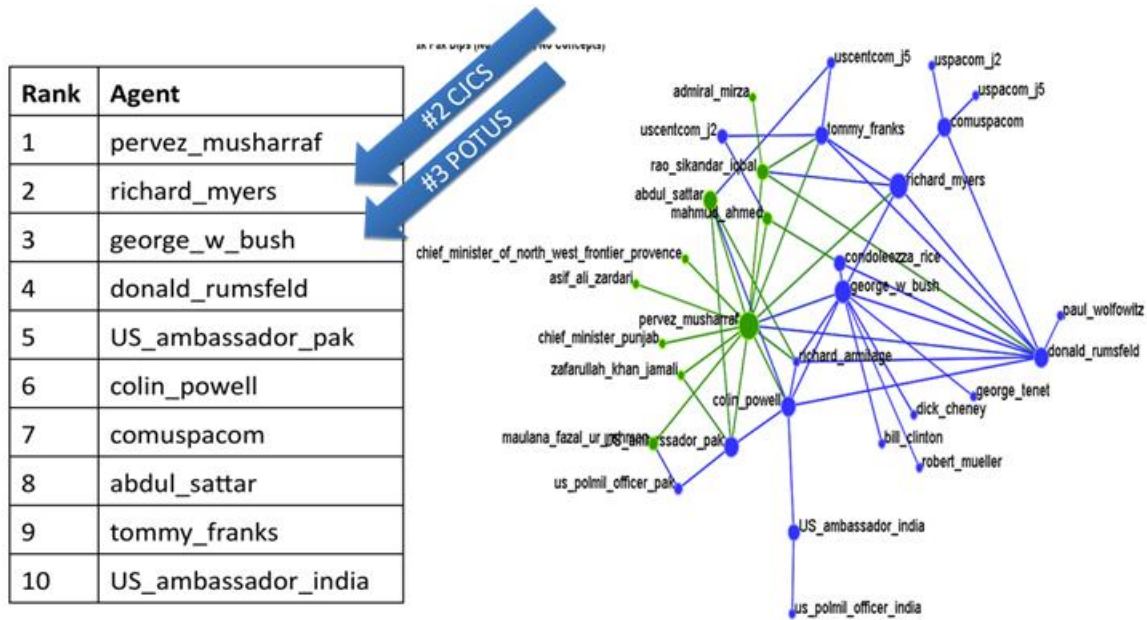


Fig. 14.25 Agent x Agent network of Pakistani and US agents

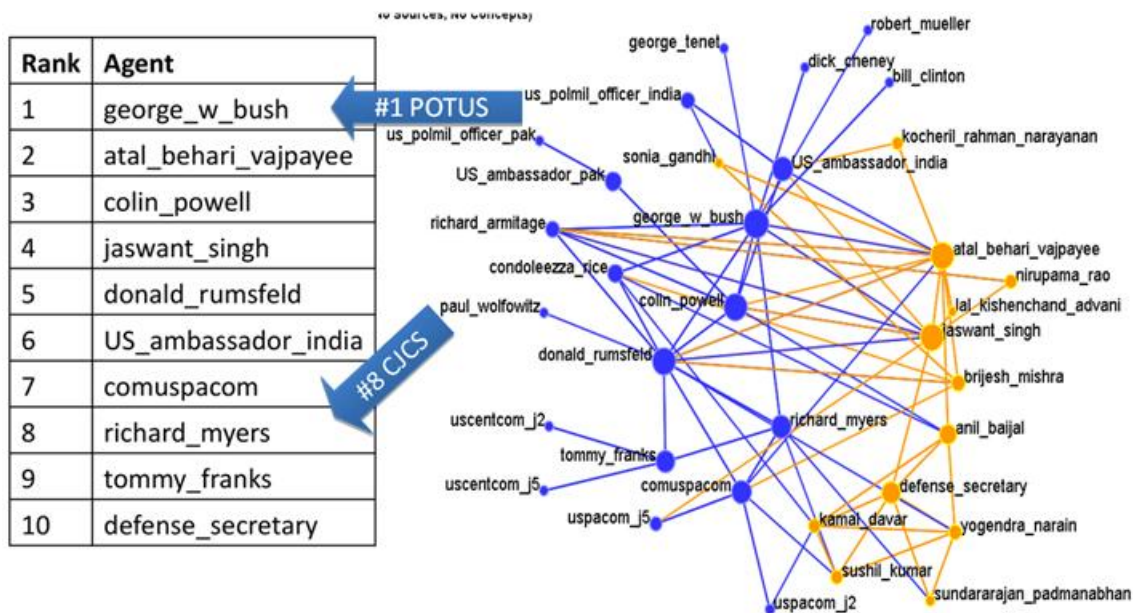


Fig. 14.26 Agent x Agent network of Indian and US agents

Figure 14.27 (a and b), from a CENTCOM (Pakistan) perspective, reflect CENTCOM's awareness that Pakistan is very aware of demonstrations in India, while at the same time Pakistan is issuing warnings to India to deescalate. A surprise to researchers was the relative importance of US actions in the midst of the rising tensions. This was our first indicator, that Construct later substantiated, that US military actions, even as innocuous as military intelligence, surveillance,

and reconnaissance flights, could inflame passions on both sides of the border crisis. The bottom figure is reflective of a common concern for military and security related knowledge, but more importantly, eight of the top nine top concepts of knowledge in the data set reflect non-militaristic concerns. This is consistent with the generalized statement that there are always many reasons to avoid war.




Fig. 14.27 CENTCOM Perspective of the situation

Figure 14.28 (a and b) reflect the PACOM (India) perspective. Of note is the mutual awareness of the other country's demonstrations as well as the mutual awareness of US military operations. There is another interesting observation in that the reports are not identical across the two combatant command's perspectives. Recall that the researchers had scraped common data, and spent the majority of time 'cleaning' the agent networks extracted from the unstructured text. Yet despite the common origins of the data, there was enough unique structure in each meta-network to derive slightly different versions of key events and knowledge.

Sample Key Entity Reports for Vignette B

With distinct time periods and a series of intervening events, there is an apparent need to compare and contrast data-driven network measures. The most striking observation for the top ranked leader figures is the rising importance of US military representatives (Fig. 14.29). This is not surprising given the close ties the US has attempted to forge with the Pakistani's since we invaded Afghanistan. Also given those close ties, it is not surprising to see the Director General of the ISI prominently figuring across the metrics listed in Section 14.3.2.




Fig. 14.28 PACOM Perspective of the situation




Fig. 14.29 Relative importance of top-ranked leaders

The US diplomatic corps, from the PACOM (India) perspective, is still the dominant force when engaging with Indian national decision makers. It is clear however that the level of authority being brought to bear is significantly higher, as the Secretary of Defense has increased to the second ranking and the Secretary of State has risen into the top ten while his deputy dropped four positions. The concurrent decrease in importance of the opposition party can also be reflective of a hardening Indian viewpoint, with the opposition leader no longer in a central advisory role to the Indian Prime Minister.

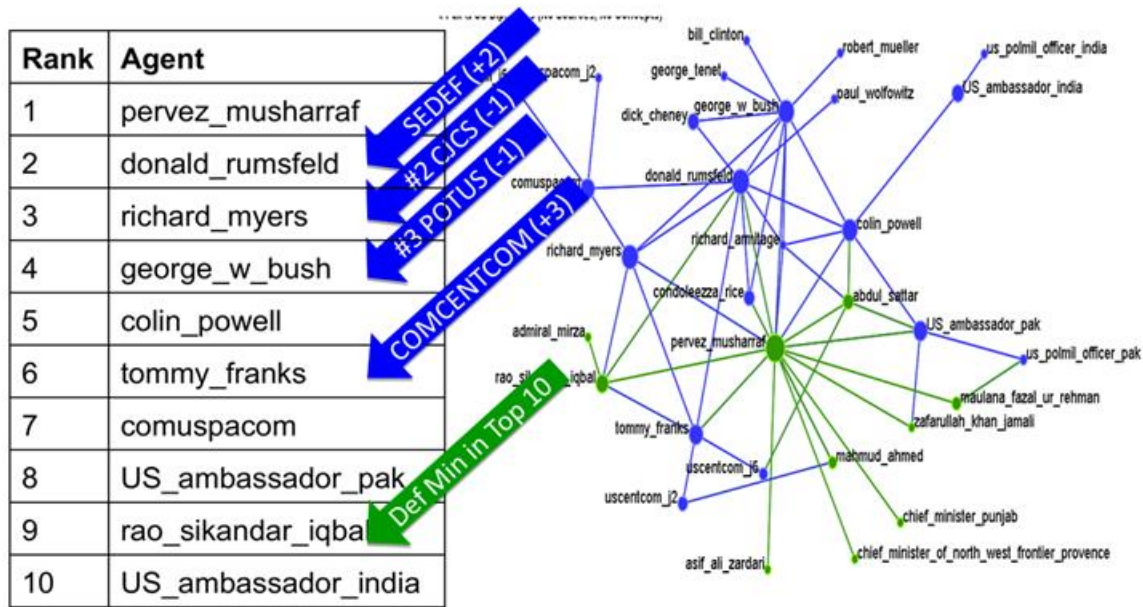


Fig. 14.30 Agent x Agent network of US and Pakistani agents

As noted in Vignette A, Fig. 14.30 and associated table display the agent x agent network of US and Pakistani agents though now it is applicable to Vignette B. Note, the increased prominence of the Secretary of Defense as an agent high in betweenness centrality and therefore a potentially influential agent.

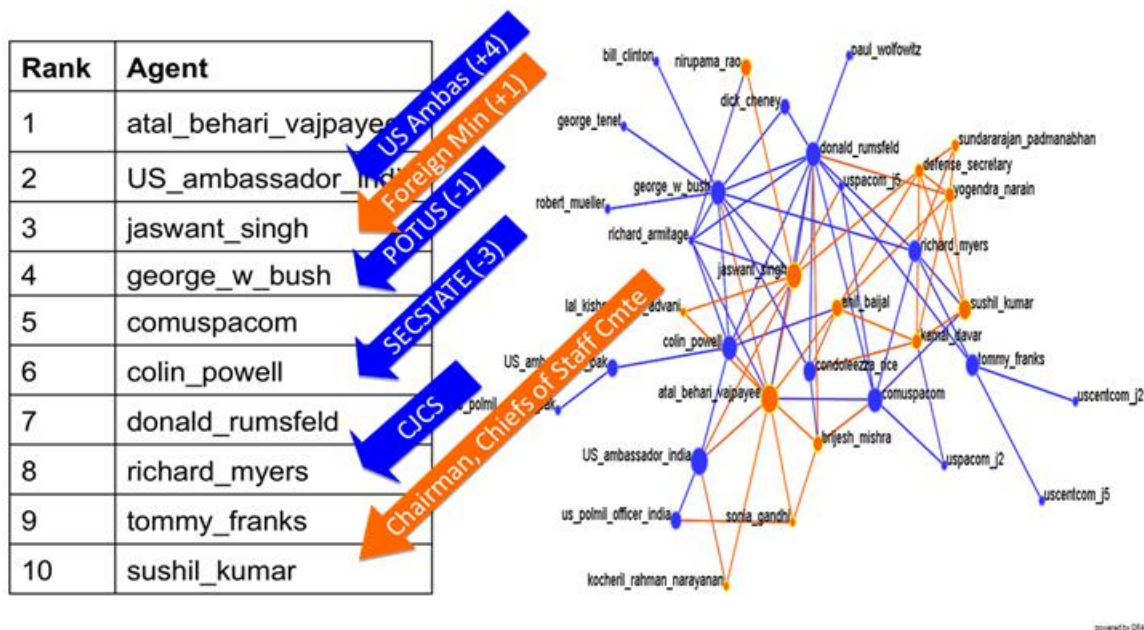


Fig. 14.31 Agent x Agent Network of US and Indian agents

Also for Vignette B, Fig. 14.31 and the associated table and figure reflect the changes from the PACOM (India) perspective. Here, we see that the relative importance of the Indian Foreign Minister, and the continued presence of the Secretary of State can be an indicator that military and military-minded agents are not yet dominating the Indian viewpoint in this crisis.



Fig. 14.32 Key events from the CENTCOM (Pakistan) perspective

As Vignette B proceeds, these key events and key knowledge figures show that the Indian Air Forces deployment has become a key event from the CENTCOM (Pakistan) perspective. Additionally, even though military related knowledge is now two of ten top knowledge events, economics is still frequently mentioned concern in the data sources. (Fig. 14.32)



Fig. 14.33 Key Events from the PACOM (India) perspective

The figures (Fig. 14.33) from the PACOM (India) perspective show a rising awareness of Pakistani troop movements. Of note, is the scenario-based warnings by the Indian government to its citizens in Pakistan that they should leave Pakistan for their own safety. US ISR flights are still in the top seven entries for key events—reinforcing the idea that US reconnaissance activities may actual increase tensions between the two countries.

14.4.2 Construct Model

For Vignette B, Indian and Pakistani decision makers were modeled as Construct agents (a brief introduction of Construct was included in Vignette A, above). Each of the provocations (as shown in Table 14.12) represented a source of knowledge that contributed to a pro-war belief. Provocations represented a source of knowledge that contribute to a pro-peace (that is, anti-war) mentality. Agents started with random distributions of both pro-war and pro-peace knowledge, and were allowed to interact (their interactions were constrained by the found social network in the data sources). Additional knowledge bits that did not influence pro-war belief were used to provide points of contrast for similarity measures. At various points through the time-course simulation, the number of agents who had the pro-war belief was measured.

Different from the goals in Vignette A, in Vignette B the interest was in the impact of response time on the number of decision-makers with the pro-war belief. Four different response times were used, and it was assumed that all the responses were coordinated for that chosen time-point. It was also assumed that all provocations detailed in the scenario occurred, even if tensions were very low or non-existent. As such, these are conservative projections of impact. The experimental design is shown in Table 4.12.

TABLE 14.12 Construct experiment design, Vignette B

Parameter	Number of Values	Values
Agent Networks	2	PACOM, CENTCOM
Responses	4	None, Day 1, Day 14, Day 20
Total Knowledge Facts	1	450
Total “Similarity” Facts	1	300
Total “Pro War” Facts	1	50
Total “Pro Peace” Facts	1	100

Each of the condition sets were run for one hundred trials, for a total of 800 simulation runs. In the set of graphs in Fig. 14.34, the dashed lines represents the average value of the “no response” case over the time-course, meaning that no actions were taken to defuse the situation in those runs. The solid lines represent the average value of the response case shown, each figure represents a different response-time. The marks show the values of individual runs over the simulation’s time-course for the response case. The general assessment of these charts presents no surprise, the earlier the response, the better. Early injections of “pro-peace” knowledge causes agents to share that knowledge early and often, presenting an outsized impact on the overall trends.

14.4.3 Organizational Models

In Vignette B, the situation between India and Pakistan intensifies. Therefore, US levers with high-level contacts with the government officials are brought in to deescalate the impending crisis. The organization models of the two countries' security structures remain the same as depicted in Figs. 14.9 and 14.10. The only refinement made to these models in Vignette B is the addition of new high-level US levers and their interactions with the government officials of the two countries. Figure 14.35 displays the organization model of Pakistan with a new US lever (Chairman, Joint Chief of Staff). The organization model of Pakistan NSC itself remains unchanged. Similarly, Fig. 14.36 shows the model of India NSC with new US levers (Chairman, Joint Chiefs of Staff and National Security Advisor).

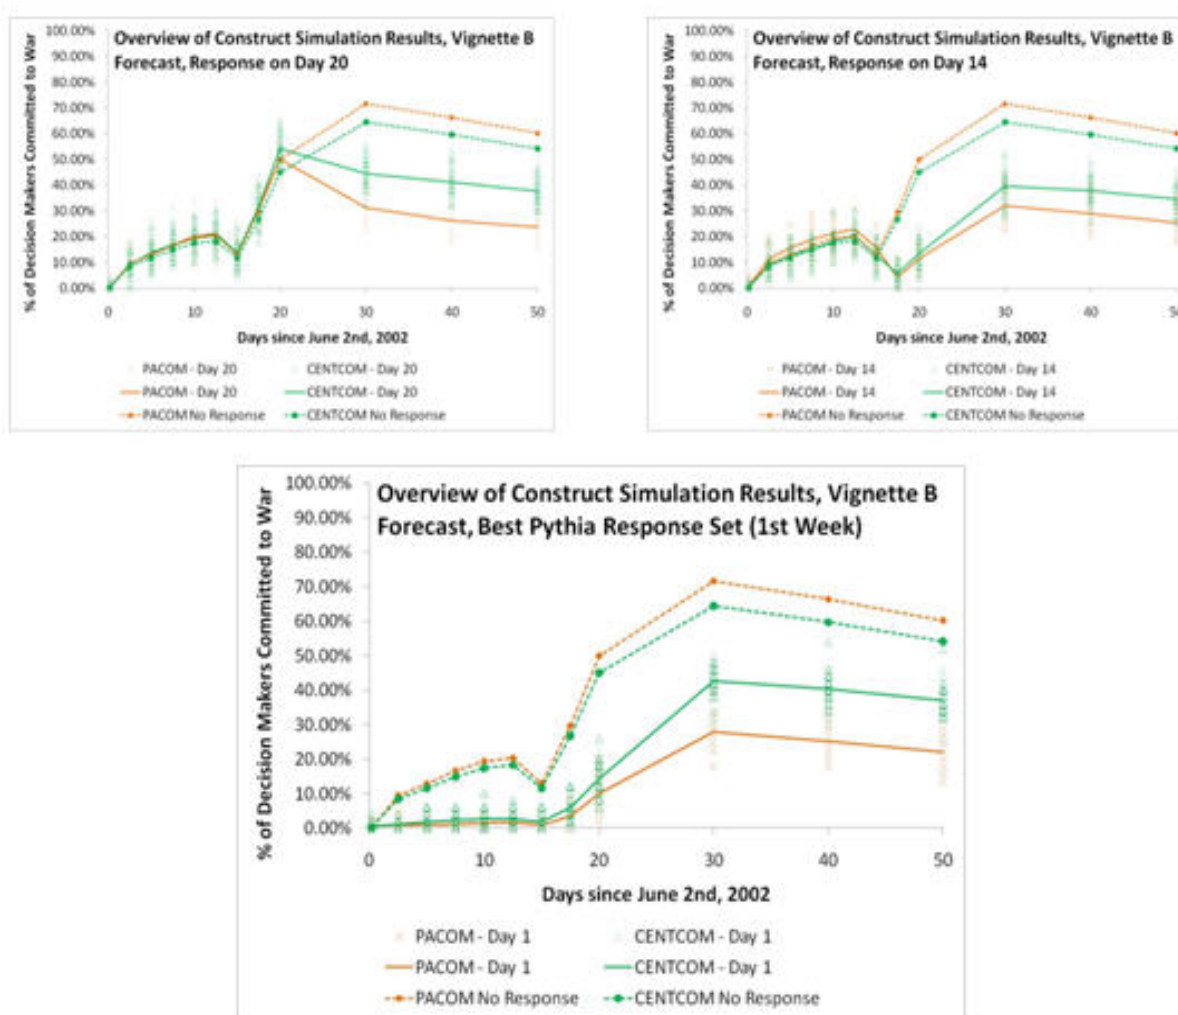


Fig. 14.34 Comparison of the base case (“No Reponse”) to reponses occurring at specific points in the simulation's time-course

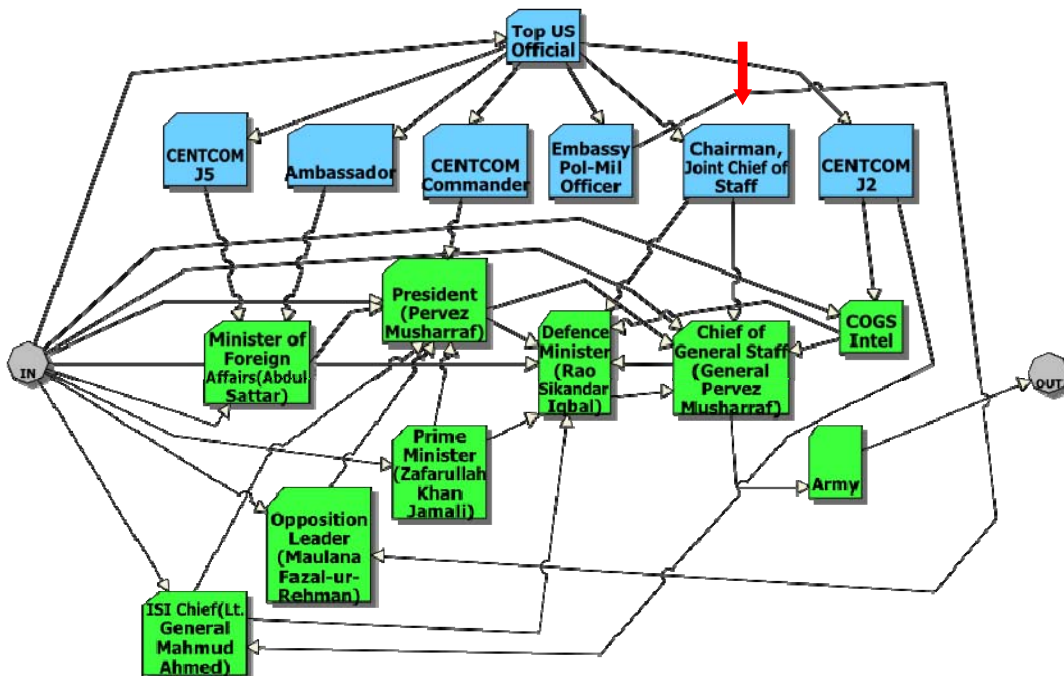


Fig. 14.35 Pakistani Government organization model for Vignette B

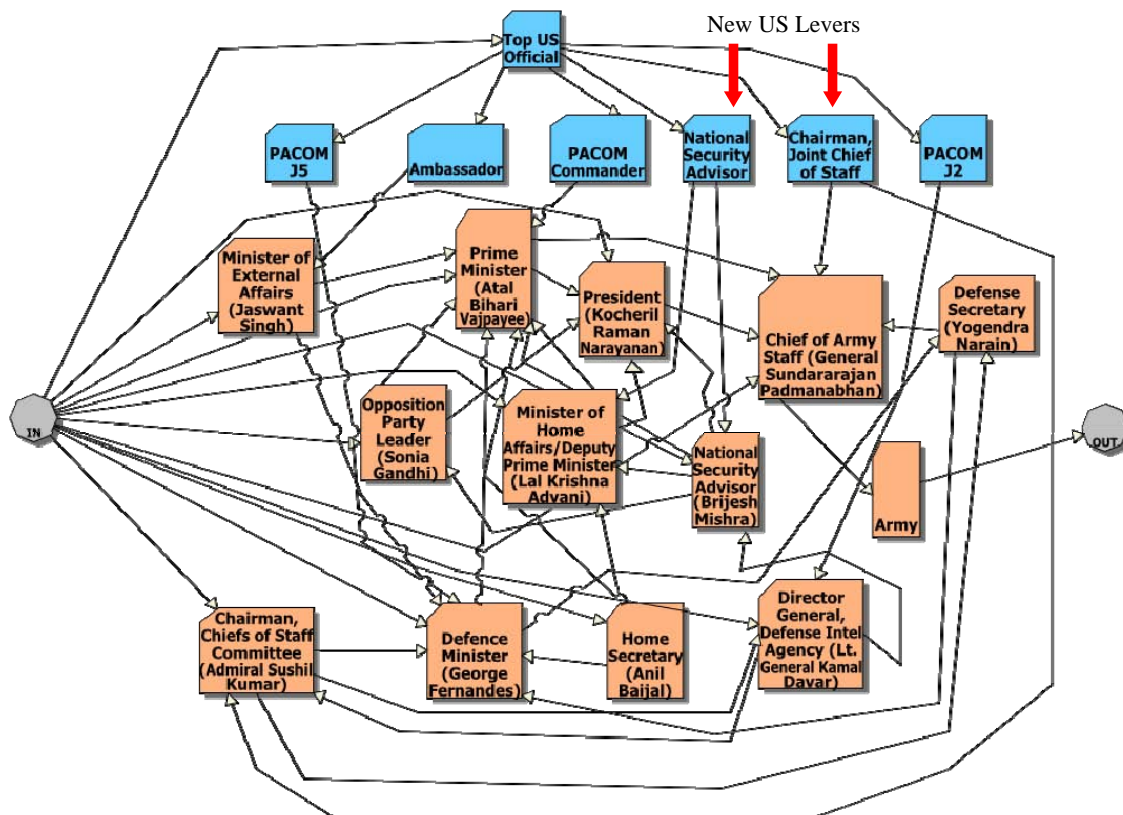


Fig. 14.36 Indian Government organization model for Vignette B

14.4.4 Analysis of Organization Models

Tables 14.13 and 14.14 show the influence of new levers into the Pakistani and Indian Government models.

TABLE 14.13 CENTCOM sphere of influence report for new US lever

Lever	Influenced DMs
Chairman, Joint Chief of Staff	{Defense Minister, Chief of General Staff} → Army

TABLE 14.14 PACOM sphere of influence report for new US levers

Lever	Influenced DMs
National Security Advisor	{National Security Advisor,, Minister of Home Affairs/Deputy Prime Minister} → {Opposition Party Leader, Prime Minister} → {President, Chief of Army Staff} → Army
Chairman, Joint Chief of Staff	{Chairman Chiefs of Staff Committee, National Security Advisor} → {Defense Minister, Director General DIA} → Defense Secretary → Minister of Home Affairs/Deputy Prime Minister → Opposition Party Leader → Prime Minister → {President, Chief of Army Staff} → Army

Table 14.15 displays the sphere of influence report for the common US levers across both organization models developed by the two commands.

TABLE 14.15 Sphere of influence of common levers

Lever	Influenced DMs
Ambassador (CENTCOM)	Minister of Foreign Affairs → President → {Defense Minister, Chief of General Staff} → Army
Chairman, Joint Chief of Staff (CENTCOM)	{Defense Minister, Chief of General Staff} → Army
Ambassador (PACOM)	Minister of External Affairs → Prime Minister → {President, Chief of Army Staff} → Army
Chairman, Joint Chief of Staff (PACOM)	{Chairman Chiefs of Staff Committee, National Security Advisor} → {Defense Minister, Director General DIA} → Defense Secretary → Minister of Home Affairs/Deputy Prime Minister → Opposition Party Leader → Prime Minister → {President, Chief of Army Staff} → Army

The Pythia models for both CENTCOM and PACOM were updated on the scenario date of June 30, 2002 to reflect the new situation in Vignette B. The possibility that both nations could begin preparation of their nuclear forces for possible use was added to the models as well as the possibility of having high level talks between the two nations. New levers were added to each model (Fig. 14.37 and 14.38). The CENTCOM Pythia model is shown in Fig. 14.37. Key events include “Pakistan Religious organizations continue protests” starting 07/01/2002, “Indian officials continue to make accusations about Pakistan” on the same day, “Indian Political Leaders call for action against Pakistan” immediately thereafter, and “India and Pakistan patrols and forces exchange fire along border” during 07/06/2002. Key effects are “Pakistan is moving to Escalate”, “Pakistani forces engage in large scale skirmishes along the border”, and “Pakistani nuclear facilities begin preparation for the use of Nucs (nuclear weapons)”. The messages of the Levers are upgraded and the US Chairman of the Joint Chiefs of Staff is added as a lever.



112

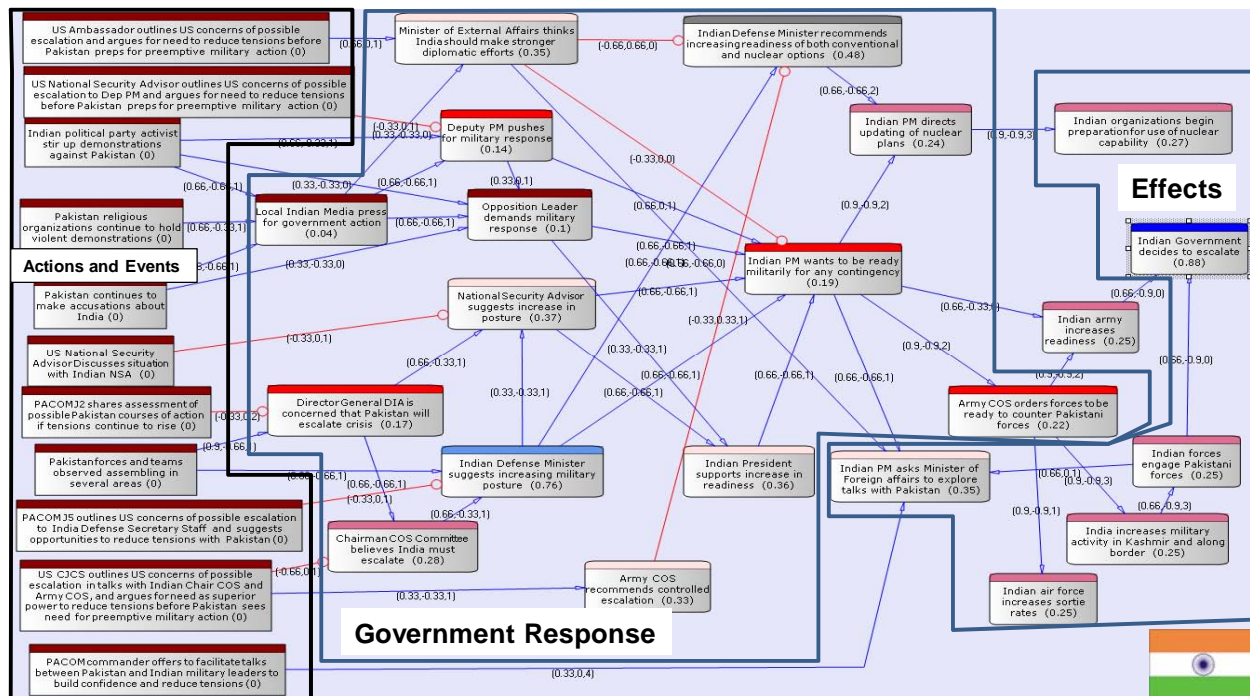
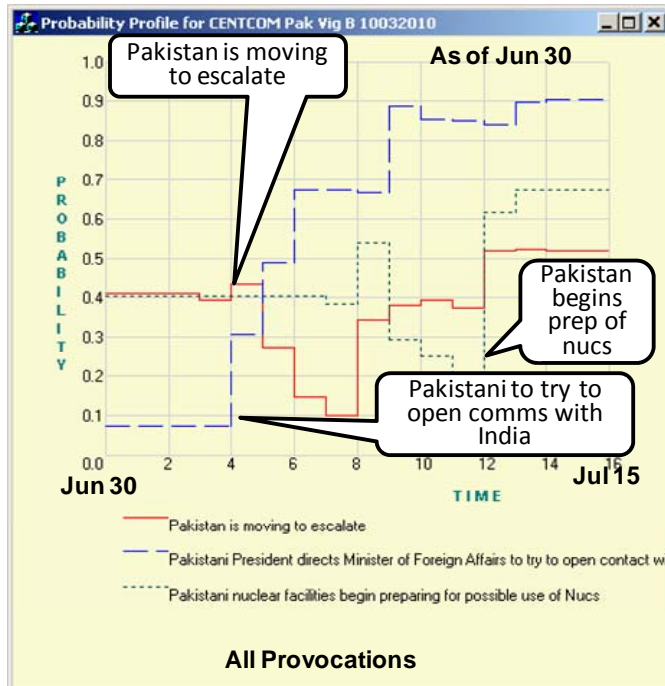


Fig. 14.38 PACOM Pythia model as of June 30, 2002

14.4.6 Assess Situation and COA Analysis

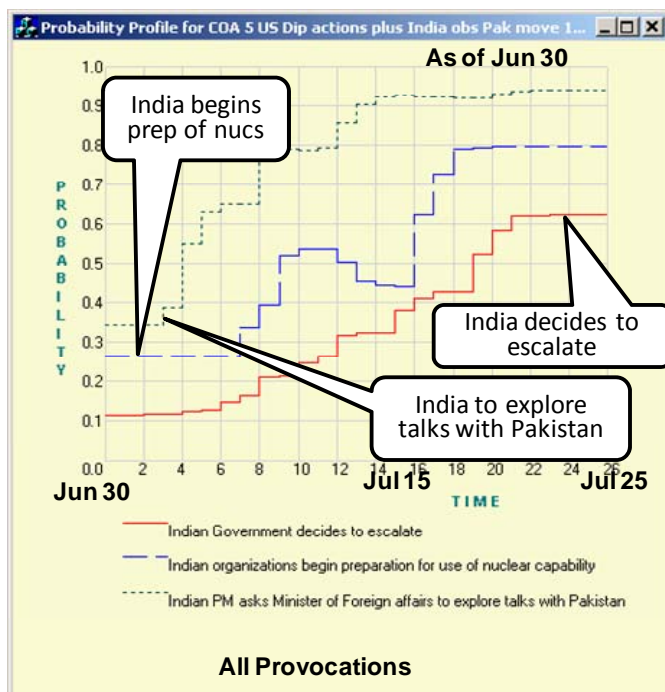
An analysis that was similar to the one performed for Vignette A was conducted for both models. The final probability profiles generated along with the suggested COA for the CENTCOM and PACOM models are shown in Figs. 14.39 and 14.40. Both analyses indicate that early use of levers may cause both sides to be willing to open talks with one another, but also show that movement of troops will likely cause both sides to escalate preparation for possible nuclear confrontation.

At this point CENTCOM and PACOM agree to merge their models for a combined analysis. The combined model is shown in Fig. 14.41 with the PACOM model on the top and CENTCOM model on the bottom. In this model, the node in the CENTCOM model that represents Pakistani Chief of General Staff ordering Pakistani army to move forces to the border is connected to the PACOM model node that represents the event of India observing Pakistani forces assembling in several areas (see darken arrow). A new node (circled) is added that represents the Pakistani and Indian ministers of foreign affairs agreeing to have talks. The type of analysis that was done for the individual models was done for the fused model. The messaging and timing of the actions by the levers was synchronized and the model was run showing the probability of key effects over time. The result after adjusting the timing of the events to create the “best” probability profiles is shown in Fig. 14.42



- Suggested COA for Worst Case**
- US Pol-Mil outlines US concerns of possible escalation and suggests opportunities to reduce tensions @ (07/01/2002)
 - CENTCOM J5 gives same message to MoFA staff @ (07/01/2002)
 - CENTCOM J2 shares assessment of possible India courses of action if tensions continue to rise @ (07/02/2002)
 - CENTCOM Commander suggests India and Pakistan hold meeting to resolve situation @ (07/03/2002)
 - US Ambassador outlines US concerns of possible escalation and argues for need to reduce tensions before India takes its own military action @ (07/03/2002)
 - US CJCS outlines US concerns of possible escalation to Minister of Defence (Pakistan) and Chief, General Staff, and argues for need to reduce tensions before India takes action to blunt their perception of imminent preemptive attack @ (07/05/2002)

Fig. 14.39 CENTCOM analysis for Vignette B



- Suggested COA for Worst Case**
- Adjusted Possible Actions and Timing**
- PACOM J2 shares assessment of possible Pakistan COAs if tensions continue to rise @ (07/02/2002)
 - PACOM J5 outlines US concerns of possible escalation to India Defense Secretary Staff and suggests opportunities to reduce tensions with Pakistan @ (07/02/2002)
 - PACOM Commander offers to facilitate talks between Pakistan and Indian military leaders @ (07/04/2002)
 - US National Security Advisor discusses situation with Indian NSA @ (07/05/2002)
 - US Ambassador outlines US concerns of possible escalation and argues for need to reduce tensions before Pakistan preps for preemptive military action @ (07/05/2002)
 - US CJCS outlines US concerns of possible escalation in talks with Indian Chair COS and Army COS, and argues for need to reduce tensions before Pakistan sees need for preemptive military action @ (07/06/2002)
 - US National Security Advisor re-enforces US CJCS message with Deputy Prime Minister @ (07/09/2002)

Fig. 14.40 CENTCOM analysis for Vignette B

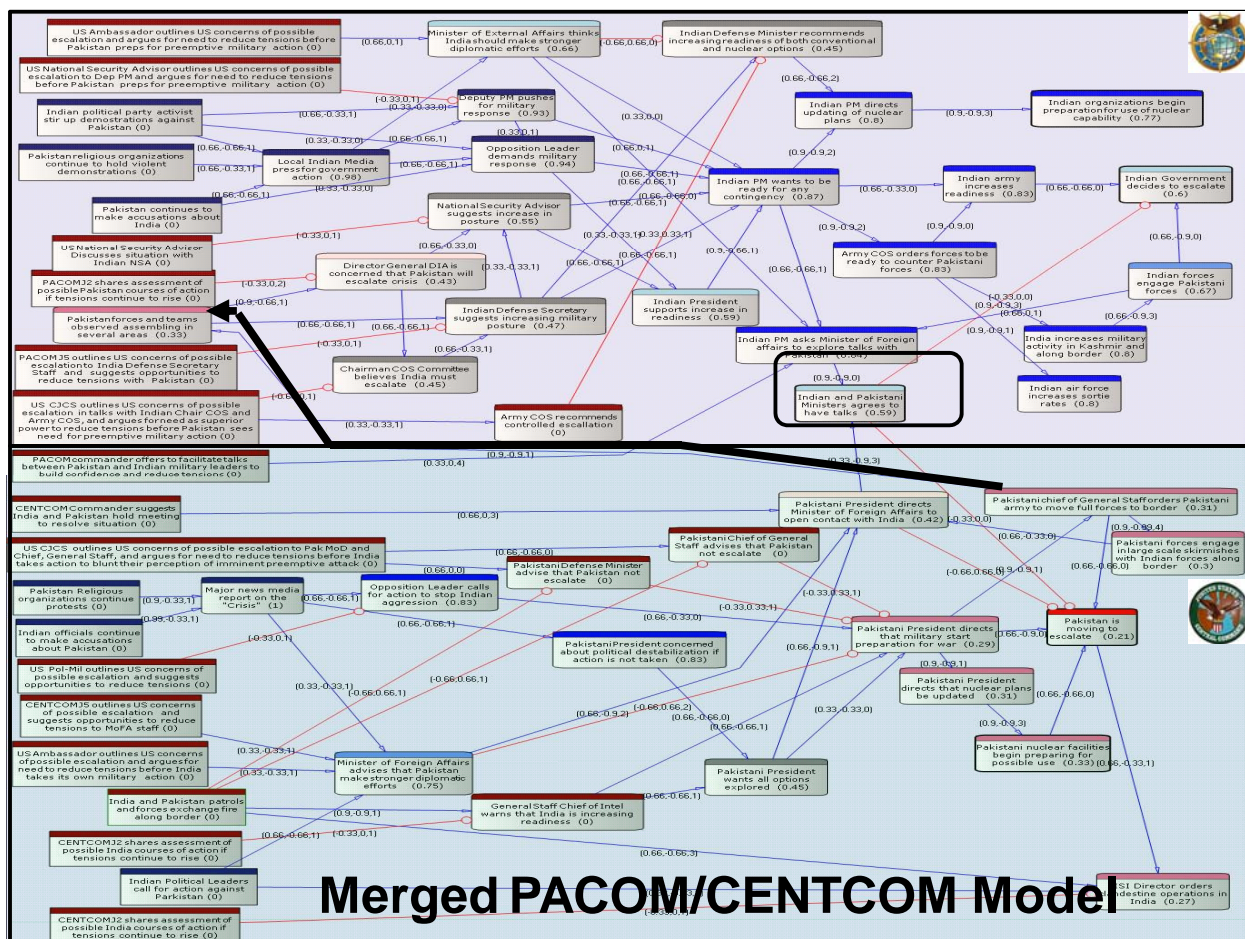
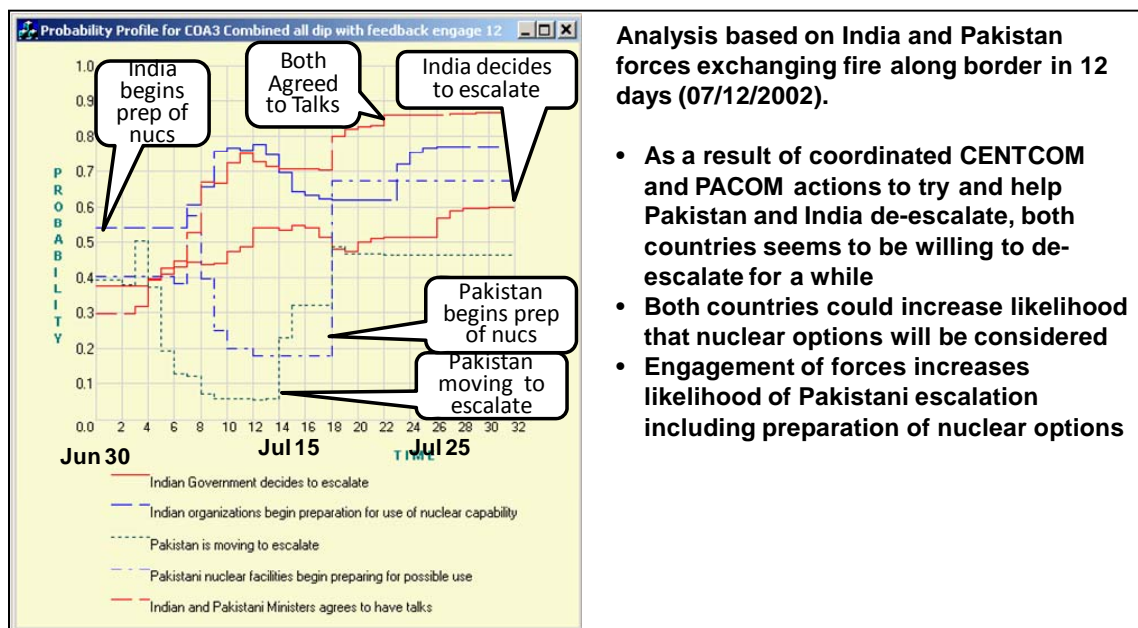


Fig. 14.41 Combined Pythia model



Analysis based on India and Pakistan forces exchanging fire along border in 12 days (07/12/2002).

- As a result of coordinated CENTCOM and PACOM actions to try and help Pakistan and India de-escalate, both countries seems to be willing to de-escalate for a while
- Both countries could increase likelihood that nuclear options will be considered
- Engagement of forces increases likelihood of Pakistani escalation including preparation of nuclear options

Fig. 14.42 Probability profiles for combined model

14.4.7 Merged COA

The analysis indicated that the best COA be based on the following concepts: (1) To use the lower level levers first, (2) Attempt to have the same type of interaction occur simultaneously in both countries, (3) Military to Military contacts occur first, (4) Pol Mil officer can also make early contact and the Ambassadors follow, and (5) then the more senior military such as Chairman Joint Chiefs and US National Security Advisor have their interactions.

The final COA based on the analysis with the combined Pythia model is shown In Table 14.16.

TABLE 14.16 Final COA for combined CENTCOM PACOM actions

Action	Day
US Pol-Mil outlines US concerns of possible escalation and suggests opportunities to reduce tensions to Pakistani Opposition Leader	30
CENTCOM J5 outlines US concerns of possible escalation and suggests opportunities to reduce tensions to Minister of Foreign Affairs staff	31*
CENTCOM J2 shares assessment of possible India courses of action if tensions continue to rise with staff of the Chief of General Staff for Intel and the Staff of ISI	31
PACOM J5 outlines US concerns of possible escalation to India Defense Secretary Staff and suggests opportunities to reduce tensions with Pakistan	31
PACOM J2 shares assessment of possible Pakistan courses of action if tensions continue to rise with the Director General of the Defense Intelligence Agency	31
CENTCOM Commander suggests India and Pakistan hold meeting to resolve situation to Pakistani President	33*
PACOM commander offers to facilitate talks between Pakistan and Indian military leaders to build confidence and reduce tensions	33
US CJCS outlines US concerns of possible escalation to Pakistan Minister of Defence and Chief, General Staff, and argues for need to reduce tensions before India takes action to blunt their perception of imminent preemptive attack	34
US CJCS outlines US concerns of possible escalation in talks with Indian Chair COS and Army COS, and argues for need as superior power to reduce tensions before Pakistan sees need for preemptive military action	34*
US Ambassador outlines US concerns of possible escalation and argues for need to reduce tensions before India takes its own military action to Pakistan Minister of Foreign Affairs	34*
US Ambassador outlines US concerns of possible escalation and argues for need to reduce tensions before Pakistan preps for preemptive military action to the Indian Minister of External Affairs	34*
US National Security Advisor outlines US concerns of possible escalation to Deputy Prime Minister and argues for need to reduce tensions before Pakistan preps for preemptive military action @9	36*

* Means the timing has been adjusted from the individual model analysis

The following summary of the assessment based on the analysis of the combined model was created. Both COCOMs should take coordinated actions to use all levers in a synchronized sequence as soon as possible. Messages include (a) US concerns of possible escalation and suggestion of opportunities to reduce tensions, (b) US will help facilitate meetings between appropriate parties from the two countries (c) CJCS should outline US concerns of possible escalation with both countries and with Pakistani senior official and argue for need to reduce tensions before India takes action to blunt their perception of imminent preemptive attack, (d) the CJCS should discuss with India the need as superior power to reduce tensions before Pakistan sees the need for preemptive military action, (e) US Ambassadors should make appropriate arguments with

their contacts, and (f) the US National Security Advisor should outline US concerns of possible escalation to each country. Furthermore the models indicate US only actions may be insufficient; therefore the US should pursue support from other countries who are friends with India and Pakistan.

14.5 Conclusions

The scenario increased the understanding of the utility of Pythia and timed influence net modeling and analysis to support Situation Assessment and Course of Action (COA) Evaluation in situations where deterrence is a key objective. The exercise demonstrated how Pythia could provide insights regarding possible COAs to maintain stability and control escalation. The multi-model enhanced Pythia Timed Influence Net models were used to identify potential adversary influence levers, their associated activities and timing, and gaps in information. Pythia enabled modeling of the influence levers and their potential effects on organizational perceptions and decisions of key individuals within both governments. The models included consideration of the messages that would be used by the levers and the impact of combined messages on two adversaries over time. Computational experimentation with varied timing enabled the ordering of multi-state activities for maximum combined effect. In the scenario the models demonstrated the value of early intervention (despite potential costs). In the end the models showed that the international community, particularly friends of the two adversaries, may be required to effectively influence behaviors. It appears that Pythia, along with the other models used in the two vignettes, could enhance the analysis of alternatives in command centers that are engaged in deterrence operations.

References

- [1] L.W. Wagenhals, A. H. Levis, and Maris McCrabb, "Effects Based Operations: a Historical Perspective for a Way Ahead," Proc. 8th Int'l Command and Control Research and Technology Symposium, NDU, June 2003
- [2] Jensen, F. V., *Bayesian Networks and Decision Graphs*, Springer-Verlag, 2001.
- [3] Neapolitan, R. E., *Learning Bayesian Networks*, Prentice Hall, 2003.
- [4] Chang, K. C., Lehner, P. E., Levis, A. H., Zaidi, S. A. K., and Zhao, X., "On Causal Influence Logic", Technical Report, George Mason University, Center of Excellence for C3I, Dec. 1994.
- [5] Rosen, J. A., and Smith, W. L., "Influence Net Modeling with Causal Strengths: An Evolutionary Approach", Proceedings of the Command and Control Research and Technology Symposium, Naval Post Graduate School, Monterey CA, Jun. 1996.
- [6] Agosta, J. M., "Conditional Inter-Causally Independent Node Distributions, a Property of Noisy-OR Models", In Proceedings of the 7th Conference on Uncertainty in Artificial Intelligence, 1991.
- [7] Drudzel, M. J., and Henrion, M., "Intercausal Reasoning with Uninstantiated Ancestor Nodes", In the Proceedings of the 9th Conference on Uncertainty in Artificial Intelligence, 1993.
- [8] Pearl, J. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, CA: Morgan Kaufmann, 1988.
- [9] J. A. Rosen and W. L. Smith, "Influencing Global Situations: A Collaborative Approach," US Air Force Air Chronicles, 1996.
- [10] J. A. Rosen, W. L. Smith, E. S. Smith, and M. A. Maldony, "Planning with Influence Net Modeling: An Architecture for Distributed, Real-Time Collaboration," Military Operations Research Symposium (MORS), 1998.
- [11] J. A. Rosen and W. L. Smith, "Influence Net Modeling for Strategic Planning: A Structured Approach for Information Operations," Phalanx, vol. 33, 2000.
- [12] M. Faraz Rafi , A. K. Zaidi , P. Papantoni-Kazakos and A. H. Levis "Optimization of Actions in Activation Timed Influence Nets" Informatica, 2009.
- [13] L. W. Wagenhals, A. H. Levis, "Course of Action Development and Evaluation," in Proceedings of the 2000 Command and Control Research and Technology Symposium, 2000.
- [14] L. W. Wagenhals, I. Shin, and A. H. Levis, "Creating Executable Models of Influence Nets with Coloured Petri Nets," Int. J. STTT, Spring-Verlag, Vol. 1998, No. 2, 1998.
- [15] L. W. Wagenhals and A. H. Levis, "Modeling Effects Based Operations in Support of War Games," in Proc. Of the 15th International Symposium on Aerospace / Defense Sensing, Internal Society for Optical Engineering, Proceedings of SPIE, Vol. # 4367, 2001.

- [16] L. W. Wagenhals and A. H. Levis, "Modeling Support of Effects-Based Operations in War Games," in Proc. 2002 Command and Control Research and Technology Symposium, Monterey, CA, June 2002.
- [17] L. W. Wagenhals and L. K. Wentz, "New Effects-Based Operations Models in War Games," in Proceedings of the 2003 International Command and Control Research and Technology Symposium, National Defense University, Washington, DC, June 2003.
- [18] L. W. Wagenhals, T. J. Reid, R. J. Smillie, & A. H. Levis, "Course of Action Analysis for Coalition Operations," in Proceedings of 6th International Command and Control Research and Technology Symposium, Annapolis, Maryland, Jun. 2001.
- [19] S. Haider, "On Computing Marginal Probability Intervals in Probabilistic Inference Networks," in Proceedings of IEEE-SMC Conference, October, 2003.
- [20] S. Haider, "A Hybrid Approach for Learning Parameters of Probabilistic Networks from Incomplete Databases," Design and Application of Hybrid Intelligent Systems, Proceedings of the Third International Conference on Hybrid Intelligent Systems (HIS'03), IOS Press, Amsterdam, The Netherlands, ISBN 158603-394-8, 2003.
- [21] S. Haider and A. H. Levis, "An Approximation Technique for Belief Revision in Timed Influence Nets," in Proceedings of Command and Control Research and Technology Symposium, 2004.
- [22] S. Haider and A. K. Zaidi, "Transforming Timed Influence Nets into Time Sliced Bayesian Networks," in Proceedings of Command and Control Research and Technology Symposium, 2004.
- [23] S. Haider and A. H. Levis, "Dynamic Influence Nets: An Extension of Timed Influence Nets for Modeling Dynamic Uncertain Situations," in Proc. 10th International Command and Control Research and Technology Symposium, Washington DC, June 2005.
- [24] S. Haider, A. K. Zaidi, and A. H. Levis, "A Heuristic Approach for Best Set of Actions Determination in Influence Nets," in Proc. IEEE International Conference on Information Reuse and Integration, Las Vegas, Nov. 2004.
- [25] S. Haider and A. H. Levis, "Effective Courses of Action Determination to Achieve Desired Effects" in IEEE Trans. on Systems, Man and Cybernetics, Part A: Systems and Humans, Vol. 37, No. 6, November 2007.
- [26] S. Haider and A. K. Zaidi, "On Temporal Analysis of Timed Influence Nets using Point Graphs," in Proc. of the 18th International FLAIRS Conference, FL, 2005.
- [27] A. K. Zaidi, L W. Wagenhals, and S. Haider, "Assessment of Effects Based Operations Using Temporal Logic," in Proc. of the 10th International Command and Control Research and Technology Symposium, 2005.
- [28] R. Howard and J. E. Matheson, "Influence diagrams," in Readings on the Principles and Applications of Decision Analysis, vol. 2, R. Howard and J. Matheson, Eds. Menlo Park, CA: Strategic Decisions Group, 1981, pp. 721–762.
- [29] R. A. Howard and J. E. Matheson, "Influence diagram retrospective," Decis. Anal., vol. 2, no. 3, pp. 144–147, Sep. 2005.

- [30] R. A. Howard, J. E. Matheson, M. W. Merkhofer, A. C. Miller, and D. W. North, "Comment on influence diagram retrospective," *Decis. Anal.*, vol. 3, no. 2, pp. 117–119, Jun. 2006.
- [31] G. F. Cooper, "The Computation Complexity of Probabilistic Inference Using Bayesian Belief Networks," *Artificial Intelligence*, 42(2-3): 393-405, March, 1990.
- [31] R. D. Shachter, "Evaluating influence diagrams," *Oper. Res.*, vol. 34, no. 6, pp. 871–882, Nov./Dec. 1986.
- [32] F.R. Kschischang and B.J. Frey, "Iterative decoding of compound codes by probability propagation in graphical models," *IEEE Journal on Selected Areas in Communication*, 16:219-230, 1998.
- [33] R. J. McEliece, McKay D.J.C., and Cheng J.F., "Turbo decoding as an instance of Pearl's belief propagation algorithm," *IEEE Journal on Selected Areas in Communication*, 16(2):140--152, February 1998.
- [34] K.P. Murphy, Y. Weiss, and M. I. Jordan, "Loopy belief propagation for approximate inference: an empirical study," in *Proceedings of Uncertainty in AI*, 1999.
- [35] L. W. Wagenhals and A. H. Levis, "Course of Action Analysis in a Cultural Landscape Using Influence Nets," in *Proc. IEEE Symp. on Computational Intelligence for Security and Defense Applications*, Honolulu, HI, April 2007
- [36] L. Wagenhals, "Course of action development and evaluation using discrete event system models of influence nets" PhD Dissertation, George Mason University, Fairfax, VA, 2000.
- [37] SIAM: Influence Net modeler (SIAC), <http://www.inet.saic.com/inet-public/siam.htm>.
- [38] Pythia: Timed Influence Net Modeler, <http://sysarch.gmu.edu/main/software/>. SAL-GMU.
- [39] D. L. McGuinness and F. van Harmelen, "OWL web ontology language overview," *World Wide Web Consortium*, 2004.
- [40] "SPARQL query language for RDF," <http://www.w3.org/TR/rdf-sparql-query/>.
- [41] S. Evren, P. Bijan, G. Bernardo Cuenca, K. Aditya and K. Yarden, "Pellet: A practical OWL-DL reasoner," *Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 5, pp. 51-53, 2007.
- [42] W. Keller., "Anatomy of a terrorist attack: an in-depth investigation into the 1998 bombings of the U.S. embassies in Kenya and Tanzania," Working Paper, The Matthew B. Ridgway Center for International Security Studies at the University of Pittsburgh, 2005.
- [43] Marschak, J. & Radner, R. (1972). *Economic Theory of Teams*, New London, CT: Yale Univ. Press.
- [44] Drenick, R. F. (1986). *A Mathematical Organization Theory*, New York: North-Holland.
- [45] Simon, H. A. (1982). *Models of Bounded Rationality*. Cambridge, MA: MIT Press.

- [46] Levis, A. H. (1988). Quantitative models of organizational information structures. In A. P. Sage, (Ed.) *Concise Encyclopedia of Information Processing in Systems and Organizations*, Oxford: Pergamon Ltd.
- [47] Jensen, K. (1990). *Coloured Petri Nets, Vols. I, II, and III*, Berlin: Springer-Verlag.
- [48] Levis, A. H. (1995). "Human Interaction with Decision Aids: A Mathematical Approach. In W. B. Rouse (Ed.) *Human/Technology Interaction in Complex Systems*, 7. Greenwich, CT: JAI Press.
- [49] Levis, A. H. (1993). A Colored Petri Net Model of Command and Control Nodes. In Carl R. Jones (Ed.), *Toward a Science of Command Control and Communications*, Washington, DC: AIAA Press.
- [50] Hofstede, G. (2001). *Culture's Consequences: Comparing Values, Behaviors, Institutions, and Organizations Across Nations*, 2nd Edition, Thousand Oaks, CA: Sage Publications.
- [51] Perdu, D. M., and Levis, A. H. (1998). Adaptation as a morphing process: A methodology for the design and evaluation of adaptive organization structures. *Computational and Mathematical Organization Theory* 4(1), 5-41
- [52] Remy, P. and Levis, A. H., " On the Generation Of Organizational Architectures Using Petri Nets," in *Advances in Petri Nets 1988, Lecture Notes in Computer Science*, G.Rozenberg Ed. Springer-Verlag, Berlin, FRG 1988.
- [53] Mooij, M. (1998). *Global Marketing and Advertising: Understanding Cultural Paradoxes*. Thousand Oaks, CA: Sage Publications.
- [54] Klein, H. A., Pongonis, A., & Klein, G. (2000). Cultural barriers to multinational C2 decision making. *Proc. 2000 Command and Control Research Symposium*, Monterey, CA.
- [55] Handley, H. A. H. & Levis, A. H. (2001). Incorporating heterogeneity in command center interactions. *Information Knowledge Systems Management* 2(4).
- [56] Zaidi, A. K., & Levis, A. H. (1995). Algorithmic design of distributed intelligence system architectures. In Gupta, M. M., & Sinha, N. K.(Eds.), *Intelligent Control Systems: Theory and Practice*, New York,
- [57] Olmez, A. E., *Exploring the Culture-Performance Link in Heterogeneous Coalition Organizations*, Ph.D. Thesis, George Mason University, Fairfax, VA, 2006.
- [58] Kathleen M. Carley, "On the Evolution of Social and Organizational Networks." In Steven B. Andrews and David Knoke (Eds.) Vol. 16 special issue of *Research in the Sociology of Organizations*. on "Networks In and Around Organizations." Greenwich, CN: JAI Press, Inc. Stamford, CT, pp. 3-30, 1999.
- [59] A. H. Levis, "Executable Models of Decision Making Organizations," in *Organizational Simulation*, William B. Rouse and Ken Boff, Eds., Wiley, NY, 2005
- [60] F. Mansoor, A. K. Zaidi, L. Wagenhals and A. H. Levis, "Meta-modeling the Cultural Behavior Using Timed Influence Nets," *Second International Workshop on Social Computing, Behavior Modeling and Prediction*, Phoenix, AZ, April, 2009

- [61] H. Neema, G. Hemingway, J. Sztipanovits, G. Karsai: The C2 Wind Tunnel: Rapid Synthesis of HLA-Based Heterogeneous Simulations via a Model-Based Integration Approach, Simulation (in press)
- [62] G. Karsai, M. Maroti, A. Lédeczi, J. Gray and J. Sztipanovits, "Composition and Cloning in Modeling and Meta-Modeling," IEEE Transactions on Control System Technology, Vol. 12, No. 2, pp. 263-278, March 2004
- [63] World Wide Web Consortium: <http://www.w3.org/>
- [64] Unified Modeling Language – UML: <http://www.uml.org/>
- [65] Resource Description Framework: <http://www.w3.org/RDF/>
- [66] Web Ontology Language - OWL: <http://www.w3.org/TR/owl-features/>
- [67] Protégé: <http://protege.stanford.edu/>
- [68] Wang, C., K. He, Y. He, W. Qian, Mappings from OWL-s to UML for Semantic Web Services, in Research and Practical Issues of Enterprise Information Systems, IFIP Intl. Federation for Information Processing, Vol. 205, 2006.
- [69] Meta-Object Facility: <http://www.omg.org/mof/>
- [70] Meta-Object Facility Self-Reference: http://en.wikipedia.org/wiki/Meta-Object_Facility
- [71] Kermeta: <http://www.kermeta.org/>
- [72] Moon, I., Carley, K. M. and Alexander H. Levis, Vulnerability Assessment on Adversarial Organization: Unifying Command and Control Structure Analysis and Social Network Analysis, SIAM International Conference on Data Mining Workshop on Link Analysis, Counterterrorism and Security, Atlanta, Georgia, Apr. 26, 2008
- [73] A. H. Levis, S. K. Kansal, A. E. Olmez, A. M. AbuSharekh, "Computational models of Multi-national Organizations" in Social Computing, Behavioral Modeling and Prediction, H. Liu, J. L. Salerno and M. J. Young, Eds., Springer, 2008
- [74] A. K. Zaidi and A. H. Levis, "TEMPER: A Temporal Programmer for Time-Sensitive Control of Discrete Event Systems," IEEE Trans. on Systems, Man, and Cybernetics, vol. 31, No. 6, Nov. 2001, pp. 485-96
- [75] A. K. Zaidi and A. H. Levis, "Validation and Verification of Decision Making Rules," Automatica, Vol. 33, No. 2, Feb. 1997, pp. 155-169.
- [76] Kathleen M. Carley, Jeffrey Reminga, Jon Storrick, and M. DeReno, "ORA User's Guide 2009", \ Carnegie Mellon University, School of Computer Science, Institute for Software Research, Technical Report CMU-ISR-09-115.
- [77] Kathleen M. Carley, "Destabilizing Terrorist Networks. Proc. of the 8th International Command and Control Research and Technology Symposium,, Washington DC 2003.
- [78] <http://www.issinc.com/solutions/webtas-overview.html>
- [79] "IHMC CMap Tools", <http://cmap.ihmc.us/download/index.php>

- [80] M. Emerson and J. Sztipanovits, "Techniques for Metamodel Composition", OOPSLA – 6th Workshop on Domain Specific Modeling, Portland, Oregon, October 22 2006, pp. 123-139
- [81] J. Ernst, What is metamodeling, 2002. (draft, online
<http://infogrid.org/wiki/Reference/WhatIsMetaModeling>)
- [82] J. P. van Gigch, System Design Modeling and Metamodeling, Plenum Press, New York, 1991.
- [83] C. Gonzalez-Perez and B. Henderson-Sellers, Metamodelling for Software Engineering. Chichester (UK), Wiley, 2008.
- [84] M. A. Jeusfeld, M. Jarke, and J. Mylopoulos, Metamodeling for Method Engineering. Cambridge (USA), The MIT Press, 2009.
- [85] F. Mansoor, A. K. Zaidi, L. Wagenhals and A. H. Levis, Meta-modeling the Cultural Behavior Using Timed Influence Nets, Second International Workshop on Social Computing, Behavior Modeling and Prediction, Phoenix, AZ, April, 2009.
- [86] R. S. Michalski and G. Tecuci, Machine learning: a multistrategy approach, Morgan Kaufman Publishers, Inc., San Francisco, 1994.
- [87] N. L. Cassimatis, A Cognitive architecture for integrating multiple representation and inference schemes, PhD Dissertation, Massachusetts Institute of Technology, February, 2002.
- [88] N. L. Cassimatis, J. G. Trafton, M. D. Bugajska, and A. C. Schultz, Integrating cognition, perception and action through mental simulation in robots, Robotics and Autonomous Systems 49 (2004) 13-23, Elsevier.
- [89] E. DeGregorio, R. Jenssen, L. Wagenhals, and R. Messier, Integrating Effects Based and Attrition Based Modeling, Proc. 2004 International Command and Control Research and Technology Symposium, June 15-17, 2004, San Diego, CA.
- [90] Paulo C. G. Costa and K. B. Laskey, Multi-Entity Bayesian Networks without Multi-Tears. USA 2005. (draft, online
http://ite.gmu.edu/~klaskey/papers/Costa_Laskey_MEBN_wo_Tears.pdf)
- [91] Paulo C. G. Costa and K. B. Laskey, Takikawa, Masami; Pool, Michael; Fung, Francis; and Wright, Edward J., MEBN logic: A Key Enabler for Network Centric Warfare, Proc. 10th International Command and Control Research and Technology Symposium (10th ICCRTS), June 13-16, 2005, McLean, Virginia, USA: CCRP publications.
- [92] Rommel N. Carvalho, Marcelo Ladeira, Laécio L. Santos, Matsumoto, Shou; and Costa, Paulo C. G. A GUI Tool for Plausible Reasoning in the Semantic Web Using MEBN. *Innovative Applications in Data Mining*. Studies in Computational Science, Vol. 169/2009, pp. 17-45. Springer-Verlag: Berlin/Heidelberg, Germany. ISBN: 978-3-540-88044-8. [o](#)
- [93] K. D. Jarman, A. J. Brothers, P. D. Whitney, J. Young and D. A. Niesen, Integrating System Dynamics and Bayesian Networks with Application to Counter-IED Scenarios, Proc. 10th International Probabilistic Safety Assessment & Management Conference, June 2010, Seattle.

- [94] P. D. Whitney and S. J. Walsh, Calibrating Bayesian Network Representations of Social-Behavioral Models, Proc. 2010 International Conference on Social Computing, Behavioral Modeling, and Prediction, Washington DC, March 30-31.
- [95] S. J. Walsh, A. C. Dalton, P. D. Whitney, and A. M. White, Parameterizing Bayesian network Representations of Social-Behavioral Models by Expert Elicitation, Proc. IEEE Intelligence and Security Informatics Conference, Vancouver, BC, May 23-26, 2010.
- [96] S. K. Kansal, A. M. AbuSharekh, and A. H. Levis, Computationally Derived Models of Adversary Organizations, Proc. IEEE Symp. on Computational Intelligence for Security and Defense Applications, Honolulu, HI, April 2007.
- [97] S. Haider, A. K. Zaidi, A. H. Levis, On Temporal Analysis of Timed Influence Nets using Point Graphs, Proc. of 18th International FLAIRS Conference, Clearwater Beach, Florida, May 2005.
- [98] I. Moon, K. M. Carley and A. H. Levis, Vulnerability Assessment on Adversarial Organization: Unifying Command and Control Structure Analysis and Social Network Analysis, SIAM International Conference on Data Mining Workshop on Link Analysis, Counterterrorism and Security, Atlanta, Georgia, Apr. 26, 2008.
- [99] F. Mansoor, A. K. Zaidi, L. W. Wagenhals and A. H. Levis, Meta-modeling the Cultural Behavior Using Timed Influence Net, Social Computing and Behavioral Modeling, Editors: Huan Liu, John J. Salerno and Michael J. Young, Springer, 2009.
- [100] G. Kappel, E. Kapsammer, H. Kargl, G. Kramler, T. Reiter, W. Retschitzegger, W. Schwinger and M. Wimmer, Lifting Metamodels to Ontologies: A Step to the Semantic Integration of Modeling Languages”, Model Driven Engineering Languages and Systems, Springer Berlin-Heidelberg, 2006, pp. 528-542.
- [101] M. Saeki and H. Kaiya, On Relationships among Models, Meta Models and Ontologies, Procc 6th OOPSLA Workshop on Domain-Specific Modeling, Portland, OR, 2006.
- [102] J. D. Novak & A. J. Cañas, The Theory Underlying Concept Maps and How to Construct and Use Them, Technical Report IHMC CmapTools 2006-01 Rev 01-2008.
- [103] F. Baader, D. McGuinness and D. Nardi, Description Logic Handbook: Theory and Applications, Cambridge University Press, 2003.
- [104] S. Bechhofer, OWL Reasoning Examples, University of Manchester.
<http://owl.man.ac.uk/2003/why/latest/>
- [105] Diesner, J. and K. M. Carley (2004). AutoMap 1.2 : extract, analyze, represent, and compare mental models from texts. [Technical report] / Carnegie Mellon University School of Computer Science Institute for Software Research International CMU-ISRI-04-100. Pittsburgh, Pa.,
- [106] Tsvetovat, M., J. Reminga, et al. (2003). DyNetML: Interchange Format for Rich Social Network Data. NAACSOS Conference, Pittsburgh, PA.
- [107] Tsvetovat, M., J. Reminga, et al. (2004). DyNetML: Interchange Format for Rich Social Network Data. [Technical report] / Carnegie Mellon University School of Computer Science Institute for Software Research International CMU-ISRI-04-105. Pittsburgh, Pa.,

- [108] Carley, Kathleen M. 1986. An approach for relating social structures to cognitive structure. *Journal of Mathematical Sociology*. 12: 137-189.
- [109] Carley, K. M. (1991) A Theory of Group Stability, *American Sociological Review*, Vol. 56, No. 3, pp. 331-354
- [110] Remy, P. A., Jin, V. Y., & Levis, A. H. (1988). On the Design of Distributed Organization Structures. *Automatica* , 24 (1), 1-18. 2000 Command and Control Research and Technology Symposium, 2000.
- [111] Wagenhals, L. W. (2000). Course of Action Development and Evaluation Using Discrete Event System Models of Influence Nets, PhD Dissertation, George Mason University, Fairfax, Va. 2000.
- [112] Wagenhals, L. W., I. Shin, and A. H. Levis (1998). Creating Executable Models of Influence Nets with Coloured Petri Nets. *Int. J. STTT*, Springer-Verlag. Vol. 1998, No. 2, 1998.
- [113] Cioffi-Revilla, C. 1985. Political reliability theory and war in the international system. *American Journal of Political Science* 29(1): 657–678.
- [114] Cioffi-Revilla, C. 1998. Politics and uncertainty. Theory, models and applications. Cambridge University Press.
- [115] Fearon, J. D. 2004. Why do some civil wars last so much longer than others? *Journal of Peace Research* 41(3): 275–301.
- [116] Bennett, D. S., and A. Stam. 2006. Predicting the length of the 2003 U.S. Iraq War. *Foreign Policy Analysis* 2(2): 101–116.
- [117] Bennett, D. S., and A. Stam. 1996. The duration of interstate wars, 1816–1985. *American Political Science Review* 43: 239–257.
- [118] Richardson, L. F. 1945. The distribution of wars in time. *Journal of the Royal Statistical Society (Series A)* 107(34): 242–250.
- [119] Richardson, L. F. 1960. Arms and insecurity: A mathematical study of the causes and origins of war. Boxwood Press.
- [120] Rapoport, A. 1957. Lewis F. Richardson’s mathematical theory of war. *Journal of Conflict Resolution* 1(3): 249–300.
- [121] Wilkinson, D. 1980. Deadly quarrels: Lewis F. Richardson and the statistical study of war. University of California Press.
- [122] Cioffi-Revilla, C., and M. I. Midlarsky. 2004. Highest-magnitude warfare: Power laws, scaling, and fractals in the most lethal international and civil wars. In *The scourge of war: New extensions on an old problem*, ed. P. F. Diehl, 3–27. University of Michigan Press.
- [123] W. Feller, *An Introduction to Probability Theory and its Applications*, Wiley, New York, 1968
- [124] O’Grady, E. 2006a. Diyala event data since March 12, 2003. Technical report, Lawrence Livermore Laboratory.
- [125] O’Grady, E. 2006b. Updated Diyala event data since March 12, 2003. Technical report, LawrenceLivermore Laboratory.

- [126] Boccara, N. 2004. Modeling complex systems. Springer.
- [127] Merritt, R. L., R. G. Muncaster, and D. A. Zinnes. 1993. International event-data developments: DDIR phase II. University of Michigan Press.
- [128] Cox, D. R. 1962. Renewal theory. Methuen.
- [129] Bartholomew, D. J. 1973. Stochastic models for social processes. Wiley.
- [130] Greene, W. H. 2000. Econometric analysis (4th ed.). Prentice Hall.
- [131] Kleiber, C., and S. Kotz. 2003. Statistical size distributions in economics and actuarial sciences.
- [132] Lewin, K. 1952. Field theory in social science: Selected theoretical papers by Kurt Lewin. Tavistock.
- [133] Vansteenkiste, M., and K. Sheldon. 2006. "There is nothing more practical than a good theory: Integrating motivational interviewing and self-determination theory. *British Journal of Clinical Psychology* 45: 63–82.
- [134] Kline, M., *Mathematics and the Search for Knowledge*, Oxford University Press, New York, 1985
- [135] Cioffi-Revilla, C., and S. O'Brien. 2007. Computational analysis in U.S. foreign and defense policy. In *Proceedings of the First International Conference on Computational Cultural Dynamics*. University of Maryland, College Park, MD, 27–28 August 2007.
- [136] Allison, G. 2004. Nuclear terrorism: The ultimate preventable catastrophe. Henry Holt and Company.
- [137] Schrodtt, P. A. 1989. Short term prediction of international events using a Holland classifier. *Mathematical and Computer Modelling* 12: 589–600.
- [138] King, G., and W. Lowe. 2003. An automated information extraction tool for international conflict data with performance as good as human coders: A rare events evaluation design. *International Organization* 57: 617–642.
- [139] O'Brien, S. P. 2002. Anticipating the good, the bad, and the ugly: An early warning approach to conflict and instability analysis. *Journal of Conflict Resolution* 46(6): 808–828.
- [140] Cioffi-Revilla, C. A. 1990. The scientific measurement of international conflict: Handbook of datasets on crises and wars, 1495–1988. A. D. L. Rienner.
- [141] Tsvetovat, M., and K. Carley. 2005. Structural knowledge and success of antiterrorist activity: The downside of structural equivalence. *Journal of Social Structure* 6(2): online .
- [142] Giuliano, C., Lavelli, A., & Romano, L. (2007). Relation extraction and the influence of automatic named-entity recognition. *ACM Transactions on Speech and Language Processing* 5(1), 2:1-2:26.
- [143] Roddick, J. F. & Lees, B. G. (2009). Spatio-temporal data mining paradigms and methodologies. In H.J. Miller and J. Han (Eds.) *Geographic data mining and knowledge discovery*, 2nd ed., (pp. 27-44). New York: CRC Press.
- [144] Battenfield, B., Gahegan, M., Miller, H. Yuan, M. (2001?). Geospatial data mining and knowledge discovery. Retrieved April 7, 2010 from

http://www.ucgis.org/priorities/research/research_white/2000%20Papers/emerging/gkd.pdf

- [145] Kozareva, Z. (2006). Bootstrapping named entity recognition with automatically generated gazetteer lists. Proceedings of the Eleventh Conference of the European Chapter of the Association for Computational Linguistics. Association for Computational Linguistics, Morristown, NJ, USA, 15-21.
- [146] Popescu, A. & Grefenstette, G. (2010). Spatiotemporal mapping of Wikipedia concepts. JCDL '10, June 21-25, 2010, Gold Coast, Queensland, Australia, 129-138.
- [147] Carley, K. M. (1993). Coding Choices for Textual Analysis: A Comparison of Content Analysis and Map Analysis. In Marsden P. (Ed), *Sociological Methodology*, 23 (pp. 75-126). Oxford: Blackwell.
- [148] Carley, K. M. (1997). Network Text Analysis: The Network Position of Concepts. In C. Roberts (Ed.), *Text Analysis for the Social Sciences: Methods for Drawing Statistical Inferences from Texts and Transcripts*. Ch. 4 (pp. 79-100). Hillsdale, NJ: Lawrence Erlbaum Associates.
- [149] Zhang, H., Yen, J., Giles, C. Lee, Mombaster, B., Spiliopoulou, M. Srivastava, J., Nasraoui, E., & McCallum, A. (2007). WebKDD/NAKDD 2007 - Web Mining and Social Network Analysis Post-Workshop Report. *SIGKDD Explorations* 9(2), 87-92.
- [150] Srivastava, J. (2008). Data mining for social network analysis. *IEEE International Conference on Intelligence and Security Informatics*, 2008, Taipei, 17-20 June 2008, pp. xxxiii-xxxiv.
- [151] Carley, Kathleen M., Dave Columbus, Michael Bigrigg and Frank Kunkel, 2010 "Auto-Map User's Guide 2010," Carnegie Mellon University, School of Computer Science, Institute for Software Research, Technical Report, CMU-ISR-10-121.
- [152] Carley, Kathleen M. 2002, "Smart Agents and Organizations of the Future," *The Handbook of New Media*. Edited by Leah Lievrouw and Sonia Livingstone, Ch. 12, pp. 206-220, Thousand Oaks, CA, Sage.
- [153] Lafferty, J., McCullum, A., & Perieira, F. (2001). Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In proceedings of the Eighteenth International Conference on Machine Learning (ICML-01).
- [154] Sha, F., & Pereira, F. (2003). Shallow parsing with conditional random fields. In Proceedings of HLT-NAACL, 213–220. Association for Computational Linguistics.
- [155] Carley, K., Frantz, T., & Diesner, J. (2006). Social and Knowledge Networks from Large Scale Databases. 56th Annual Conference of the International Communication Association (ICA). Dresden, Germany, June 19-23, 2006.
- [156] Diesner, J., & Carley, K.M. (2005). Revealing and Comparing the Organizational Structure of Covert Networks with Network Text Analysis. XXV Sunbelt Social Network Conference, Redondo Beach, CA, February 16-20, 2005.
- [157] Weil, S.A., Carley, K.M., Diesner, J., Freeman, J., and Cooke, N.J. (2005). Measuring Situational Awareness through Analysis of Communications: A Preliminary Exercise.

Submitted to the Command and Control Research and Technology Symposium 2006, San Diego, CA.

- [158] Pazienza, M. T., Pennacchiotti, M., & Zanzotto, F. M. (2006). Discovering verb relations in corpora: Distributional versus non-distributional approaches. In A. Ali & R. Dapoigny (Eds.), IEA/AIE 2006, LNAI 4031, 1042-1052.
- [159] Xu, X., Mete, M., & Yuruk, N. (2005). Mining concept associations for knowledge discovery in large textual databases. SAC '05, March 13-17, 2005, Santa Fe, New Mexico, U.S.A., 549-550.
- [160] Danowski, J. A. (2009). Automatic mapping of social networks of actors from text corpora: Time series analysis. 2009 Advances in Social Network Analysis and Mining, 137-142.
- [161] Child, J. "Predicting and Understanding Organization Structure." Administrative Science Quarterly 18.2 (1973): 168-185.
- [162] Meyer, J. W. and B. Rowan. "Institutionalized Organizations: Formal Structure as Myth and Ceremony." The American Journal of Sociology 83.2 (1977): 340-363.
- [163] Shetty, J. and J. Adibi. Discovering Important Nodes through Graph Entropy: The Case of Enron Email Database. Conference on Knowledge Discovery in Data archive, Proceedings of the 3rd International Workshop on Link Discovery. (2005): 74 – 81.
- [164] Borgatti, S. P., K. M. Carley, and D. Krackhardt. "On the Robustness of Centrality Measures under Conditions of Imperfect Data." Social Networks 28. 2 (2006): 124-136.
- [165] Rabasa, A., et al. Beyond al-Qaeda Part 1. City of publication: Rand Corp., 2006. 27-29.
- [166] Morel, B. and R. Ramanujam. "Through the Looking Glass of Complexity: The Dynamics of Organizations as Adaptive and Evolving Systems." Organization Science 10.3, Special Issue: Application of Complexity Theory to Organization Science (1999).
- [167] Grant, R. M. "Toward a Knowledge-Based Theory of the Firm." Strategic Management Journal 17 Special Issue: Knowledge and the Firm (1996): 109-122.
- [168] Krebs, V. E. "Mapping Networks of Terrorist Cells." CONNECTIONS 24.3 (2002): 43-52
- [169] Mayntz, R. "Organizational Forms of Terrorism: Hierarchy, Network, or a Type Sui Generis?" Discussion Paper. Max-Planck-Institut für Gesellschaftsforschung Köln, MPIfG, 2004.
- [170] Carley, K. M., "Destabilization of Covert Networks." Computational & Mathematical Organization Theory 12.1 (2006): 51-66.
- [171] Sageman, M. Understanding Terror Networks. Philadelphia: University of Pennsylvania Press, 2004.
- [172] Thompson
- [173] Nonaka, I. "A New Organizational Structure." Knowledge in Organizations. Ed. L. Prusak. Boston: Butterworth-Heinemann, 1997. 110-119.

- [174] Schilling, M. A. and C. C. Phelps. "Interfirm Collaboration Networks: The Impact of Large-Scale Network Structure on Firm Innovation." *Management Science* 53.7 (2007): 1113-1126.
- [175] Elliott, E. and L. D. Kiel. "A Complex Systems Approach for Developing Public Policy toward Terrorism: An Agent-Based Approach." *Chaos, Solitons & Fractals* 20. 1 (April 2004): 63-68.
- [176] Levis, A. H. "Executable Models of Decision Making Organizations," in *Organizational Simulation*, William B. Rouse and Ken Boff, Eds., Wiley, NY, 2005
- [177] Smircich, L. "Concepts of Culture and Organizational Analysis." *Administrative Science Quarterly* 28.3 (1983): 339-358.
- [178] Floyd, S. W. and B. Wooldridge. "Knowledge Creation and Social Networks in Corporate Entrepreneurship: The Renewal of Organizational Capability." *Entrepreneurship: Theory and Practice* 23 (1999).
- [179] Gabbay, S. and R. Leenders. "Social Capital of Organizations." *Research in the Sociology of Organizations*. Vol. 18. Amsterdam: JAI Press, 2001.
- [180] Reagans, R. and E. W. Zuckerman. "Networks, Diversity, and Productivity: The Social Capital of Corporate R&D Teams." *Organization Science* 12.4 (2001): 502-517.
- [181] Carley, K. M., Summary of Key Network Measures for Characterizing Organizational Architectures. Unpublished Document: CMU, 2002.
- [182] Carley, Kathleen M. 1997, "Network Text Analysis: The Network Position of Concepts," in C. Roberts (Ed.), *Text Analysis for the Social Sciences: Methods for Drawing Statistical Inferences from Texts and Transcripts*. Hillsdale, NJ: Lawrence Erlbaum Associates, Chapter 4, pp. 79-102.
- [183] Freeman, L. C. "Centrality in Social Networks: Conceptual Clarification." *Social Networks* 1.3(1979): 215-239.
- [184] Bonacich, P. "Techniques for Analyzing Overlapping Memberships." *Sociological Methodology*. Ed. Herbert L. Costner. San Francisco: Jossey-Bass, 1972. 176-85.
- [185] Reminga, J. and K. M. Carley. "ORA: Organization Risk Analyzer, Tech Report, CMU-ISRI-04-106." CASOS. Carnegie Mellon University. Pittsburgh, PA. 2004. <<http://www.casos.cs.cmu.edu/projects/ora/index.html>>.
- [186] Borgatti, S. P. and M. G. Everett. "Notions of Position in Social Network Analysis." *Sociological Methodology* 22 (1992): 1-35.
- [187] Krackhardt, D., and K. M. Carley. A PCANS Model of Structure in Organization, In *Proceedings of the 1998 International Symposium on Command and Control Research and Technology*. 1998: 113-119.
- [188] Carley, K. 1990. Group stability: A socio-cognitive approach. In Lawler E., Markovsky B., Ridgeway C., and Walker H. (Eds.) *Advances in group processes: Theory & research*. Vol. VII. (pp. 1-44). Greenwich, CN:JAI Press.
- [189] Carley, K.M. 1995. Communication technologies and their effect on cultural homogeneity, consensus, and the diffusion of new ideas. *Sociological Perspectives*, 38(4):547-571.

- [190] Hirshman and Carley
- [191] Schreiber, C. & K.M. Carley, 2004. Going Beyond the Data: Empirical Validation Leading to Grounded Theory, Computational and Mathematical Organization Theory, 10, 155-164.
- [192] Giddens, A. (1984). The constitution of society: Outline of the theory of structuration. Berkeley, CA: University of California Press.
- [193] Salancik & Pfeffer, 1978
- [194] Manis , J. G. & Meltzer, B. N. 1978. Symbolic interaction: A reader in social psychology. Boston: Allyn & Bacon.
- [195] Stryker S., 1980. Symbolic Interactionism. Menlo Park, CA: Benjamin Cummings.
- [196] Friedkin, N., 1998, A Structural Theory of Social Influence, New York, NY, Cambridge University Press.
- [197] Festinger, L., 1957. A theory of cognitive dissonance. Evanston, IL: Row, Peterson.
- [198] Festinger, L., 1954. A Theory of Social Comparison Processes, Human Relations, 7: 117-140.
- [199] Wegner, D. M., 1986. Transactive memory: A contemporary analysis of the group mind, In B. Mullen & G. R. Goethals (Eds), Theories of group behavior, (pp. 185-205), New York: Springer-Verlag.
- [200] Hirshman, B.R., K.M. Carley & M.J. Kowalchuck, 2007a, "Specifying Agents in Construct," Carnegie Mellon University, School of Computer Science, Institute for Software Research, Technical Report, CMU-ISRI-07-107.
- [201] Hirshman, B.R., K.M. Carley & M.J. Kowalchuck, 2007b. "Loading Networks in Construct," Carnegie Mellon University, School of Computer Science, Institute for Software Research, Technical Report, CMU-ISRI-07-116.
- [202] Wilkins, J. S. 1998. What's in a meme? Journal of Memetics-Evolutionary Models of Information Transmission, 2 (1).
- [203] Carley, K. M. & D. Krackhardt. 1996. Cognitive inconsistencies and non-symmetric friendship, Social Networks, 18: 1-27.
- [204] Schreiber, C. & K.M. Carley, 2004. "Construct - A Multi-agent Network Model for the Co-evolution of Agents and Socio-cultural Environments." Carnegie Mellon University, School of Computer Science, Institute for Software Research International, Technical Report CMU-ISRI-04-109.
- [205] Carley, K. M. & V. Hill. 2001. Structural change and learning within organizations. In A. Lomi and E. R. Larsen (eds), Dynamics of Organizations: Computational Modeling and Organization Theories. Menlo Park, CA: MIT Press/AAAI.
- [206] Schreiber, C. & K.M. Carley, 2008, "Dynamic Network Leadership: Leading for Learning Adaptability," pps. 291 – 332 in Mary Uhl-Bien and Russ Marion (eds.): Complexity Leadership Part I: Conceptual Foundations, Charlotte, NC: Information Age Publishing, Inc.

- [207] Kaufer, D.S. & K.M. Carley. 1993. *Communication at a Distance: The Effect of Print on Socio-Cultural Organization and Change*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- [208] Hirshman, B.R., M.K. Martin, M. Bigrigg & K.M. Carley, 2008, "The Impact of Educational Media by Socio-Demographic Attribute," Carnegie Mellon University, School of Computer Science, Institute for Software Research, Technical Report, CMU-ISR-08-118.
- [209] Hirshman, B.R., M.K. Martin, A. Birukou, M. Bigrigg & K. M. Carley, 2008b, "The Impact of Educational Media on Real and Stylized Cities," Carnegie Mellon University, School of Computer Science, Institute for Software Research, Technical Report, CMU-ISR-08-114.
- [210] Blau, P.M. . 1977. *Inequality and Heterogeneity*. New York: The Free Press of Macmillan Co.
- [211] Carley, K.M., J. Lee & D. Krackhardt, 2001, Destabilizing Networks, *Connections* 24(3):31-34.
- [212] McPherson, J. M. & L. Smith-Lovin. 1987. Homophily in Voluntary Organizations: Status Distance and the Composition of Face-to-Face Groups. *American Sociological Review* 52: 370-79.
- [213] Schreiber, C., 2006. "Human and Organizational Risk Modeling, Critical Personnel and Leadership in Network Organizations," Ph.D. Thesis, Carnegie Mellon University, School of Computer Science, Institute for Software Research, Computation, Organizations and Society, Pittsburgh, PA.
- [214] Valente, T.H., P.R. Poppe & A.P. Merritt, 1996. Mass-Media-Generated Interpersonal Communication as Sources of Information About Family Planning, *Journal of Health Communication: International Perspectives*, 1(3): 247-266.
- [215] Smith-Lovin, L., M. McPherson & J. Cook. 2001. Birds of a feather: Homophily in social networks. *Annual Review of Sociology* , 27: 415-44.
- [216] Hill, V. & K.M. Carley, 2008, "Win Friends and Influence People: Relationships as Conduits of Organizational Culture in Temporary Placement Agencies," *Journal of Management Inquiry*, Vol. 17 (4), 369-379.
- [217] Erdos, P. & A. Renyi, 1959. On random graphs i. *Publicationes Mathematicae Debrecen*, 6:290–297, 1959.
- [218] Frantz, T. & K. M. Carley, 2005, "A Formal Characterization of Cellular Networks," Carnegie Mellon University, School of Computer Science, Institute for Software Research International, Technical Report CMU-ISRI-05-109.
- [219] Watts, D.J. & S.H. Strogatz, 1998. Collective dynamics of 'small-world' networks. *Nature* 393 (6684): 409–10.
- [220] Newman, M., 2000. Models of the small world. *Journal of Statistical Physics*, Vol. 101, Nos. 3/4, 819-841.
- [221] Hirshman, B.R., M.K. Martin & K.M. Carley, 2008, "Modeling Information Access in Construct," Carnegie Mellon University, School of Computer Science, Institute for Software Research, Technical Report, CMU-ISR-08-115.

- [222] Epstein, J., Steinbruner, J. D., and Parker, M. T. (2001) Modeling Civil Violence: An Agent-Based Computational Approach. Washington, D.C., Center of Social and Economic Dynamics, Brookings Institute.
- [223] Bergkvist, M., Davidson, P., Persson, J. A., and Ramstedt, L. 2004. A Hybrid Micro-Simulator for Determining the Effects of Governmental Control Policies on Transport Chains. Joint Workshop on Multi-Agent and Multi-Agent Based Simulation, New York, NY, Springer
- [224] Louie, M. A., and Carley, K. M. (2006) The Role of Multi-Agent Models of Socio-Political Systems in Policy, working paper, CASOS, Carnegie Mello University
- [225] Carley, K. M. (2006) A Dynamic Network Approach to the Assessment of Terrorist Groups and The Impact of Alternative Courses of Action, In Visualising Network Information. Meeting Proceedings RTO-MP-IST-063. Neuilly-sur-Seine, France: RTO, 12 - 14. Available from: http://www.vistg.net/documents/IST063_PreProceedings.pdf
- [226] Davis, J. P., Bingham, C. B., and Eisenhardt, K. M. (2006) Developing Theory Through Simulation Methods, Academy of Management Review.
- [227] Cohen, M. D., March, J., and Olsen, J. P. (1972) A Garbage Can Model of Organizational Choice. Administrative Science Quarterly, 17(1): 1-25.
- [228] March, J. G. (1991) Exploration and Exploitation in Organizational Learning. Organization Science, 2(1): pp 71-87.
- [229] Freeman, L. C. (1979). Centrality in social networks: Conceptual clarification. Social Networks, 1, 215-239.
- [230] Bonacich, P. (1972) Factoring and weighting approaches to clique identification. Journal of Mathematical Sociology, 2 (1972), 113-120.
- [231] Carley, Kathleen M. (2002) Smart Agents and Organizations of the Future. The Handbook of New Media. Edited by Leah Lievrouw & Sonia Livingstone, Ch. 12 pp. 206-220, Thousand Oaks, CA, Sage.
- [232] Felter J. and Fishman, B. (2007) Al-Qai'da's Foreign Fighters in Iraq, Harmony Project Reports, US Military Academy at Westpoint, NY
- [171] Sageman, M (2004) Understanding terror networks, Philadelphia: University of Pennsylvania Press
- [233] Champagne, B., Cobb, W. B., Kaplan, S., and Morrison, D. (eds.) (2005) Anatomy of a Terrorist Attack: An In-Depth Investigation into the 1998 Bombings of the US Embassies in Kenya and Tanzania, Ridgway Center for International Security Studies, Issue 17, University of Pittsburgh, US

APPENDIX A

Proof of Lemma 3.2

Applying the Bayes Rule, we write

$$P(B | x_1^{n-1}) = \frac{P(x_1^{n-1} | B)P(B)}{P(x_1^{n-1})} \quad (\text{A.1})$$

Where, due to the Theorem of Total Probability, we have

$$P(x_1^{n-1} | B) = \sum_{x_n=0,1} P(x_1^n | B) \quad (\text{A.2})$$

From the Bayes Rule we also have,

$$P(x_1^n | B) = \frac{P(B | x_1^n)P(x_1^n)}{P(B)} \quad (\text{A.3})$$

Substituting (A.2) and (A.3) in (A.1), we obtain:

$$P(B | x_1^{n-1}) = \sum_{x_n=0,1} P(B | x_1^n)P(x_1^n)P^{-1}(x_1^{n-1}) = \sum_{x_n=0,1} P(B | x_1^n)P(x_n | x_1^{n-1}) \quad (\text{A.4})$$

Where in (A.4) we used the definition of conditional probability $P(x_n | x_1^{n-1}) = P(x_1^n)P^{-1}(x_1^{n-1})$. Using expression (4) in Section III and substituting in (A.4), we obtain.

$$\begin{aligned} & \{1 + h_{n-1}(x_1^{n-1})[1 - P(B)]P^{-1}(B)\}^{\text{sgn } h_{n-1}(x_1^{n-1})} \bullet \{1 + h_{n-1}(x_1^{n-1})\}^{1 - \text{sgn } h_{n-1}(x_1^{n-1})} = \\ & \sum_{x_n=0,1} P(x_n | x_1^{n-1}) \{1 + h_n(x_1^n)[1 - P(B)]P^{-1}(B)\}^{\text{sgn } h_n(x_1^n)} \bullet \{1 + h_n(x_1^n)\}^{1 - \text{sgn } h_n(x_1^n)} \stackrel{\Delta}{=} Q_n + 1 \end{aligned} \quad (\text{A.5})$$

Observing (A.5), we notice that if $(Q_n + 1) \in [0, 1]$, then $h_{n-1}(x_1^{n-1})$ must be necessarily negative, reducing the left part of the equality in (A.5) to $1 + h_{n-1}(x_1^{n-1})$. If, on the other hand, $(Q_n + 1) \in [1, P^{-1}(B)]$ then the left part of (A.5) must necessarily reduce to $1 + h_{n-1}(x_1^{n-1})[1 - P(B)]P^{-1}(B)$, with $h_{n-1}(x_1^{n-1})$ positive. The above observations clearly lead to the result in the lemma.

Proof of Lemma 3.3

Due to the Bayes Rule, we have:

$$P(B | x_1^n) = \frac{P(x_1^n | B)P(B)}{P(x_1^n)} \quad (\text{A.6})$$

Due to the independence assumption, we have:

$$\frac{P(x_1^n | B)}{P(x_1^n)} = \prod_{i=1}^n \frac{P(x_i | B)}{P(x_i)} \quad (\text{A.7})$$

where

$$\frac{P(x_i | B)}{P(x_i)} = \frac{P(B | x_i)}{P(B)} \quad (\text{A.8})$$

Substituting (A.8) in (A.7) and then (A.6), we obtain:

$$P(B | x_1^n) = P^{-(n-1)}(B) \prod_{i=1}^n P(B | x_i) \quad (\text{A.9})$$

Substituting expression (10) in (A.9) we obtain the expression in the lemma.

Proof of Lemma 3.4

Due to the Bayes Rule, we have:

$$\frac{P(B | x_1^n)}{P(B)} = \frac{P(x_1^n | B)}{P(x_1^n)} \quad (\text{A.10})$$

Then, due to (13), we obtain from (A.10),

$$\frac{P(B | x_1^n)}{P(B)} = \frac{P(x_1 | B)}{P(x_1)} \prod_{i=2}^n \frac{P(x_i | x_{i-1}, B)}{P(x_i | x_{i-1})} \quad (\text{A.11})$$

Applying the Bayes Chain Rule, we have:

$$P(x_i | x_{i-1}, B)P(x_{i-1} | B)P(B) = P(B | x_i, x_{i-1})P(x_i | x_{i-1})P(x_{i-1}) \quad (\text{A.12})$$

where,

$$P(x_{i-1} | B)P(B) = P(B | x_{i-1})P(x_{i-1}) \quad (\text{A.13})$$

Substituting (A.13) in (A.12) we then obtain:

$$\frac{P(x_i | x_{i-1}, B)}{P(x_i | x_{i-1})} = \frac{P(B | x_i, x_{i-1})}{P(B | x_{i-1})}; i \geq 2$$

$$\frac{P(x_1 | B)}{P(x_1)} = \frac{P(B | x_1)}{P(B)} \quad (\text{A.14})$$

Where, directly from the results in Lemma 1, we have:

$$P(B | x_i) = \left\{ 1 + h_1^{(i)}(x_i) [1 - P(B)] P^{-1}(B) \right\}^{\text{sgn } h_1^{(i)}(x_i)} \left\{ 1 + h_1^{(i)}(x_i) \right\}^{1 - \text{sgn } h_1^{(i)}(x_i)} \quad (\text{A.15})$$

where

$$h_1^{(i)}(x_i) = \begin{cases} Q_{i,i+1} - 1 & ; \text{ if } Q_{i,i+1} \in [0,1] \\ P(B)[1 - P(B)]^{-1}[Q_{i,i+1} - 1] & ; \text{ if } Q_{i,i+1} \in [1, P^{-1}(B)] \end{cases} \quad (\text{A.16})$$

$$Q_{i,i+1} = \sum_{x_{i+1}=0,1}^{\Delta} P(x_{i+1} | x_i) \left\{ 1 + h_2^{(i,i+1)}(x_i, x_{i+1}) [1 - P(B)] P^{-1}(B) \right\}^{\text{sgn} h_2^{(i,i+1)}(x_i, x_{i+1})} \bullet \left\{ 1 + h_2^{(i,i+1)}(x_i, x_{i+1}) \right\}^{1 - \text{sgn} h_2^{(i,i+1)}(x_i, x_{i+1})} \quad (\text{A.17})$$

Substitution of expression (A.14) and (A.17) in (A.11), in conjunction with expression (14), give the result in the lemma.

Proof of Lemma 3.5

Due to the Bayes Rule and the Theorem of Total Probability, we have:

$$P(B) = \sum_{x_1^n} P(B | x_1^n) P(x_1^n) \quad (\text{A.18})$$

Substituting in (A.18) the expression (27) for the conditional probability $P(B | x_1^n)$, we obtain:

$$\begin{aligned} & \sum_{x_1^n : \sum_{i=1}^n h_1(x_i) > 0} P(x_1^n) \left\{ 1 + P^{-1}(B) [1 - P(B)] \max_{1 \leq i \leq n} h_1(x_i) \right\} + \\ & \sum_{x_1^n : \sum_{i=1}^n h_1(x_i) = 0} P(x_1^n) + \sum_{x_1^n : \sum_{i=1}^n h_1(x_i) < 0} P(x_1^n) \left\{ 1 + \min_{1 \leq i \leq n} h_1(x_i) \right\} = 1 \end{aligned} \quad (\text{A.19})$$

which gives after simplification:

$$[1 - P(B)] \sum_{x_1^n : \sum_{i=1}^n h_1(x_i) > 0} P(x_1^n) \max_{1 \leq i \leq n} h_1(x_i) + P(B) \sum_{x_1^n : \sum_{i=1}^n h_1(x_i) < 0} P(x_1^n) \min_{1 \leq i \leq n} h_1(x_i) = 0 \quad (\text{A.20})$$

APPENDIX B

Pythia

Peter W. Pachowicz, Lee W. Wagenhals, John Pham, and Alexander H. Levis

Pythia, a timed influence net modeling and simulation tool for course of action development, evaluation and selection in the context of effects based planning, [2] and (3) *Temper*, a temporal planning and temporal reasoning tool using time interval logic. [3]

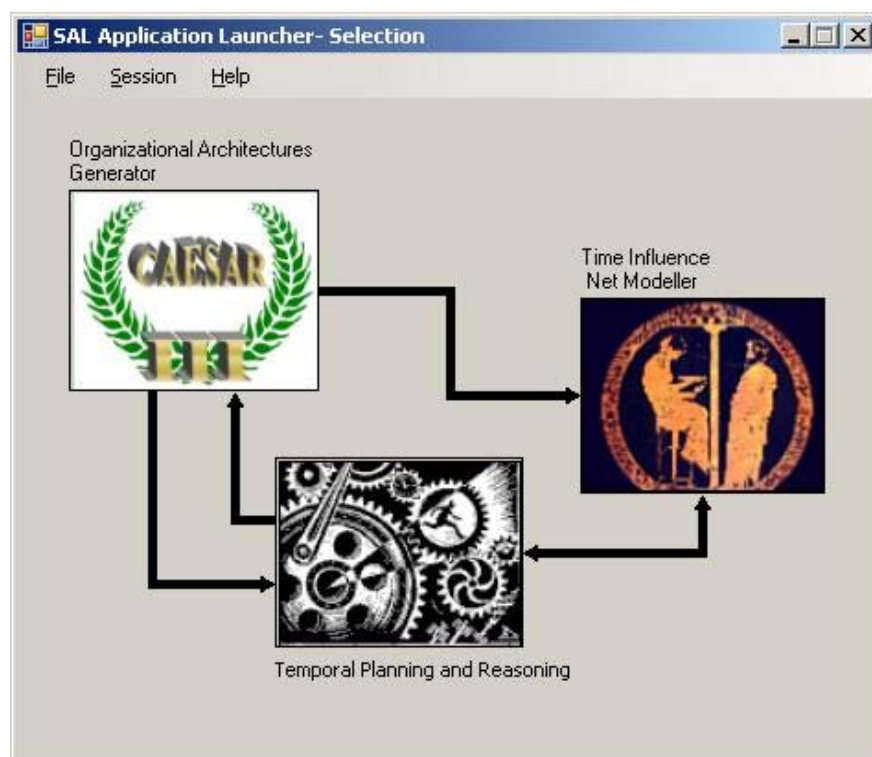


Figure 1: The Interface to the Analyst's Toolbox

The above-mentioned software has been re-engineered from being a stand-alone applications into the state-of-the-art client-server computing environment. At this time, only *Pythia 1.5* has been brought into the distributed infrastructure, with the remaining tools to follow shortly. Each of these tools has been converted to a server-centric application for multi-user and multi-process computing; the environment relies on the Citrix Presentation Server for integration, security, and maintenance.

Figure 2 shows the Presentation Services approach incorporated in the development of the analyst's suite. While each application is run on a server, the input/output services are controlled and displayed through a client machine that typically is a PC or laptop remotely located from the

server. An individual user must obtain an account and install a few system components on a client machine. Login and all interactions with analyst tools are performed using a client-GUI. A server performs all computations and generates displays presented virtually on a client machine. A user does not have to know the location of the computing resources and is not limited by the computational power of his/her client machine, meaning that operators in the field can use the tool with relatively small devices with limited computing power and memory. Moreover, a user can open several concurrent processes of the same (or different) application with different data sets. Only the power of the server-side computing infrastructure and the number of Citrix licenses limits the number of application processes that can be open concurrently. Our current efforts are also focused on providing a cross-application interfacing through model and data exchange mechanisms. This will allow us to apply complementary tools to the same modeling and simulation problem.




Figure 2: Presentation Services Approach

A Citrix Presentation Services architecture has been chosen for the development of the client-server computing infrastructure. In short, Citrix Presentation Services technology (i) provides virtual delivery of an application through a centralized and secure distributed architecture, (ii) arranges a virtual view of system resources (file system, registry, named objects, etc.), and (iii) supports a variety of hardware/software/OS technologies no matter the infrastructure diversity. The infrastructure on the server-side can involve one or more farms of application servers depending on their geographic distribution. Each server farm can be divided into zones depending on the complexity of the local network and the number of connected servers. Incoming tasks are automatically routed to application servers depending on their resources and current load. Connectivity with external databases is arranged. Citrix provides higher-level middleware (APIs, and SDKs) for a middleware customization in order to arrange total control over the client, network delivery, and server services. Such a centralized computing architecture, allows diverse users to develop and maintain their application models, perform simulations, and apply different analytical tools, while the application developer can independently continue to enhance the capabilities of these tools.

Pythia's Time Influence Net Modeling (TIN)

Pythia is using a Time Influence Net modeling extension to the Influence Net modeling paradigm. An Influence Net (IN) is a Directed Acyclic Graph where nodes in a graph represent random variables and the edges between nodes represent causal relationships. While mathematically IN are similar to Bayesian Nets (BN) [4], there are key differences. The most important is that IN uses CAST Logic [5,6] to enhance knowledge elicitation from subject matter experts in defining precise a-priori conditional probabilities used by BN. IN modeling is accomplished by creating a series of cause and effect relationships between desired (and undesired) effects and the set of actions that might impact their occurrence. The actionable events in IN are drawn as root nodes (nodes without incoming edges). Desired effects, or objectives the decision maker is interested in, are modeled as leaf nodes (nodes without outgoing edges). In some cases, internal nodes can be effects of interest, as well.

Figure 3 shows an example IN that has been created with the Pythia tool. The actionable nodes (nodes on the left) have been assigned marginal probability values – a probability indicating whether or not an action will be taken. The other nodes have assigned baseline probability in the CAST Logic – a probability indicating whether the random variable the node represents will be true on its own without the influence of any parent nodes in the model. The edges (links between nodes) have casual strength values (first two values) indicating the degree of influence that a parent node has on its child in the CAST Logic. The first strength value indicates the effect on the child node when the parent is ‘True.’ The second strength value indicates the effect on the child node when the parent is ‘False.’

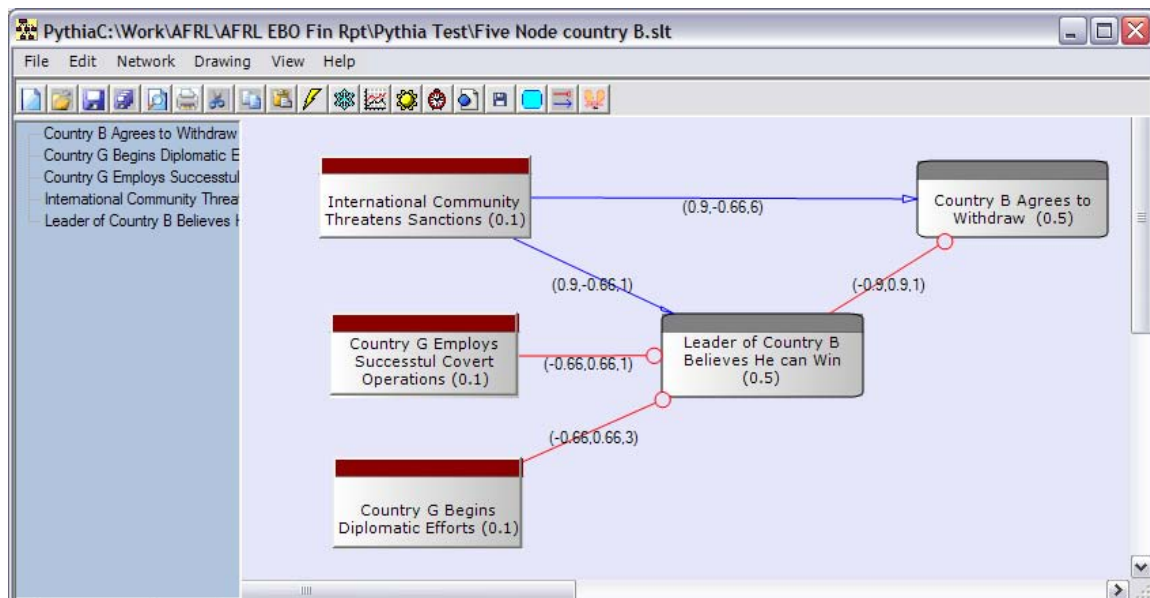


Figure 3: Example Time Influence Net

The TIN modeling extension allows the modeler to allocate time delays associated with nodes and edges, representing an impact of events (actions or effects) that takes some time to reach and

be processed by the affected events or conditions. For example, the third value assigned to edges (see Figure 3) represents time delays in time units. Consequently, Time Stamps are associated with each node (including the action nodes). Hence, a user can specify a Coarse of Action (COA), as a time sequence, on the action nodes, which are propagated through the network and trigger changes to the probability values of the effected nodes. The change in probability value of desired effect (leaf node) can be observed over time. Figure 4 illustrates the change in probability value for the effect node ‘Country B Agrees to Withdraw’ for the following COA on the action nodes:

COA = [‘Country G Begins Diplomatic Efforts’ at time 1, &
‘International Community Threatens Sanctions’ at time 2, &
‘Country G Employs Successful Covert Operations’ at time 5]

Such a probability profile in TIN modeling provides important complementary information about the probability of success of a desirable effect when studied over time. First, the final probability level is given. Second, an unwanted drop in the probability level over time can be detected. And third, the time required to reach the final probability level is determined.




Figure 4: Example Probability Profile

Pythia’s Capabilities

On top of IN/TIN model building capabilities, Pythia has two tools developed to analyze and simulate IN/TIN models. The *Static Propagation* tool allows for computing the likelihood of occurrence of corresponding events in a static situation – when time data is not taken into account. A user can also perform two types of *Sensitivity Analysis* on a given IN model: (i) sensitivity of input – sensitivity of effect to actionable events, and (ii) sensitivity of influence – sensitivity of an effect to the CAST Logic parameters associated with edges.

In addition to TIN modeling, Pythia has a wide range of tools dedicated to TIN model analysis. For example, once a COA is defined, Pythia generates *probability profiles* of selected events over time, as described in Section 3. Pythia allows for *comparing COAs* based on computed probability profiles. A SAF tool (*Sets of Actions Finder*) finds various combinations of actions that cause the probability of a desired effect to be above a certain threshold. An *ECAD-EA* tool (Effective Course of Action Determination using Evolutionary Algorithms methodology) allows for finding the best COAs according to user specified metrics, temporal and casual constraints, and windows of observation and opportunity. Pythia also used APIs from the *TEMPER* temporal logic tool. Specifically the Point-Interval Logic (*PIL*) engine, together with a *What-If* tool provides an opportunity to transform a TIN into a corresponding point-graph and subsequently answer temporal queries and perform what-if analysis. Other tools allow for model transformation and more advanced model analysis.

References

- [1] Kansal, S. K., AbuSharekh, A., and Levis, A. H., “Computationally Derived Models of Adversary Organizations,” Proc. IEEE Symp. On Computational Intelligence for Security and Defense Applications 2007, Honolulu, HI, April 2007.
- [2] Haider, S., and Levis, A.H., “Dynamic Influence Nets: An Extension of Timed Influence Nets for Modeling Dynamic Uncertain Situations”. Proc. 10th International Command and Control Research and Technology Symposium, Washington DC, June 2005
- [3] Zaidi, A. K., and Levis, A. H., “TEMPER: A Temporal Programmer for Time-Sensitive Control of Discrete Event Systems,” IEEE Trans. on Systems, Man, and Cybernetics, vol. 31, No. 6, Nov. 2001, pp. 485-96
- [4] Jensen, F. V., *Bayesian Networks and Decision Graphs*, Springer-Verlag, 2001.
- [5] Chang, K. C., Lehner, P. E., Levis, A. H., Zaidi, S. A. K., and Zhao, X., “On Causal Influence Logic”, Technical Report, George Mason University, Center of Excellence for C3I, Dec. 1994.
- [6] Rosen, J. A., and Smith, W. L., “Influence Net Modeling with Causal Strengths: An Evolutionary Approach”, Proceedings of the Command and Control Research and Technology Symposium, Naval Post Graduate School, Monterey CA, Jun. 1996.

Appendix C

THE C2 WIND TUNNEL

The C2WT is an integrated, multi-modeling simulation environment. Its framework uses a discrete event model of computation as the common semantic framework for the precise integration of an extensible range of simulation engines, using the Run-Time Infrastructure (RTI) of the High Level Architecture (HLA) platform. The C2WT offers a solution for multi-model simulation by decomposing the problem into model integration and experiment or simulation integration tasks.

Model Integration: Integrated experiments or simulations are specified by a suite of domain specific models, including for instance: human organizations (expressed using the Colored Petri Net modeling language and social networks), networks (OMNET++ network simulation language), physical platforms (Matlab/Simulink based models), and the physical environment (e.g., Google Earth). While the individual behaviors simulated by the different simulation models are essential they must interact as specified by the workflow for the particular simulation or experiment. Their interactions need to be formally captured and the simulation of the components needs to be coordinated. This is a significant challenge, since the component models are defined using dramatically different domain specific modeling languages. The C2WT, therefore, uses the meta-modeling technology and the Vanderbilt MIC tool suite¹. The key new component is the Model Integration Layer (Fig. 2), where a dedicated Model Integration Language (MIL) is used for model integration. The MIL consists of a carefully selected collection of modeling concepts that represent the domain-specific simulation tools.

Model-based Experiment Integration: C2WT uses the MIC model interpretation infrastructure for the generators that automatically integrate heterogeneous experiments on the HLA platform deployed on a distributed computing environment. After finalizing the component models, the integration models, and setting the parameters, the MIL model interpreters generate all the necessary configuration information and run-time code. The architecture of the C2WT is shown in Fig. 2. Each modeling language is depicted as a federate on which models built using that language run.

Time Management in C2WT: Time Management is critical to preserve causality with simulations operating at different timescales. The C2WT builds upon the time management features of the underlying HLA standard, which has provision for both discrete time and discrete event models.

¹ MIC is a meta-programmable Model-Integrated Computing (MIC) tool suite for integrating models, to manage the configuration and deployment of scenarios in the simulation environment, and to generate the necessary interface code for each integrated simulation platform. It has evolved over two decades of research at the Institute for Software Integrated Systems at Vanderbilt University and is now used in a wide range of government and industry applications.

The main elements of time management in HLA are: a) a Logical Timeline, b) Time ordered delivery of interactions between simulations, and c) a protocol for advance of Logical Time. In a causality preserving execution (note that HLA supports untimed executions as well), the underlying RTI maintains a logical time, and interaction messages generated by simulations are time stamped with the logical time and delivered to their destinations in a timed order. The logical time is advanced by a cooperative Time Advance Request and Grant protocol. A similar protocol is supported for event driven simulation in which the event driven simulation requests the Next Event to the RTI. The simulation logical time is advanced either to the earliest available interaction or to the time stamp of the next event local to the requesting simulation

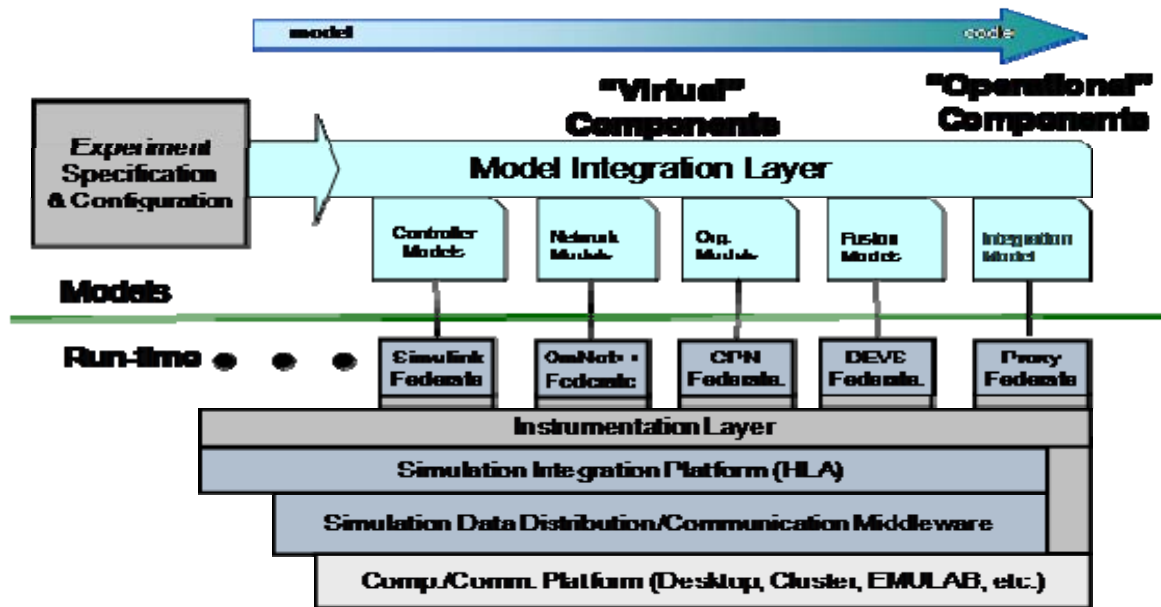


Fig. 1: The C2WT architecture

Modeling and Simulating Terrorist Networks in Social and Geospatial Dimensions

Il-Chul Moon and Kathleen M. Carley, *Carnegie Mellon University*

A simple theoretical multiagent model reasons about the criticality of terrorists and regions as terrorist interactions coevolve in geographical and social spaces.

Where we are influences who we know, and vice versa. As we move to new cities or countries, our contacts change. For instance, when a company relocates its employees, they develop new working relations with others while they perform assigned tasks. In theory, these relocations should improve company performance.¹

However, performance also depends on individuals knowing who to ask about what—that is, on transactive memory.² Moving disrupts transactive memory and the social relations by which information flows. So, the question arises whether performance can improve when social and geospatial distributions change simultaneously.

Social and spatial relations evolve over time. Estimating their evolutions is important for management, command and control structures, and intelligence analysis research. By knowing future agent social and spatial distributions, an analyst can identify emergent leaders, hot spots, and organizational vulnerabilities. Historically, such estimations have depended heavily on qualitative data analyses by subject-matter experts.³ A few researchers approached the issue using multiagent models and simulation. The models addressed the complex nature of the organization and task assignments, resource distributions, or agent locations. The simulations addressed the near-term organizational changes. This research came from two perspectives: the effects of change in the social network^{4,5} and the effects of geospatial change.^{6,7} Both perspectives can project aspects of emerging organizational structure and future performance, but they can't examine the interaction between physical and social movements.

We've developed a simple theoretical multiagent

simulation model to show how changes in the coevolution of social and geospatial dimensions affect group behavior. Our model overcomes the limitations of isolated social and spatial models (see the "Related Work in Social and Geospatial Modeling" sidebar). To illustrate the model's potential for reasoning, we examine its implications here for a real-world terrorist network, using data extracted from open source texts. Although a full validation would require additional field data, the model's output reveals important aspects of complex organizational evolution that apply beyond the counterterrorism domain.

Input data set

The model's input is a network representation of an organizational structure in the social and geospatial dimensions. It includes knowledge and task information: who knows what and who is using that knowledge. For the terrorist network, we extracted relevant data from unclassified documents, using the AutoMap text analysis tool.⁸ The documents included newspaper articles and unclassified intelligence reports from subject-matter experts. We hand-coded the corresponding latitudes and longitudes for the relevant data.

Figure 1a is an overall visualization of the resulting network consisting of four node types: agents, knowledge, tasks, and locations; figure 1b is a




Figure 1. (a) The overall visualization of the example terrorist network represents agents, knowledge bits, tasks, and locations as red, yellow, blue, and orange nodes, respectively. (b) The agent-to-agent network of the data set consists of three disconnected subnetworks.

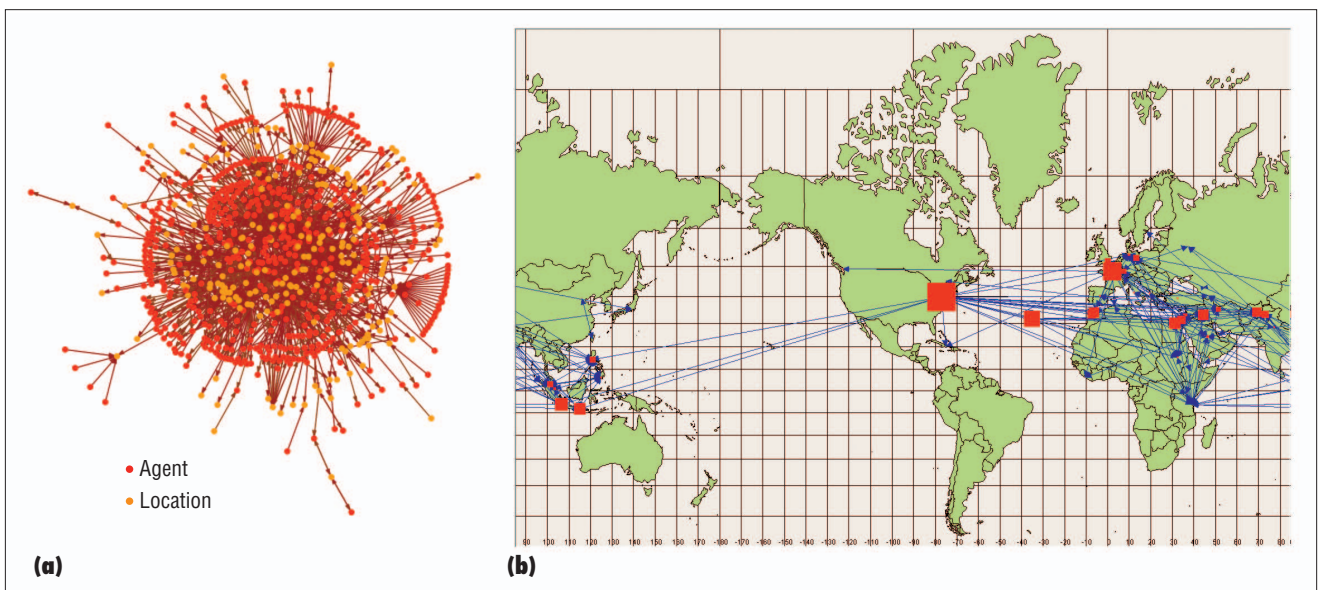


Figure 2. (a) Agent-to-location network. The red nodes represent agents; the orange nodes represent agent locations, which include latitude and longitude coordinates. (b) The AL network overlaid on a world map. To suggest how many agents are clustered at a specific location, the different-sized squares correspond to the number of agents in different regions. The blue edges display the agent-to-agent network links.

visualization of the data set's agent-to-agent (AA) links. Table 1 shows the input network's adjacency matrix, or *metamatrix*, across these nodes. This multimode, multilink network data represents the organization's current structural characteristics in our model.

We model the social dimension by using an algorithm that specifies the interaction probability between two agents. We model the geospatial dimension by using an agent-relocation mechanism that interprets agent movement in a geospatial location network's data set. For instance, if two agents have interactions or formal relations, we assume

that an AA link exists between them. Similarly, if an agent possesses a knowledge bit, we assume an agent-to-knowledge (AK) link between the nodes. If two locations appear in the same context, we regard the two locations as related (LL). This topological location network constitutes the agent-relocation dimension. The other subnetworks, such as an agent-to-task network, knowledge-to-location network, and task-to-location network, have their own intuitive meanings based on the connected node types and the data coder's perspective.

Figure 2a shows the agent-to-location (AL) network; figure 2b

Related Work in Social and Geospatial Modeling

Researchers who study people's movements concurrently through social relations and space mainly use two techniques: data mining and simulation. Data mining can uncover patterns such as an organization's network structure, entity properties, and entity clusters. For instance, in a summary of data mining's impact on the counterterrorism community, Jeff Jonas and Jim Harper claim that the 9/11 attack plan was available before the attack.¹ Uncovering the plan would have required extensive data mining on available databases, but the US government might have disrupted the plan by pursuing available leads. Although the authors make a counterterrorism case for data mining, they also note that high false positives, or incorrect predictions, could waste valuable resources.

Link analysis and discovery is another data mining technique applied to counterterrorism. Raymond Mooney and his colleagues use it with an inductive-logic-programming method to discover implied rules in multirelational data.² They describe a powerful tool for approximating a complete organizational network from an incomplete one.

Vandana P. Janeja, Vijayalakshmi Atluri, and Nabil R. Adam focus a modeling and simulation approach to detecting anomalous geospatial trajectories on the basis of spatiosemantic associations.³ They create basic spatial analysis units, or *spatial units*, and cluster them into a microneighborhood that shares similar characteristics across subspatial units. Their analysis of spatial and social characteristics at the same time is similar to our correlation between spatial and social dimensions.

Hsinchun Chen, Fei-Yue Wang, and David Zeng describe the development of an intelligence and informatics security model that depends heavily on network and link analysis.⁴ They examine three interesting uses of the model: cross-jurisdiction information sharing, terrorism information collection, and smart-border and bioterrorism applications. One of their applications, the West Nile Virus-Botulism Portal, includes hot-spot analysis and a prediction function.

Organizational-behavior research has benefited from agent-based modeling techniques. For instance, Kathleen Carley has made the efforts to model sociotechnical systems as networked multiagent structures.⁵ She introduces exemplary multiagent models such as OrgaHead⁶ and Construct.⁷ These models take networked organizational structures as an input and generate

the estimated performance of task accuracy and information diffusion over time as well as the evolved structures after simulation. This approach might be difficult to validate, but it represents an effort to create more complex, realistic models that can automatically generate hypotheses forecasting organizational behavior.⁸ Researchers could then use these hypotheses to estimate domain features or trends of interest and subsequently use other statistical analysis tools, such as data mining, to validate the hypotheses.

References

1. J. Jonas and J. Harper, *Effective Counterterrorism and the Limited Role of Predictive Data Mining*, policy analysis 584, Cato Inst., Dec 2006.
2. R.J. Mooney et al., "Relational Data Mining with Inductive Logic Programming for Link Discovery," *Data Mining: Next Generation Challenges and Future Directions*, H. Kargupta et al., eds., AAAI Press, 2004, pp. 239–254.
3. V.P. Janeja, V. Atluri, and N.R. Adam, "Detecting Anomalous Geospatial Trajectories through Spatial Characterization and Spatiosemantic Associations," *Proc. 2004 Ann. Nat'l Conf. Digital Government Research*, 2004, pp. 1–10.
4. H. Chen, F. Wang, and D. Zeng, "Intelligence and Security Informatics for Homeland Security: Information, Communication, and Transportation," *IEEE Trans. Intelligent Transportation Systems*, vol. 5, no. 4, 2004, pp. 329–341.
5. K.M. Carley, "Computational Organization Science: A New Frontier," *Proc. Nat'l Academy of Science*, vol. 99, no. 3, 2002, pp. 7314–7316.
6. K.M. Carley and D.M. Svoboda, "Modeling Organizational Adaptation as a Simulated Annealing Process," *Sociological Methods and Research*, vol. 25, no. 1, 1996, pp. 138–168.
7. C. Schreiber and K.M. Carley, "Going beyond the Data: Empirical Validation Leading to Grounded Theory," *Computational and Mathematical Organization Theory*, vol. 10, 2004, pp. 155–164.
8. K.M. Carley, "On Generating Hypotheses Using Computer Simulations," *Systems Eng.*, vol. 2, no. 2, 1999, pp. 69–77.

Table 1. A metamatrix of the input data set for a terrorist network organizational structure.*

Nodes	Agent	Knowledge	Task	Location
Agent (916)	Social network (AA, 0.0024)	Knowledge network (AK, 0.0093)	Assignment network (AT, 0.0070)	Deploy network (AL, 0.0026)
Knowledge (614 bits)	—	Not used	Needs network (KT, 0.0961)	Regional knowledge network (KL, 0.0692)
Task (258)	—	—	Not used	Regional task network (TL, 0.1042)
Location (387)	—	—	—	Proximity network (LL, 0.0799)

*The number of nodes and network densities are in parentheses.

overlays it on a geographic map. Details on the coding process are available elsewhere.⁹

Model summary

The model simulates each agent and its interaction with others to estimate changes over time in organizational performance and structure. As agents interact and learn, their behavior eventually changes the performance and structure. The following algorithm outlines the interaction and relocation mechanisms for agent A's behavior:

1. A searches for locations within its vision range (VR), looking for unknown

Table 2. Model input, output, parameters, and internal variables.*

Type	Name	Implication
Input	A networked organizational structure	A network including agents, knowledge bits, tasks, and locations. The network represents the target domain's complex organizational structure.
Output	An evolved network organization	A network organization with a recreated agent-to-agent (AA) network and an agent-to-location (AL) network, both of which reflect interactions and relocations.
	Knowledge diffusion	A performance metric showing how fast information can diffuse across the network.
	Energy task accuracy	A performance metric showing how accurately information is distributed to agents who require it to complete their tasks.
	Gini coefficient for AA and AL	Coefficients indicating the extent of unequal distribution of AA and AL network criticalities.
Parameters	Simulation runtime step (default = 30 steps)	The total simulation runtime.
	Number of replications (default = 3)	The number of model runs (required because the model is stochastic, not deterministic).
	Move radius (MR)	The radius on the spatial-route network specifying the maximum distance an agent can move in one time step.
	Vision range (VR) (default = 1)	The range on the spatial-route network specifying an agent's ability to gather a knowledge bit or interact with another agent.
	Sphere of influence (SI) (default = 2)	The number of social links that an agent can cross for an interaction.
	Relative-similarity (RS) weight, w_1 ; relative-expertise (RE) weight, w_2 ; social-distance (SD) weight, w_3 (default = 0.5); and spatial-proximity (SP) weight, w_4 (default = 0.5)	The weights to calculate four interaction probabilities.
	Learning rate from an agent (default = 0.05)	The possibility that an agent can learn a knowledge bit from an interaction with another agent.
	Learning rate from a location (default = 0.025)	The possibility that an agent can gather a knowledge bit by observing a knowledge node within vision range.
Internal variables	Relative similarity (RS_{ij})	The likelihood of interactions caused by homophily between i and j (passive information seeking).
	Relative expertise (RE_{ij})	The likelihood of interactions caused by expertise between i and j (active information seeking).
	Social distance (SD_{ij})	The likelihood of interactions over multiple social links.
	Spatial proximity (SP_{ij})	The likelihood of interactions from spatial distance.
	Interaction candidate set (ICS_i)	The agent set with which agent i can interact.
	Probability of interaction ($P_{ij}^{Interaction}$)	The likelihood of agent i 's interaction with agent j , calculated as the weighted linear sum of RS, RE, SD, and SP.
	Probability of relocation ($P_{il}^{Relocation}$)	The likelihood of agent i 's moving to location l , determined by the number of available knowledge bits required to perform the agent's assigned tasks.

* We use model default values here except for the move radius (MR), relative-similarity weight (w_1), and relative-expertise weight (w_2).

- but necessary knowledge bits.
2. A moves to a found location.
3. A learns the unknown knowledge at its location.
4. A selects an agent from those that qualify as communication candidates.
5. A exchanges the unknown knowledge with the selected agent.

Basically, agents can interact and relocate at each simulation time step. They select a location to move to and an agent to interact with according to probabilistic values for each interaction and relocation opportunity. Exactly which agents interact with which, when they interact, what choices they make, and what they communicate and learn are defined probabilistically. Consequently, the model is sto-

chastic and, as such, requires multiple replications to generate stable results and to define the space of outcomes.

Table 2 lists several factors that drive an agent's behavior and so the network's evolution and organization's structure. For example, agent behavior depends on the given input data set, which sets the initial environment. The input determines the initial probability of interaction among agents according to what they know and where they are located. The model's parameters include the relocation (move) radius in the geospatial dimension, the interaction (sphere of influence) radius in the social dimension, and the probability of learning after a knowledge exchange with an agent or a knowledge gathering at a certain location. Finally, the internal variables reflect behaviors calculated from the defined inputs and parameters, according to various model formulas.

We tested this model by varying important parameters in the agent interactions and relocations. First, we changed the agent move radius (MR) by 0, 1, and 2. If an agent's MR is 0, it's stationary to its initial location. If its MR is 2, the agent can search locations linked by two LL links from its initial location. Next, we varied the weight of relative similarity (RS) and relative expertise (RE) contributing to the probability interaction. If the RS weight is high, the agents interact mainly with agents sharing similar backgrounds, beliefs, and knowledge. This imitates agents as passive information receivers. In contrast, a higher RE weight makes the agents active information seekers. Finally, we tested the input data's sensitivity by randomly dropping or adding links in the AA or LL networks.

Agent-interaction mechanism

Agents have the opportunity to interact during each time period. They select an agent to interact with according to a probability of interaction, P , that's a weighted sum of four different factors: RS, RE, social distance (SD), and spatial proximity (SP). The theory for these factors comes from sociology, communication theory, and counterterrorism analysis.

Because the model is stochastic, an agent will usually interact with agents that it has a higher probability of choosing but will occasionally end up with a less likely choice. Like humans, these simulated agents can't always talk to their first choice. The model thus captures interactions that reflect less-than-optimal connections between intention and action as well as the rare unexpected interaction.

After an agent chooses another agent to interact with, the two agents will exchange knowledge bits. For each exchanged knowledge bit, the model draws a number from a uniform distribution ranging from 0 to 1. If the number is within the receiving agent's learning rate, that agent will have a new link to the communicated knowledge piece in the AK network.

Relative similarity and relative expertise. RS is a ratio reflecting similarity in the choosing and chosen agents' knowledge. It's based on the sociological principle of *homophily*,¹⁰ which describes the increased likelihood of a person interacting with another person who shares similar education, beliefs, or race. RS represents the probability of a terrorist interacting with other terrorists that share the same religion or nationality. RE is a ratio reflecting the amount of knowledge the chosen agent has that the chooser doesn't have, and it's based on transactive memory.² RE captures why a Middle Eastern terrorist interacts with a South American drug cartel to exchange weapons expertise or information about funding sources. At first glimpse, the two factors might seem contradictory, but they're just two metrics capturing different aspects of terrorist knowledge-acquisition attitudes:

$$RS_{ij} = \frac{\sum_{k=0}^K AK_{ik} AK_{jk}}{\sum_{k=0}^K AK_{ik}}, RE_{ij} = \frac{\sum_{k=0}^K AK_{jk} (1 - AK_{ik})}{\sum_{k=0}^K AK_{ik}} \quad (1)$$

where K is the number of knowledge bits.

Social distance. SD is another factor affecting agents' probability of interaction—if two agents must cross many social links, then the probability should be low, and vice versa.¹¹ We compute it by finding the shortest path between two agents and then dividing one by the number of links in that path.

$$SD_{ij} = \frac{1}{|AA_{ij}|} \quad (2)$$

$$|AA_{ij}| = \begin{cases} \text{No. of links on shortest path from } i \text{ to } j \\ (\text{no. of links} \leq SI) \\ SI + 1 \text{ (no. of links} > SI) \end{cases}$$

If SD is larger than the maximum number of links in the sphere of influence, SI, then SD is set to one plus the maximum for social-interaction perimeter modeling. An agent can recognize and distinguish the closeness of other agents within the SI perimeter, but it

can't differentiate the closeness when the interacting agent is outside the perimeter. In this case, an agent regards the interacting agents as just $SI + 1$ links away, though the real SD might differ.

Spatial proximity. Intuitively, two persons at the same location are more likely to talk than are two at different locations.^{12–14} Some might argue that SP isn't significantly correlated with interaction frequency in the Internet age. However, in the terrorism domain, attending the same training camp or the same mosque is a critical interaction indicator.¹⁴ The SP model is similar to SD but indicates the probability of being at the same location, rather than having a social link:

$$SP_{ij} = \frac{1}{\sum_{l_1=0}^L \sum_{l_2=0}^L (|LL_{l_1 l_2}| + 1) AL_{il_1} AL_{jl_2}} \quad (3)$$

$$|LL_{l_1 l_2}| = \begin{cases} \text{No. of links on shortest path from } l_1 \text{ to } l_2 \\ (\text{no. of links} \leq VR) \\ VR + 1 \text{ (no. of links} > VR) \end{cases}$$

As with SD, if SP is greater than VR, which is a maximum communication range across the geospatial dimension, and chosen by the user, the model sets SP to one plus the maximum for computing convenience. The rationale for using VR in the geospatial-domain calculation is the same as the rationale for using SI in the social dimension.

Probability of interaction. Using the four different factors we've described, we can express the probability that agents will select another agent to interact with as a weighted sum:

$$P_{ij}^{\text{Interaction}} = w_1 RS_{ij} + w_2 RE_{ij} + w_3 SD_{ij} + w_4 SP_{ij} \quad (4)$$

Although the model can calculate the probability for any pair of agents, we limit the number of possible interaction candidate agents according to two distances, SD and SP. This restriction assumes that

Agents have the opportunity to interact during each time period. They select an agent to interact with according to a probability of interaction that's a weighted sum of four different factors.

Table 3. Virtual-experiment design: Sensitivity analysis and parameter-space exploration.

Input parameters	Value	Implication
Move radius (MR)	0, 1, or 2 (3 cases)	Parameter-space exploration, examining the results' sensitivity according to the agent-movement perimeter (MR parameter)
Weights for RS (w_1)/RE (w_2)	0/1, 0.25/0.75, 0.6/0.4, 0.75/0.25, 1/0 (5 cases)	Parameter-space exploration, examining the agent-interaction attitudes and their affect on the results, from passive information gathering to active information gathering
Density of the organizational-structure network (AA and LL densities)	75%, 100%, 125% (3 cases)	Sensitivity analysis, examining the sensitivity of results according to the density changes of the AA and LL networks corresponding to the social and geospatial dimensions, respectively
Total virtual experiment cells	45 cells ($3 \times 5 \times 3$ cases)	—

a person will interact with others in his or her neighborhood—either social or geographic. Formally, the model defines the interaction candidate set as

$$ICS_i = \left\{ A_j \mid \left(\left| AA_{ij} \right| \leq SI \right) \vee \left(\left| LL_{i_1 i_2} \right| \mid AL_{i_1} AL_{j_2} \leq VR \right) \right\} \quad (5)$$

An agent can communicate only with its candidate agents, so the probability of interaction is calculated between each agent and its candidate agents.

Agent-relocation mechanism

Our model lets agents relocate themselves to adjacent locations. The MR parameter defines the sphere of relocation, but the probability of choosing a certain location is more complicated:

$$P_{il}^{Relocation} = \frac{1}{\sum_{t=0}^T \sum_{k=0}^K AT_{it} \times KT_{kt} \times |KL_{kl}|} \quad (6)$$

$$|KL_{kl}| = \begin{cases} \text{No. of links on shortest path from } l \text{ to } k \\ \text{(no. of links} \leq VR) \\ VR + 1 \text{ (no. of links} > VR) \end{cases}$$

In essence, the agents choose a location that, on average, guarantees the shortest path to their required knowledge bits. In other words, the agents try to put themselves at the optimal location to collect the knowledge they want. However, like the AA interaction model, this is a stochastic model that determines location choices probabilistically. So, it's possible to choose a nonpreferred location with lower probability.

After selecting a location, the model changes the AL network by removing the edge from the agent to the old location and adding an edge to the new location. Additionally, the agent will gather knowledge bits linked to locations in its VR. This knowledge gathering is similar to the knowledge exchange between agents, except it uses a different learning rate. Some might argue that this regional knowledge acquisition isn't necessarily true, especially in the real world where terrorists can learn new knowledge from Web sites. However, many terrorists go to training sites and organization headquarters to receive specific, detailed training. These relocations are an important issue in the counterterrorism field,¹⁴ and we're specifically examining them in this example.

Output measures

We use two performance metrics to evaluate an evolving organization over time: *knowledge diffusion* and *energy task accuracy*. KD gauges the dispersion of the knowledge bits across the agents

as follows:

$$KD = \frac{\sum_{i=0}^A \sum_{j=0}^K AK_{ij}}{K \times A} \quad (7)$$

But KD considers only who knows what. ETA calculates the extent to which the agents have the knowledge they need to do their assigned tasks. This calculation introduces the agent-to-task (AT) and knowledge-to-task (KT) networks:

$$ETA = \frac{100}{T} \sum_{t=0}^T \frac{\sum_{k=0}^K \left(KT_{kt} \times \sum_{a=0}^A AK_{ak} \right)}{\sum_{a=0}^A AT_{at} \times \sum_{k=0}^K KT_{kt}} \quad (8)$$

Furthermore, we define two criticality metrics for the agents and locations. For agents, we count the number of agents that an agent interacts with during the simulation. This represents the number of agents that the agent knows and influences. For locations, we count the number of agents in a location at the end time. If the location harbors more agents, it might have higher terrorist activities.

Results

We used the model to analyze the terrorist network in the meta-matrix format (see table 1) and to generate estimates on agent relocation, geospatial clustering, agent interaction, and social-network evolution. We performed a sensitivity analysis first, then visualized and analyzed the model output in two dimensions.

We replicated the simulation three times for 30 simulation time steps. The sensitivity analysis showed significant p-values for some independent factors. Specifically, MR is a significant predicting factor for ETA, KD and the Gini coefficient are significant for location-criticality distribution, and RS is important for explaining the Gini coefficient of the agent-criticality distribution. (The Gini coefficient comes from economics and describes a property's distribution across a population.) These p-values indicate that three replications are sufficient for identifying the critical factors for each performance metric. The stabilized distribution of the results requires further examination, but that work is outside this article's scope.

Sensitivity analysis

We analyzed the model's sensitivity by varying the input parameters in table 3. After running the model with varied parameters, we performed a regression analysis. The independent variables are the varied parameters of a *virtual-experiment cell*—for example, a com-

Table 4. Regression for sensitivity analysis.

Dependent variable	Energy task accuracy	Knowledge diffusion	Gini coefficient of location-criticality distribution	Gini coefficient of agent-criticality distribution
Standardized coefficients				
Move radius	0.748*	0.780*	-0.956*	-0.088
Relative similarity	0.008	0.004	0.020	0.131†
Possible density	0.010	0.009	-0.114*	-0.865*
Adjusted R-square	0.506	0.555	0.925	0.765

*p-value < 0.001

†p-value < 0.01

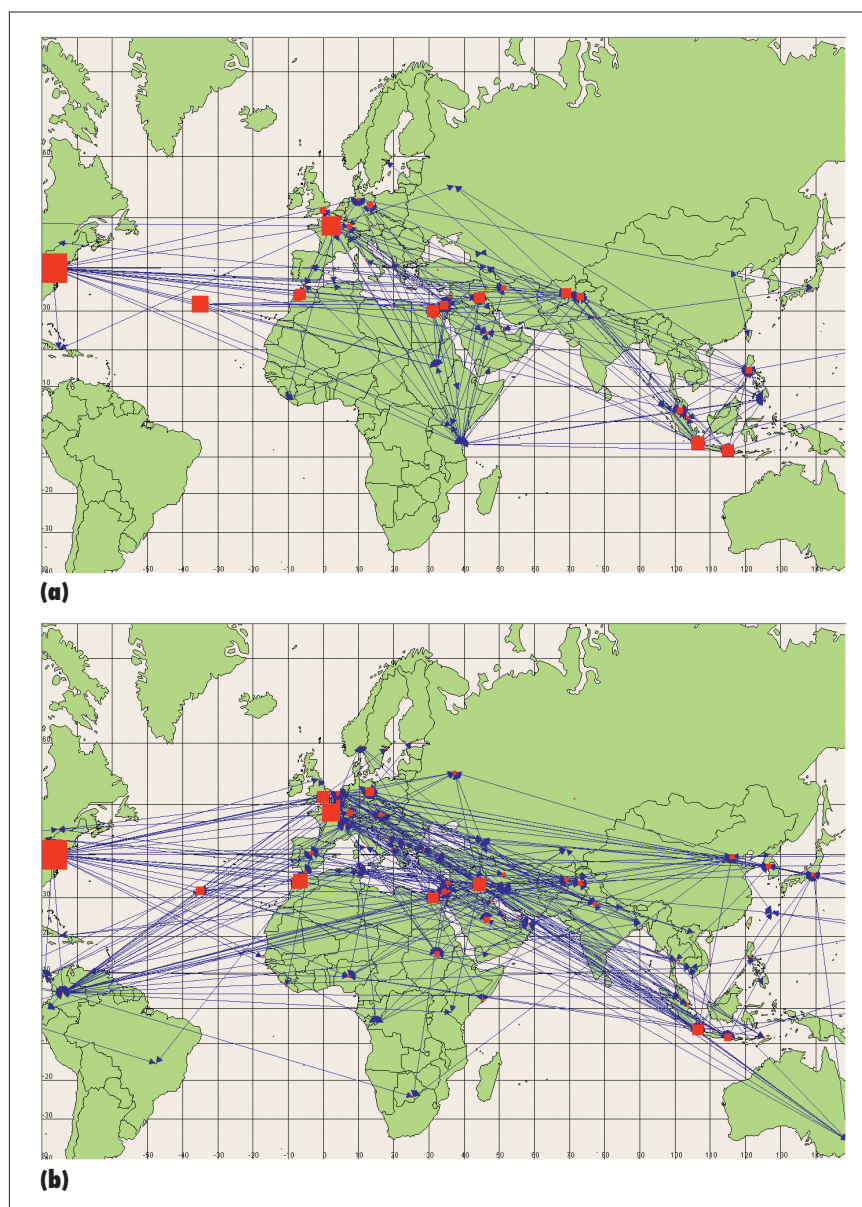


Figure 3. The agent geospatial distribution changes over time with a move radius of 1: (a) the distribution at time 0, (b) the distribution at time 30. The size of the squares corresponds to the number of agents in the region, and the lines are the interlocation agent-to-agent communication links.

bination of 0 MR, 0/1 RS/RE ratio, and 75 percent density rate. The dependent variables are the two performance metrics and the Gini coefficients of the agent- and location-criticality distributions.

Table 4 is the regression-analysis result. First, as MR increases, the network's performance improves. The terrorists in the model tend to relocate to regions where they can collect more information, rather than stay in their current location. Furthermore, these relocations increase task performance by increasing the information feed. Next, higher MR and higher possible density decreases the Gini coefficient of location criticalities. This indicates that terrorists will disperse more if they can relocate more easily and the input network is denser. Finally, lower RS will induce a more centralized terrorist network. Particularly, the input network density has a great impact on the agent-criticality distribution compared with its impact on the location-criticality distribution.

Location-criticality analysis

Agent movement creates segregation patterns over time (see figure 3). Figure 4 shows an accumulated agent distribution across the locations. The distribution implies that agents will disperse more if we increase MR: the fewer places harboring terrorists, the greater the MR, which should help the terrorists find the places to cluster. However, our model indicates the opposite scenario: the terrorists in our model can't find the places to cluster densely. Rather than gathering in a few regions, the terrorists will disperse around the world.

Table 5 lists the top 10 locations harboring terrorists after the simulations. Although the accumulated distribution and its Gini coefficient in figure 4 showed terrorist dispersion, the top 10 locations are fairly consistent across three different MR levels. This implies that the hot regions with frequent terrorist

activity will remain at the top after the relocations, even though some terrorists in those regions move to regions with less activity. In detail, the northwest African regions—that is, Morocco and Casablanca—become important locations as well as some European regions, such as France. The south Asian regions of Indonesia and Bali and the areas of frequent activity—US and Israel—will remain the same.

Agent-criticality analysis

We analyzed the important agents after the simulation. According to the sensitivity analysis, RS changes impact the distribution of the agent criticality. Figure 5 visualizes the accumulated agents' social link coverage across the RS levels. It shows some slight differences in terms of Gini coefficients, but the link-coverage distribution doesn't change much. This implies that the terrorist social network's evolution is stable regardless of the parameter change. In spite of the small changes, the increase in Gini coefficient with higher RS suggests that fewer terrorists will control the social links if the terrorists gather information more passively. For instance, one terrorist group often has different backgrounds from another group. In that case, under a strong RS interaction weight, only terrorists with backgrounds similar to both groups will be able to communicate with the groups' members. A strong homophily trend means that agents will have fewer possible agents within their ICS and that fewer agents will control more social links.

As with location-criticality analysis, we identified the top 10 terrorists who control the most links after simulation. Table 6 shows that the top terrorists, such as Bin Laden and Riduan Isamuddin, have similar power after simulations in spite of varying parameters. This is because they're already the center of terrorist social networks, so they appear frequently in ICSs. Additionally, they have fairly comprehensive backgrounds and knowledge, so most agents can find high RS and RE with the top-ranking agents. On the other hand, Mohammad Atta shows higher ranks under the passive-information-gathering assumption, because his background was common across the agents.

Our analysis indicates that the agents become more dispersed around the world but that critical agents themselves don't change much. Obviously, the analysis method has its limitations. First, validating the simulation model is very challenging and involves open research questions, such as matching the simulated time step




Figure 4. An accumulated distribution of agents across the locations. The whereabouts of 570 agents are known, and there are 387 locations.

Table 5. The top 10 critical locations.

Rank	MR = 0 (stationary)	MR = 1 (adjacent move)	MR = 2 (farther move)
1	US	US	US
2	Israel	France	France
3	France	Morocco	Morocco
4	Bali	Israel	Casablanca
5	Morocco	Bali	Bali
6	Egypt	Casablanca	Egypt
7	Afghanistan	Egypt	Israel
8	Casablanca	Iraq	Strasbourg
9	Iraq	Indonesia	Gaza
10	Indonesia	Strasbourg	Indonesia

to the real-time flow. Also, incorrect input data sets can misdirect the model's output. Complete and correct real-world data sets are rare, but we expect to resolve some concerns by adding more realistic agent-behavior mechanisms. As the subjects' behaviors become more complex, adding more salient features to the model will increase its usability. A recent book addresses defense modeling, simulation, and analysis issues further.¹⁵

Despite some concerns, this complex multiagent model generates several estimates that are useful for policy making and theory building. Furthermore, the formula-based, agent-behavior design can be updated easily as findings from other disciplines become available. These two points provide incentives for using the model in the real world and for updating and developing it on the basis of future findings. ■




Figure 5. The accumulated distribution of the percentage of covered social links across the agents.

Table 6. The top-10 critical agents.

Rank	RS = 0.0 (active information gathering)	RS = 0.6	RS = 1.0 (passive information gathering)
1	Bin Laden	Bin Laden	Bin Laden
2	Riduan Isamuddin	Riduan Isamuddin	Riduan Isamuddin
3	Abdul Aziz	Abdul Aziz	Mohammed Atta
4	Yasser Arafat	Yasser Arafat	Bakar Bashir
5	Bakar Bashir	Yaacov Perry	Yasser Arafat
6	Mohammed Atta	Mohammed Atta	Zacarias Moussaoui
7	Yaacov Perry	Bakar Bashir	Yaacov Perry
8	Imam Samudra	Zacarias Moussaoui	Abdul Aziz
9	Zacarias Moussaoui	Mohamedou Slah	Yazid Sufaat
10	Abdullah Sungkar	Abdullah Sungkar	Mohamedou Slah

Acknowledgments

Support for this work came in part from the US Office of Naval Research (ONR N00014-06-1-0921, N00014-06-1-0104, and N00014-06-1-0772), National Science Foundation (SES-0452487), Army Research Lab, and Air Force Office of Sponsored Research (MURI: Cultural Modeling of the Adversary, 600322) for research in dynamic network analysis. Additional support came from Carnegie Mellon University’s Center for Computational Analysis of Social and Organizational Systems. The views and conclusions in this article are those of the authors and shouldn’t be interpreted as representing the official policies, either expressed or implied, of the Office of Naval Research, the National Science Foundation, the Army Research Lab, or the US government.

References

1. D.C. Feldman and H.B. Thompson, “Expatriation, Repatriation, and Domestic Geographical Relocation: An Empirical Investigation of Adjustment to New Job Assignments,” *J. Int’l Business Studies*, vol. 24, 1993, pp. 507–529.

2. D.M. Wegner, “Transactive Memory: A Contemporary Analysis of the Group Mind,” *Theories of Group Behavior*, B. Mullen and G.R. Goethals, eds., Springer, 1986, pp. 185–205.

3. J. Arquilla and D. Ronfeldt, eds., *Networks and Netwars: The Future of Terror, Crime, and Militancy*, tech. report MR-1382-OSD, Rand, 2001; www.rand.org/publications/MR/MR1382.

4. K.M. Carley, J.S. Lee, and D. Krackhardt, “Destabilizing Networks,” *Connections*, vol. 24, no. 3, 2001, pp. 31–44, 2001.

5. T.A.B. Snijders, C. Steglich, and M. Schweinberger, “Modeling the Co-evolution of Networks and Behavior,” *Longitudinal Models in the Behavioral and Related Sciences*, K. Van Montfort, H. Oud, and A. Satorra, eds., Lawrence Erlbaum, 2006, pp. 41–71.

6. J. Epstein, J.D. Steinbruner, and M.T. Parker, “Modeling Civil Violence: An Agent-Based Computational Approach,” working paper no. 20, Center on Social and Economic Dynamics, Brookings Inst., 2001.

7. M. Bergkvist et al., “A Hybrid Micro-Simulator for Determining the Effects of Governmental Control Policies on Transport Chains,” *Proc. Joint Workshop Multiagent and Multiagent-Based Simulation*, LNAI 3415, Springer, 2005, pp. 236–247.

8. J. Diesner and K.M. Carley, “Revealing and Comparing the Organizational Structure of Covert Networks with Network Text Analysis,” *Program and Abstracts: XXV Int’l Sunbelt Social Network Conf. Program*, 2005, p. 138; www.socsci.uci.edu/~ssnconf/conf/SunbeltXXVProgram.pdf.

9. K.M. Carley et al., “Toward an Interoperable Dynamic Network Analysis Toolkit,” to be published in *Decision Support Systems*, vol. 43, no. 4, 2007, pp. 1324–1347.

10. M. McPherson, L. Smith-Lovin, and J.M. Cook, “Birds of a Feather: Homophily in Social Networks,” *Ann. Rev. Sociology*, vol. 27, 2001, pp. 415–444.

11. D.J. Watts, P.S. Dodds, and M.E.J. Newman, “Identity and Search in Social Networks,” *Science*, vol. 296, no. 5571, 2002, pp. 1302–1305.

12. O. Sorenson and T.E. Stuart, “Syndication Networks and the Spatial Distribution of Venture Capital Investments,” *Am. J. Sociology*, vol. 106, 2001, pp. 1546–1588.

13. C.T. Butts, “Spatial Models of Large-Scale Interpersonal Networks,” doctoral dissertation, Dept. Social and Decision Sciences, Carnegie Mellon Univ., 2002.

The Authors



Il-Chul Moon is a PhD student at Carnegie Mellon University. His theoretic research interests include social-network analysis, multiagent modeling and simulation, machine learning, and artificial intelligence. His practical research interests include military command and control structures, counterterrorism, and disaster management. He received his MSc in computation, organization and society from Carnegie Mellon University.

He's a student member of the IEEE. Contact him at the School of Computer Science, Carnegie Mellon Univ., Pittsburgh, PA 15213; icmoon@cmu.edu.



Kathleen M. Carley is a professor of computer science in the Institute of Software Research International and the director of the center for Computational Analysis of Social and Organizational Systems. Her research combines cognitive science, social networks, and computer science to address complex social and organizational problems in areas such as dynamic network analysis, computational social and organization theory,

adaptation and evolution, text mining, and the impact of telecommunication technologies and policy on communication, information diffusion, disease contagion, and response within and among groups, particularly in disaster or crisis situations. She received her PhD in sociology from Harvard University. Contact her at the School of Computer Science, Carnegie Mellon Univ., Pittsburgh, PA 15213; carley@cs.cmu.edu.

14. M. Sageman, *Understanding Terror Networks*, Univ. Pennsylvania Press, 2004, pp. 142–146.
15. S. Keller-McNulty et al., *Defense Modeling, Simulation, and Analysis: Meeting the Challenge*, Nat'l Academies Press, 2006, pp. 64–73.

For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.



Log on to our Web site, where you can

- Search our vast archives • Preview upcoming topics
- Browse our calls for papers • Submit your article for publication • Subscribe or renew

www.computer.org/intelligent

IEEE computer society

PURPOSE: The IEEE Computer Society is the world's largest association of computing professionals and is the leading provider of technical information in the field.

MEMBERSHIP: Members receive the monthly magazine *Computer*, discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others interested in the computer field.

COMPUTER SOCIETY WEB SITE: www.computer.org

OMBUDSMAN: Email help@computer.org.

Next Board Meeting: 9 Nov. 2007, Cancún, Mexico

EXECUTIVE COMMITTEE

President: Michael R. Williams*

President-Elect: Rangachar Kasturi; *** Past President:** Deborah M. Cooper;*

VP, Conferences and Tutorials: Susan K. (Kathy) Land (1ST VP); *** VP,**

Electronic Products and Services: Sorel Reisman (2ND VP); *** VP, Chap-**

ters Activities: Antonio Doria; *** VP, Educational Activities:** Stephen B.

Seidman; **† VP, Publications:** Jon G. Rokne; **† VP, Standards Activities:**

John Walz; **† VP, Technical Activities:** Stephanie M. White; *** Secretary:**

Christina M. Schober; *** Treasurer:** Michel Israel; **† 2006–2007 IEEE Divi-**

sion V Director: Oscar N. Garcia; **† 2007–2008 IEEE Division VIII Direc-**

tor: Thomas W. Williams; **† 2007 IEEE Division V Director-Elect:** Deborah

M. Cooper; *** Computer Editor in Chief:** Carl K. Chang; **† Executive**

Director: Angela R. Burgess†

* voting member of the Board of Governors † nonvoting member of the Board of Governors

BOARD OF GOVERNORS

Term Expiring 2007: Jean M. Bacon, George V. Cybenko, Antonio Doria,

Richard A. Kemmerer, Itaru Mimura, Brian M. O'Connell, Christina M.

Schober

Term Expiring 2008: Richard H. Eckhouse, James D. Isaak, James W. Moore,

Gary McGraw, Robert H. Sloan, Makoto Takizawa, Stephanie M. White

Term Expiring 2009: Van L. Eden, Robert Dupuis, Frank E. Ferrante, Roger

U. Fujii, Ann Q. Gates, Juan E. Gilbert, Don F. Shafer

EXECUTIVE STAFF

Executive Director: Angela R. Burgess; **Associate Executive Director:** Anne

Marie Kelly; **Associate Publisher:** Dick Price; **Director, Administration:** Vio-

let S. Doan; **Director, Finance and Accounting:** John Miller

COMPUTER SOCIETY OFFICES

Washington Office. 1730 Massachusetts Ave. NW, Washington, DC 20036-1992

Phone: +1 202 371 0101 • Fax: +1 202 728 9614

Email: hq.ofc@computer.org

Los Alamitos Office. 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-1314

Phone: +1 714 821 8380 • Email: help@computer.org

Membership and Publication Orders:

Phone: +1 800 272 6657 • Fax: +1 714 821 4641

Email: help@computer.org

Asia/Pacific Office. Watanabe Building, 1-4-2 Minami-Aoyama, Minato-ku,

Tokyo 107-0062, Japan

Phone: +81 3 3408 3118 • Fax: +81 3 3408 3553

Email: tokyo.ofc@computer.org

IEEE OFFICERS

President: Leah H. Jamieson; **President-Elect:** Lewis Terman; **Past**

President: Michael R. Lightner; **Executive Director & COO:** Jeffrey W.

Raynes; **Secretary:** Celia Desmond; **Treasurer:** David Green; **VP,**

Educational Activities: Moshe Kam; **VP, Publication Services and**

Products: John Baillieul; **VP, Regional Activities:** Pedro Ray; **President,**

Standards Association: George W. Arnold; **VP, Technical Activities:**

Peter Staecker; **IEEE Division V Director:** Oscar N. Garcia; **IEEE Division**

VIII Director: Thomas W. Williams; **President, IEEE-USA:** John W.

Meredith, P.E.



revised 25 June 2007

Optimization of Actions in Activation Timed Influence Nets

M. Faraz Rafi, Abbas K. Zaidi and Alexander H. Levis

System Architectures Laboratory, ECE Dept., George Mason University, Fairfax, VA USA

E-mail: {mrafi1, szaidi2, alevis}@gmu.edu

P. Papantoni-Kazakos

EE Dept., University of Colorado Denver, Denver, CO USA

E-mail: Titsa.Papantoni@cudenver.edu

Keywords: influence net, activation timed influence net, Bayesian net

Received: October 12, 2008

A sequential evolution of actions, in conjunction with the preconditions of their environment and their effects, are all depicted by Activation Timed Influence Nets. In this paper, we develop two algorithms for the optimal selections of such actions, given a set of preconditions. A special case for the two algorithms is also considered where the selection of actions is further constrained by the use of dependencies among them. The two algorithms are based on two different optimization criteria: one maximizes the probability of a given set of target effects, while the other maximizes the average worth of the effects' vector.

Povzetek: Predstavljena sta dva algoritma za optimizacijo akcij v časovno odvisnih mrežah.

1 Introduction

We consider the scenario¹ where a sequence of actions needs to be initialized towards the materializing of some desirable effects. As depicted in Figure 1, each action is supported by a set of preconditions and gives rise to a set of effects; the latter become then the preconditions of the following action(s) which, in turn, gives rise to another set of effects. Such sequential evolution of actions is termed Activation Timed Influence Nets (ATINs), where the action performers may be humans. ATINs are an extension of an earlier formalism called Timed Influence Nets (TINs) [6-12, 20-27, 30, 31] that integrate the notions of time and uncertainty in a network model. The TINs are comprised of nodes that represent propositions (i.e., pre-and post-conditions of potential actions as well as assertions of events which may indirectly describe such actions), connected via causal links that represent relationships between the nodes, without any explicit representation of actions. TINs have been experimentally used in the area of Effects Based Operations (EBOs) for evaluating alternate courses of actions and their effectiveness to mission objectives in a variety of domains, e.g., war games [20-22, 25], and coalition peace operations [24, 27], to name a few. A number of analytical tools [6-12, 23, 24, 27, 30] have also been developed over the years for TIN models to help an analyst update conditions/assertions, represented as nodes in a TIN, to map a TIN model to a Time Sliced Bayesian Network for incorporating feedback evidence, to determine best set of pre-conditions for both timed and un-timed versions of Influence Nets, and to assess temporal aspects of the in-

fluences between nodes. A recent work [31] on TINs, underlying constructs and the computational algorithms, provides a comprehensive analytical underpinning of the modeling and analysis approach.




Figure 1: Network Representation of an Activation Timed Influence Net (ATIN)

In contrast to their predecessors (i.e., TINs), ATINs explicitly incorporate as nodes the *mechanisms* and/or actions that are responsible for changes in the state of a domain; other nodes represent preconditions and effects of actions. A set of preconditions may support a number of different actions, each of which may lead to the same effects, with different probabilities and different costs/awards, however. The objective is to select an optimal set of actions, where optimality is determined via a pre-selected performance criterion. In this paper, we present two algorithms which attain such an objective. We note that an effort to develop an action selection algorithm is also presented in [1].

The organization of the paper is as follows: In Section 2, we present the core formalization of the problem,

¹ This work was supported by the Air Force Office of Scientific Research (AFOSR) under Grants FA9550-05-1-0106 and FA9550-05-1-0388.

including two different optimization criteria. In Section 3, we derive the two algorithms which address the latter criteria. In Section 4, we express the extensions of the two algorithms to the network propagation scenario. In Section 5, we include numerical evaluations while in Section 6, we draw some conclusions.

1.1 Related Work

ATINs include action planning. In the domain of action planning, classical planners assume that the effects of an action are known with certainty and generate a set of actions that will achieve the desired goals [19]. Some planners do monitor for errors as actions are executed, but no action adaptations are incorporated [29]. Other planners assign probabilities to the effects of actions [2, 13, 14, 16, 28], but provide no mechanisms for reacting to changes in the environment. Reactive planners [5, 15, 17, 18] are designed to select and execute actions in response to the current state of the world, but, with a few exceptions [3], [4], they do not use probabilistic information to determine the likelihood of success of the actions. In [1], probabilistic information is used, in an effort to deal with environmental uncertainties, but no optimal action selection strategies are considered and/or proposed.

The ATIN formalism in this paper is similar to an earlier work by Sugato Baghci et al [1] on planning under uncertainty. The similarity, however, stops with the graph representation of preconditions, actions and their effects. Similar parallels can also be drawn with other graph-based planning approaches, e.g. GraphPlan (<http://www.cs.cmu.edu/~avrim/graphplan.html>). The approach in this paper represents a new formalism and is based on well established statistical results.

2 Problem formalization – core

In this section, we consider a modular core problem. We initially isolate a single action with its supporting preconditions and its resulting effects, as depicted in Fig. 2.




Figure 2: A Single Action ATIN

$X_1^a = [X_1, \dots, X_n]^T$ The status random vector of the preconditions, where $X_i = 1$, if precondition c_i is present and $X_i = 0$ if precondition c_i is absent. X_1^a denotes binary vector value realizations of X_1^a .

$Y_1^m = [Y_1, \dots, Y_m]^T$ The status random vector of the effects, where $Y_i = 1$, if effect e_i is present and $Y_i = 0$ if effect e_i is absent. Y_1^m denotes binary vector value realizations of Y_1^m .

$p_j(x_1^a)$ The probability of success for action a_j , given that the value of the precondition status vector is x_1^a ;

$P(\text{success for action } a_j | x_1^a)$

$q_j(y_1^m)$ The probability that the value of the effects' status vector is y_1^m , given that the action a_j is taken;

$P(y_1^m | a_j \text{ taken})$

$q_0(y_1^m)$ The probability that the value of the effects' status vector is y_1^m , given that no action is taken;

$P(y_1^m | \text{no action taken})$

$U_j(y_1^m)$ The utility of the value y_1^m of the effects' status vector, when action a_j is taken.

$U_0(y_1^m)$ The utility of the value y_1^m of the effects' status vector, when no action is taken.

We note that the utility function $U_j(y_1^m)$ measures the net worth of the effects' vector value y_1^m when action a_j is taken; thus, $U_j(y_1^m)$ is computed as the worth of y_1^m minus the cost of deployment for action a_j .

Let us now assume mutually exclusive actions, which are supported by the same preconditions, to lead to the same set of effects (as shown in Fig. 3). Let $\{a_j\}_{1 \leq j \leq k}$ be this set of actions and let X_1^a and Y_1^m denote the common status random vectors of preconditions versus effects, respectively. Let the utility functions for each action in the set $\{a_j\}_{1 \leq j \leq k}$ be nonnegative; let also $U_0(y_1^m)$ be nonnegative.




Figure 3: A Single Level ATIN

We now state multiple versions of the core problem, based on two different optimization criteria. Problem 3a and 3b are the constrained versions of the first two problems.

Problem 1 (Optimal Path Problem)

Given a preconditions vector value x_1^a , given an effects vector value y_1^m , find the maximum probability action that connects them. That is, find the action that maximizes the conditional probability $P(y_1^m | x_1^a)$.

Problem 2 (Average Utility Maximization)

Given a preconditions vector value x_1^a , find the action or actions that maximize the effects' average utility.

Problem 3a (Optimal Path Problem with Constrained Actions)

Given a preconditions vector value x_1^a , given an effects vector value y_1^m , and an action dependency matrix, find the maximum probability action that connects them. That is, find the action that maximizes the conditional probability $P(y_1^m | x_1^a)$. In this case, only those action combinations are considered that are allowed by the constraints in the dependency matrix.

Problem 3b (Average Utility Maximization with Constrained Actions)

Given a preconditions vector value x_1^a , find the action or actions that maximize the effects' average utility. As in Problem 3a, only those action combinations are considered that are allowed by the constraints in the dependency matrix.

Action Dependency Matrix (ADM)

An action dependency matrix is a tool which defines dependency among actions in the network. It reduces the number of combinations of actions by considering only those allowed by the dependency matrix. It also reduces significantly the amount of calculations required to obtain the optimal path. The value of the variable a_{ij} reflects the existence or absence of dependency between actions a_i and a_j , where a_{ij} equals 1; for positive dependency and equals 0; for negative dependency, and $1 \leq i, j \leq n$, where 'n' represents the total number of actions in the network. The elements of an ADM are determined as follows:

$$a_{ij} = \begin{cases} 1 & \text{; if action } a_i \text{ is selected for execution} \\ & \text{in level } l, \text{ then } a_j \text{ refers to the action} \\ & \text{that has to be or has been selected} \\ & \text{for execution in level } k, \text{ where } l \neq k \\ 0 & \text{; if action } a_i \text{ is selected for execution} \\ & \text{in level } l, \text{ then } a_j \text{ refers to the action} \\ & \text{that must not be selected for execution} \\ & \text{in level } k, \text{ where } l \neq k \end{cases}$$

where, level l in an ATIN corresponds to a set of preconditions (C_1, C_2, \dots, C_n) followed by a set of actions (a_1, a_2, \dots, a_k) and a set of effects (e_1, e_2, \dots, e_m) (as shown in Fig. 3). The effects of this level then serve as the preconditions for the next level $l+1$ and so on.

3 Solutions to the core problems

We present the solutions to the two core problems posed in Section 2 in the form of a theorem, whose proof is in the Appendix.

Theorem 1

- a. Given x_1^a , given y_1^m , and given a set of actions $\{a_j\}_{1 \leq j \leq k}$, the conditional probability $P(y_1^m | x_1^a)$ is maximized as follows:

by action a_j ; if

$$q_{j*}(y_1^m) p_j(x_1^a) = \max_{1 \leq j \leq k} q_j(y_1^m) p_j(x_1^a) > q_0(y_1^m) \quad (1)$$

where then $\max P(y_1^m | x_1^a) = q_{j*}(y_1^m) p_j(x_1^a)$

by no action; if

$$q_0(y_1^m) > \max_{1 \leq j \leq k} q_j(y_1^m) p_j(x_1^a) \quad (2)$$

where then $\max P(y_1^m | x_1^a) = q_0(y_1^m)$

If more than one action satisfy the maximum in (1), then one of these actions may be selected randomly.

- b. Given x_1^a , given a set of actions $\{a_j\}_{1 \leq j \leq k}$, and given utility functions $\{U_j(y_1^m)\}_{1 \leq j \leq k}$ and $U_0(y_1^m)$, the average utility

$$\begin{aligned} \bar{U}(x_1^a) &\triangleq \sum_{1 \leq j \leq k} \sum_{y_1^m} P(a_j \text{ taken}, y_1^m | x_1^a) \cdot U_j(y_1^m) \\ &+ \sum_{y_1^m} P(\text{no action taken} | x_1^a) \cdot U_0(y_1^m) \end{aligned}$$

is maximized as follows:

by action a_j ; if

$$\begin{aligned} A_{j*}(x_1^a) &\triangleq p_j(x_1^a) \sum_{y_1^m} q_{j*}(y_1^m) U_j(y_1^m) = \\ &\max_{1 \leq j \leq k} p_j(x_1^a) \sum_{y_1^m} q_j(y_1^m) U_j(y_1^m) > \\ &\sum_{y_1^m} q_0(y_1^m) U_0(y_1^m) \quad (3) \end{aligned}$$

by no action; if

$$\begin{aligned} \sum_{y_1^m} q_0(y_1^m) U_0(y_1^m) &> \\ &\max_{1 \leq j \leq k} p_j(x_1^a) \sum_{y_1^m} q_j(y_1^m) U_j(y_1^m) \quad (4) \end{aligned}$$

$A_{j*}(x_1^a)$ in (3) is the award assigned to action a_j ; it is also the worth assigned to the precondition vector value x_1^a by the action a_j .

If more than one action attain the maximum award $A_{j*}(x_1^a)$ in (3), one of them is selected randomly.

4 Solutions of the network propagation problem

In this section, we generalize the core problem solutions expressed in Theorem 1, Section 3, to the sequence of actions depicted by the ATIN in Fig. 1.

Problem 1 (The Optimal Path Problem)

In the ATIN in Fig. 1, we fix the preconditions vector value $x_1^a(1)$, at time 1, and the effects' vector value $y_1^m(N)$, at time N . We then search for the sequence of actions that maximizes the probability $P(y_1^m(N) | x_1^a(1))$. The solution to this problem follows a dynamic programming approach where $x_1^a(l) = y_1^m(l-1)$, $2 \leq l \leq N$, in our notation. The proof of the step evolution is included in the Appendix.

Step 1

For each $y_1^m(1) = x_1^a(2)$ value, find

$$r(y_1^m(1)) \triangleq \max_j [q_0(y_1^m(1)), \max_j q_j(y_1^m(1)) p_j(x_1^a(1))]$$

and the action index $j^*(y_1^m(1))$ that attains $r(y_1^m(1))$.

Step l

The values $r(y_1^m(l-1)) \triangleq \max_j P(y_1^m(l-1) | x_1^a(l))$, for each $y_1^m(l-1)$ value, are in memory, as well as the actions that attain them. At step l , the values

$$r(y_1^m(l)) \triangleq \max_{y_1^m(l-1)} r(y_1^m(l-1)) \times \max_j [q_0(y_1^m(l)), \max_j q_j(y_1^m(l)) p_j(y_1^m(l-1))]$$

are maintained, as well as the sequence of actions leading to them.

The complexity of this problem is polynomial with respect to the number of links. Assume that a given ATIN model has 'N' number of levels and each level has 'k' links, then the complexity is given as $O(N \times k)$.

Problem 2 (The Average Utility Maximization)

In the ATIN in Fig. 1, we fix the value of the precondition vector at time 1, denoted $x_1^a(1)$. For each value $y_1^m(N)$ of the effects vector at time N , we assign worth functions $U(y_1^m(N))$. For each action $a_j(l)$, at time l , we assign a deployment cost $c_j(l)$. The utility of the effects' vector value $y_1^m(N)$, when action a_j is taken, is then equal to $U_j(y_1^m(N)) \triangleq U(y_1^m(N)) - c_j(N)$, while the utility of the same value, when no action is taken, equals $U_0(y_1^m(N)) \triangleq U(y_1^m(N))$. We are seeking the sequence of actions which lead to the maximization of the average utility. The evolving algorithm, from part (b) of Theorem 1, back propagates as follows. The proof is in the Appendix.

Step 1

Compute the action awards (including that to no action), with notation of Figure 1, as follows: $0 \leq j \leq r$;

$$A_j(x_1^a(N-1)) \triangleq p_j(x_1^a(N-1)) \sum_{y_1^m(N)} q_j(y_1^m(N)) U_j(y_1^m(N))$$

with $p_0(x_1^a(N-1)) \triangleq 1$

Select $A_{j^*(x_1^a(N-1))}(x_1^a(N-1)) \triangleq \max_{0 \leq j \leq r} A_j(x_1^a(N-1))$;

for each $x_1^a(N-1)$ value.

Take action $a_{j^*(x_1^a(N-1))}(N)$ for preconditions vector value $x_1^a(N-1)$ and simultaneously assign worth $A_{j^*(x_1^a(N-1))}(x_1^a(N-1))$ to $x_1^a(N-1)$. That is, assign:

$$U(x_1^a(N-1)) = A_{j^*(x_1^a(N-1))}(x_1^a(N-1)) \quad (5)$$

Step 2

Back propagate to the preconditions at $N-2$, as in Step 1, starting with the worth assignments in (5), and subsequent utilizations

$$U_j(x_1^a(N-1)) = \max[A_{j^*(x_1^a(N-1))}(x_1^a(N-1)) - c_j(N-1), 0]$$

Step n

As in Steps 1 and 2 (for subsequent levels) the above described algorithm generates the optimal sequence of actions for given initial preconditions $x_1^a(1)$. The optimal such preconditions can be also found via maximization of the utility $U_j(x_1^a(2))$, with respect to $x_1^a(1)$.

The complexity of this problem is also polynomial with respect to the number of links.

Problems 3a, 3b (Optimization with Constrained Actions)

Problems 3a and 3b impose dependency constraints on the actions in the ATIN network. As explained in Section 2, an ADM defines the dependency of one action on every other one, where positive dependency is depicted by 1 and negative dependency is depicted by 0. The dependency constraints are taken into account, when, at a certain level, an optimal action is finalized. At any given level, only positively related actions are considered in the calculations.

As described in Step 1 of Problem 1 (see Section 4), for the first level, $r(y_1^m(1))$ is calculated the same way for constrained actions also. But for the rest of the levels, it is calculated in a different manner.

Consider,

$$r(y_1^m(l)) \triangleq \max_{y_1^m(l-1)} r(y_1^m(l-1)) \times \max_j [q_0(y_1^m(l)), \max_j q_j(y_1^m(l)) p_j(y_1^m(l-1))]$$

The parameter $\max_{y_1^m(l-1)} r(y_1^m(l-1))$ corresponds to an ac-

tion selected for execution in level $l-1$. Its dependent actions can be known from the ADM. In this way, those combinations of actions which are not allowed by the ADM are eliminated from the calculation of $r(y_1^m(l))$, hence eliminating all links to and from the actions exhibiting negative dependencies. As a result of which it yields a network with lesser number of links and eases the determination of optimal sequence of actions.

5 Numerical evaluations

In this section, we focus on numerical scenarios. We first state the experimental setup. We then, evaluate and discuss a specific experimental scenario. We only state the experimental setups for Problems 1 and 2, since those of Problems 3a and 3b are straight forward modifications of the former.

5.1 Experimental Setups

Experimental Setup for Problem 1

Assign the probabilities

$\{q_j(x_1^k(l))\}$ and $\{p_j(x_1^k(l))\}$ as in problem 2.

Given these probabilities:

a. Compute first:

$$r(y_1^m(1)) \stackrel{\Delta}{=} \max_j [q_0(y_1^m(1)), \max_j q_j(y_1^m(1)) p_j(x_1^a(1))] \\ \text{and the action } j^*(y_1^m(1)) \text{ that attains } r(y_1^m(1)).$$

b. For each $l: 2 \leq l \leq N$, maintain in memory the values

$$r(y_1^m(l-1)) \stackrel{\Delta}{=} \max P(y_1^m(l-1) | x_1^a(l)), \text{ for each } y_1^m(l-1) \text{ value, and the actions that attain them. Then, compute and maintain the values:}$$

$$r(y_1^m(l)) \stackrel{\Delta}{=} \max_{y_1^m(l-1)} r(y_1^m(l-1)) \times \\ \times \max [q_0(y_1^m(l)), \max_j q_j(y_1^m(l)) p_j(y_1^m(l-1))] \\ \text{Also, maintain the actions that attain the values } r(y_1^m(l)).$$

Experimental Setup for Problem 2

Considering the network in Fig. 1, assign:

a. Worth function $U(y_1^w(N))$ for all $y_1^w(N)$ values of the effects' status vector, at level N .

b. Probabilities $q_j(x_1^k(l)) \stackrel{\Delta}{=}$

$P(x_1^k(l) \text{ occurring} | \text{action } j \text{ at step } l-1) \text{ at all levels, } 2 \text{ to } N,$

where $q_0(x_1^k(l)) \stackrel{\Delta}{=}$

$P(x_1^k(l) \text{ occurring} | \text{no action } j \text{ at step } l-1) \text{ at all levels, } 2 \text{ to } N,$

c. Probabilities $p_j(x_1^k(l)) \stackrel{\Delta}{=}$

$P(\text{action } j \text{ succeeds} | x_1^k(l) \text{ preconditions}) \text{ at all levels, from } 1 \text{ to } N-1,$

where $p_0(x_1^k(l)) = 1; \forall l$

d. Implementation/deployment costs $c_j(l)$ for all actions, at all levels 2 to N .

Given the above assignments,

a. Compute first,

$$A_j(x_1^l(N-1)) \stackrel{\Delta}{=} \\ p_j(x_1^l(N-1)) \sum_{y_1^w(N)} q_j(y_1^w(N)) U_j(y_1^w(N))$$

where,

$$p_0(x_1^l(N-1)) \stackrel{\Delta}{=} 1;$$

$$U_j(y_1^w(N)) = \max [U(y_1^w(N)) - c_j(N), 0]$$

$$A_{j^*(x_1^l(N-1))}(x_1^l(N-1)) \stackrel{\Delta}{=} \max_{0 \leq j \leq n} A_j(x_1^l(N-1)) ; \\ \text{for all } x_1^l(N-1) \text{ values.}$$

b. Take action $a_{j^*(x_1^l(N-1))}$ for each precondition vector value $x_1^l(N-1)$.

Assign worth $A_{j^*(x_1^l(N-1))}(x_1^l(N-1))$ to $x_1^l(N-1)$, as

$$U(x_1^l(N-1)) = A_{j^*(x_1^l(N-1))}(x_1^l(N-1))$$

Repeat steps (a) and (b) for level $N-1$ and back propagate to level $N-2$. Continue back propagation to level 1.

5.2 A Specific Experimental Scenario

In this section, we illustrate the use of Activation Timed Influence Nets with the help of an example ATIN, and present the results of the algorithms included in this paper, when applied to this ATIN. The model used in this section was derived from a Timed Influence Net presented in Wagenhals et al., in 2001 [27] (which was developed with the help of a team of subject matter experts) to address the internal political instabilities in Indonesia in the context of East Timor. For purposes of results illustration, we have selected a part of this network, as shown in Fig. 4.

Example ATIN:

The model provides detailed information about the religious, ethnic, governmental and non-governmental organizations of Indonesia. In this section, the propositions and actions referred are given in *italic text*. According to the model, rebel militia formed by a minority group poses the main concern which has captured a large number of people under its secured territory. Amongst these people in the community, some are against the rebels and considered to be at risk, in case the negotiations with the local government didn't work. For this example,

consider the initial conditions when *the rebels are getting local support*, the *community is in unrest* and the *local administration is losing control*. Based on the data provided, only one action can be executed from a possible set of actions at a given time i.e. either of the *Indonesian press* or *provincial authority* or the *minister of interior* would declare resolve to keep peace. Depending upon this selected action and the data provided for the effects, only a specific set of events can result. For instance, rebels may or may not start thinking that they are getting publicity, GOI (original anti-government of Indonesia) war may or may not expand, GOI chances of intervention and international attention may increase or decrease. Similarly, this specific set of events forms the set of possible pre-conditions for a later time. Depending upon which conditions actually become true, second action can be selected for execution from another set of actions, i.e. *Security Council* and *General Assembly* may or may not pass resolutions or *UN* may or may not declare resolve to keep peace. Depending upon this action and the data provided for the effects, coalition may or may not form, rebels may or may not contemplate talks, GOI support may increase or decrease or may not increase at all, or GOI may or may not allow coalition into territories. Ultimately, the coalition may authorize use of force which might compel rebels to negotiate and the humanitarian assistance (HA) may start preparing for the worst case. Depending upon which conditions meet, the coalition may declare resolve to keep peace or may declare war on rebels. This may affect the chances of military confrontation, rebels' popularity and chances of negotiated settlement which represents the final effects in the network.

Table 1 lists some of the parameters (and their values) required by the network in Fig. 4. The parameters in the table are listed by their abbreviated labels also in addition to the phrases shown inside the network nodes in the figure. For the sake of brevity, we do not list all the values.

Solutions to Problems:

Solution to Problem 1 (Optimal Path Problem):

Consider the example scenario described earlier, we need

to identify an optimal path (i.e., the sequence of actions) resulting into the final effect when, *military confrontation chances are reduced*, while *rebels start losing local support* and *negotiation chances start increasing*. This set of effects (post-conditions) leads to the following output state in the ATIN model:

- Reduction in the chances of military confrontation (i.e. $Y_{12} = 0$)
- Decrease in local support and popularity for Rebels (i.e. $Y_{13} = 1$)
- Increase in chances of negotiated settlement (i.e. $Y_{14} = 1$).

The above defined conditions lead to a post-condition vector $[0, 1, 1]^T$ at level 4, i.e. $y_{12}^{14}(4)$.

After fixing the post-condition vector, we define the initial preconditions, when *rebels have been getting local support*, the *community has been in unrest* and the *local administration has started losing control*. This set of pre conditions given by $x_1^3(1)$ results into a vector value of $[1, 1, 1]^T$, where

- $X_1 = 1$; represents the condition *Rebels are getting Local Support*
- $X_2 = 1$; represents the condition *There is unrest in the Community*
- $X_3 = 1$; represents the condition *Local Administration is losing Local Control*.

We want to find out the sequence of actions which achieves the desired effects $y_{12}^{14}(4)$ given the initial preconditions $x_1^3(1)$. Technically, we want to identify the sequence of actions which maximizes the probability $P(y_{12}^{14}(4) | x_1^3(1))$. Applying the optimal path algorithm (see Section 4) results that if the *provincial authority* and *UN declare resolve to keep peace* and *coalition does not take any action*, instead it declares resolve to keep peace, then the desired effects will be achieved which will result into *less chances of military confrontation*, *reduction in local support for rebels* and *more chances of a negotiated settlement*.

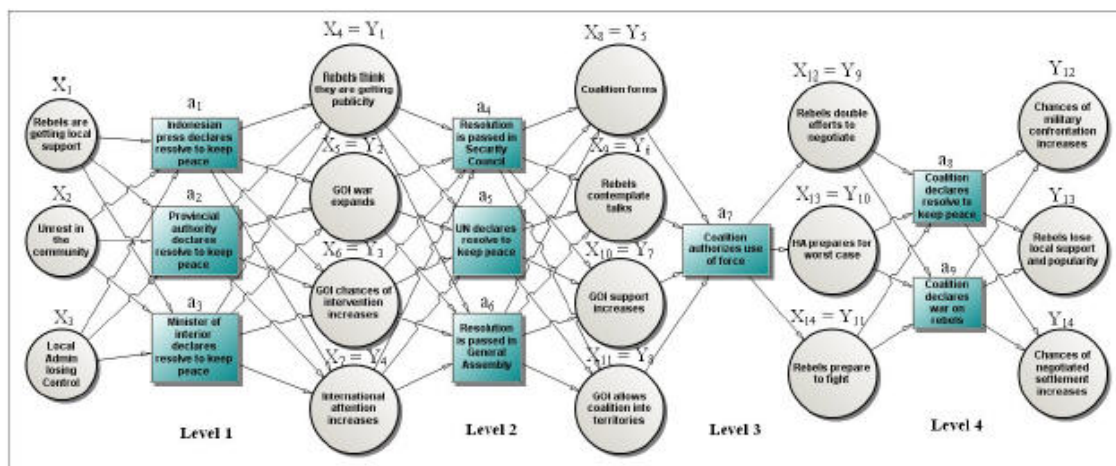


Figure 4: Example ATIN.

Table 1: Parameter values in the Example ATIN

Level 1							
Y_1^4	Action a_j	$q_j(y_1^4)$	$q_0(y_1^4)$	X_1^3	$p_j(x_1^3)$	$p_j(x_1^3) * q_j(y_1^4)$	$r(y_1^4)$
$[1,0,1,1]^T$	Indonesian press declares resolve to keep peace (a1)	68.00%	6.45%	$[1,1,1]^T$	80.00%	54.40%	63.64%
	Provincial Authority declares resolve to keep peace (a2)	74.00%			86.00%	63.64%	
	Minister of interior declares resolve to keep peace (a3)	56.00%			20.00%	11.20%	
Level 2							
Y_4^8	Action a_j	$q_j(y_4^8)$	$q_0(y_4^8)$	X_4^7	$p_j(x_4^7)$	$p_j(x_4^7) * q_j(y_4^8)$	$r(y_4^8)$
$[1,1,1,1]^T$	Resolution is passed in Security council (a4)	14.00%	0.95%	$[0,0,0,0]^T$	16.00%	2.24%	34.02%
				
				$[1,0,1,1]^T$	91.00%	12.74%	
				
				$[1,1,1,1]^T$	18.00%	2.52%	
	UN declares resolve to keep peace (a5)	42.00%		$[0,0,0,0]^T$	66.00%	27.72%	
				
				$[1,0,1,1]^T$	81.00%	34.02%	
				
				$[1,1,1,1]^T$	13.41%	5.63%	
	Resolution is passed in General Assembly (a6)	39.00%		$[0,0,0,0]^T$	43.05%	16.79%	
				
				$[1,0,1,1]^T$	48.20%	18.80%	
				
				$[1,1,1,1]^T$	24.87%	9.70%	
Level 3							
Y_9^{11}	Action a_j	$q_j(y_9^{11})$	$q_0(y_9^{11})$	X_9^{10}	$p_j(x_9^{10})$	$p_j(x_9^{10}) * q_j(y_9^{11})$	$r(y_9^{11})$
$[1,0,0]^T$	Coalition Authorizes the Use of Force (a7)	43.00%	59.00% (No Action)	$[0,0,0,0]^T$	0.00%	0.00%	59.00%
				
				$[1,0,1,1]^T$	64.00%	27.52%	
				
$[1,1,1,1]^T$	18.00%	7.74%					
Level 4							
Y_{12}^{14}	Action a_j	$q_j(y_{12}^{14})$	$q_0(y_{12}^{14})$	X_{12}^{13}	$p_j(x_{12}^{13})$	$p_j(x_{12}^{13}) * q_j(y_{12}^{14})$	$r(y_{12}^{14})$
$[0,1,1]^T$	Coalition declares resolve to keep peace (a8)	21.00%	1.50%	$[0,0,0]^T$	16.00%	3.36%	19.11%
				
				$[1,0,0]^T$	91.00%	19.11%	
				
				$[1,1,1]^T$	38.00%	7.98%	
	Coalition declares war on rebels (a9)	17.00%		$[0,0,0]^T$	67.00%	11.39%	
				
				$[1,0,0]^T$	18.48%	3.14%	
				
				$[1,1,1]^T$	30.88%	5.25%	

The details of this result are given in Table 1. It only contains the values that correspond to the selected actions at their respective levels, while a complete set of probabilities has been used to calculate the actual final sequence. The optimal actions, their corresponding state vectors and the probabilities are underlined in the table. The Optimal Path algorithm is of dynamic programming nature, so it requires two traversals to finalize the sequence of actions. During the forward traversal, $r(y_1^*)$ is calculated for each level for all possible post-condition combinations. At the last level, the post-condition vector $y_{12}^{14}(4)$ is fixed to be the desired effect of the network which is $[0, 1, 1]^T$ as determined earlier. The best action associated with this post-condition vector is identified along with its pre-condition vector $x_{12}^{14}(4)$. Using this pre-condition vector (which is the post-condition vector of the second last level), the network is traversed in reverse direction identifying actions and their corresponding preconditions, from last to the first level. The action at the first level is identified by fixing the pre-condition to the value determined earlier, i.e. $x_1^3(1)$ which is $[1, 1, 1]^T$. Completing both forward and reverse traversals gives the optimal actions which achieve the desired effects when the initial causes are given.

Solution to Problem 2 (Average Utility Maximization):

Consider a scenario where we need to identify the sequence of actions which maximizes the effects' average utility (at level 4) for the same input pre-condition as it was used in the solution of Problem 1, i.e. $[1, 1, 1]^T$. Assume, that the deployment costs for actions a_8 and a_9 are 25 and 30 units, respectively. The worth of each effect in the last level (i.e. level 4) is given by the worth function values $U(y_{12}^{14}(4))$ given in Table 2 and 3. Each effect also has a net utility which is determined by subtracting the deployment cost of the action from the worth of the effect. This net utility $U_j(y_{12}^{14}(4))$ (when action a_j is taken) and the action awards are given in Tables 2 and 3. The action award is calculated for each action corresponding to all of its pre-conditions. Similarly, these calculations are performed for the rest of the actions in ATIN model (after costs are assigned to every action in the model), but for the sake of brevity only the results for actions a_8 and a_9 are shown in Tables 2 and 3, respectively.

As described in Section 4, the action award is calculated for all actions in each level. For instance, starting from the last level, the action awards are calculated for actions a_8 and a_9 . The selected action is the one which maximizes the average utility and its action index 'j' is recorded. As each action award is calculated, it is also assigned as the worth function to the previous level effects vector. The latter worth function is used to calculate the utilities at the previous level, and calculations are repeated similarly. This procedure is back traversed from last to first levels. Table 4 summarizes the action awards of those actions which maximize the effects' average utility at their respective levels.

Table 2: Utility Functions and Action awards for Action a_8

Level 4 – Action a_8						
$X_{12}^{14}(4)$	$p_i(x_{12}^{14}(4))$	$Y_{12}^{14}(4)$	$q_j(y_{12}^{14}(4))$	$U(y_{12}^{14}(4))$	$U_j(y_{12}^{14}(4))$	$A_8(x_{12}^{14}(4))$
$[0,0,0]^T$	16.00%	$[0,0,0]^T$	37.00%	40	15	11.11
$[0,0,1]^T$	24.00%	$[0,0,1]^T$	65.00%	30	5	16.66
$[0,1,0]^T$	75.00%	$[0,1,0]^T$	53.00%	60	35	52.07
$[0,1,1]^T$	85.00%	$[0,1,1]^T$	21.00%	79	54	59.02
$[1,0,0]^T$	91.00%	$[1,0,0]^T$	19.00%	41	16	11.11
$[1,0,1]^T$	72.00%	$[1,0,1]^T$	43.00%	65	40	49.99
$[1,1,0]^T$	16.00%	$[1,1,0]^T$	29.00%	37	12	63.18
$[1,1,1]^T$	38.00%	$[1,1,1]^T$	27.00%	51	26	26.38

Table 3: Utility Functions and Action awards for Action a_9

Level 4 – Action a_9						
$X_{12}^{14}(4)$	$p_i(x_{12}^{14}(4))$	$Y_{12}^{14}(4)$	$q_j(y_{12}^{14}(4))$	$U(y_{12}^{14}(4))$	$U_j(y_{12}^{14}(4))$	$A_9(x_{12}^{14}(4))$
$[0,0,0]^T$	67.00%	$[0,0,0]^T$	41.00%	40	10	48.25
$[0,0,1]^T$	97.15%	$[0,0,1]^T$	26.00%	<u>30</u>	<u>0</u>	<u>69.96</u>
$[0,1,0]^T$	58.29%	$[0,1,0]^T$	71.00%	60	30	41.97
$[0,1,1]^T$	13.00%	$[0,1,1]^T$	17.00%	79	49	9.36
$[1,0,0]^T$	18.48%	$[1,0,0]^T$	26.00%	41	11	13.31
$[1,0,1]^T$	39.28%	$[1,0,1]^T$	54.00%	65	35	28.29
$[1,1,0]^T$	38.67%	$[1,1,0]^T$	62.00%	37	7	27.85
$[1,1,1]^T$	30.88%	$[1,1,1]^T$	58.00%	51	21	22.24

From Table 4 it can be seen that the sequence of actions that maximizes the effects' average utility, obtained as a result of applying the algorithm is given by: a_1 (i.e. *Indonesian press declares resolve to keep peace*), a_6 (i.e. *Resolution is passed in General Assembly*), a_7 (i.e. *Coalition authorizes use of Force*), a_9 (i.e., *Coalition declares war on rebels*). The underlined entries in Table 3 correspond to the worth, utility function and action award of action a_9 .

Solution to Problem 3a, 3b (Constrained Actions):

The dependencies among the actions in the example ATIN model are defined in the action dependency matrix given in Figure 5.

Most of the dependencies given in the matrix are quite evident. For instance, the *peace resolution declaration by UN* (a_5) ensures that either of *Indonesian press*, *provincial authority* or *minister of interior* must also have declared the *resolution to keep peace* (either of a_1 or a_2 or a_3 must have been executed in the past) which would represent the opinion of the locals in general. Similarly, *resolution passed by the Security Council or General Assembly* (a_4 or a_6) makes sure that *whether or not the coalition will have to authorize the use of force* (a_7), considering the resolution is in support of use of force. This infers that if the coalition authorizes the use of force, it will declare war on Rebels otherwise, it will declare resolve to keep peace. All of these dependencies can be observed from the ADM (as shown in Fig. 5).

Consider a_{25} in ADM, (as shown in Fig. 5) which

corresponds to a positive dependency between *peace declaration by the provincial authority* (a_2) and *peace declaration by UN* (a_5). The ADM suggests that there exist negative dependencies between action a_2 and actions a_4 , a_6 , a_7 and a_9 which means that if *Provincial authority declares peace resolution*, *Security Council* and *General Assembly* won't pass resolution and the *Coalition* will not authorize the use of force and hence will declare resolve to keep peace. This knowledge of dependencies from the ADM certainly reduces an extensive amount of effort in calculating the optimal path. While calculating the optimal path, during the forward traversal, only those paths are considered which satisfy the constraints defined in ADM yielding less number of combinations to consider for calculation and making it easy to back traverse and identify the optimal actions.

The same applies to the solution of the second problem of identifying sequence of actions maximizing the effects' average utility under constraints. The action awards are calculated for those actions only which satisfy constraints defined in ADM, and hence reducing the effort of calculating action awards and assignment of worth function at each level.

Table 4: Action Awards.

Level 1	Level 2	Level 3	Level 4
$A_1(x_1^3(1))$	$A_2(x_2^7(2))$	$A_7(x_7^{11}(3))$	$A_9(x_9^{14}(4))$
151.02	85.18	77.92	69.96

6 Conclusion

This paper presented an extension of a Timed Influence Net, termed ATIN (Activation Timed Influence Net). An ATIN utilizes a set of preconditions required for the undertaking of an action and produces a set of effects. These effects become then the preconditions for the next level of action(s), resulting in a sequential evolution of actions. Some other probabilistic planning techniques were also discussed. The paper identified several pre-

selected performance criteria regarding ATINs (i.e., optimal path and average utility maximization with and without constrained actions) and recommended algorithms for their satisfaction. A tool called ADM (Action Dependency Matrix) was introduced, which induces dependencies among the actions. It is represented with the help of a $m \times m$ matrix, where ' m ' represents the total number of actions in the network.

The implementation of the suggested algorithms was illustrated with the help of a real world example. The example demonstrated a politically unstable situation in Indonesia. Sets of actions preceded by preconditions and followed by sets of effects were demonstrated in the form of an ATIN Model (see Figure 4). The experiment was formulated based on a previous Timed Influence Network model for the same scenario. The experimental procedure was applied to the network with a set of probability data. Solutions of both problems were discussed in depth. The optimal path problem required the knowledge of an initial set of causes (preconditions) and the final set of effects (postconditions). With the help of the algorithm, an optimal sequence of actions was identified which maximized the conditional probability of achieving the desired effects, when the initial conditions were given. For the sake of brevity, only significant parts of the probability data used were shown in Table 1. For the same scenario, the second algorithm yielded a sequence of actions, which maximized the effects' average utility. The solution for both problems was comprehended in detail. The experiment was repeated with constrained actions considering only dependent actions as defined in the Action Dependency Matrix (see Figure 5) which produced similar results and required lesser effort to calculate than without ADM.

References

- [1] Bagchi S., Biswas G., and Kawamura K., (Nov, 2000), "Task Planning under Uncertainty using a Spreading Activation Network", IEEE Transactions

	a1 Indonesian Press declares resolve to keep peace	a2 Provincial Authority Press declares resolve to keep peace	a3 Minister of Interior Press declares resolve to keep peace	a4 Resolution is passed in Security council	a5 UN Declares resolve to keep peace	a6 Resolution is passed in General Assembly	a7 Coalition authorizes use of Force	a8 Coalition declares resolve to keep peace	a9 Coalition declares war on Rebels
a1 Indonesian Press declares resolve to keep peace	1	0	0	0	1	0	0	1	0
a2 Provincial Authority Press declares resolve to keep peace	0	1	0	0	1	0	0	1	0
a3 Minister of Interior Press declares resolve to keep peace	0	0	1	0	1	0	0	1	0
a4 Resolution is passed in Security council	0	0	0	1	0	0	1	0	1
a5 UN Declares resolve to keep peace	1	1	1	0	1	0	0	1	0
a6 Resolution is passed in General Assembly	0	0	0	0	0	1	1	0	1
a7 Coalition authorizes use of Force	0	0	0	1	0	1	1	0	1
a8 Coalition declares resolve to keep peace	1	1	1	0	1	0	0	1	0
a9 Coalition declares war on Rebels	0	0	0	1	0	1	1	0	1

Figure 5: Action Dependency Matrix

- on Systems, Man and Cybernetics, Vol. 30, No. 6, pp 639–650.
- [2] Dean T. and Wellman (1991) M.P., Planning and Control. San Mateo, CA: Morgan Kaufmann.
 - [3] Dean T., Kaelbling L.P., Kirman J., and Nicholson A., (1993), "Planning with deadlines in stochastic domains," in Proc. 11th Nat. Conf. Artificial Intelligence (AAAI-93). Chicago, IL: AAAI Press, pp. 574–579.
 - [4] Drummond M. and Bresina J., (1990), "Anytime synthetic projection: Maximizing the probability of goal satisfaction," in Proc. 8th Nat. Conf. Artificial Intelligence (AAAI-90). Chicago, IL: AAAI Press, pp.138–144.
 - [5] Firby R.J., (1987), "An investigation into reactive planning in complex domains," in Proc. 6th Nat. Conf. Artificial Intelligence (AAAI-87). San Mateo, CA: Morgan Kaufmann, vol. 1, pp. 202–206.
 - [6] Haider S., (October, 2003), "On Computing Marginal Probability Intervals in Probabilistic Inference Networks," in Proceedings of IEEE-SMC Conference.
 - [7] Haider S. (August, 2003), "A Hybrid Approach for Learning Parameters of Probabilistic Networks from Incomplete Databases," Design and Application of Hybrid Intelligent Systems, Proceedings of the Third International Conference on Hybrid Intelligent Systems (HIS'03), IOS Press, Amsterdam, The Netherlands, ISBN 158603-394.
 - [8] Haider S., Zaidi A. K., and Levis A. H. (Nov. 2004), "A Heuristic Approach for Best Set of Actions Determination in Influence Nets," in Proc. IEEE International Conference on Information Reuse and Integration, Las Vegas.
 - [9] Haider S. and Levis A. H. (2004), "An Approximation Technique for Belief Revision in Timed Influence Nets," in Proceedings of Command and Control Research and Technology Symposium.
 - [10] Haider S. and Zaidi A. K. (2004), "Transforming Timed Influence Nets into Time Sliced Bayesian Networks," in Proceedings of Command and Control Research and Technology Symposium.
 - [11] Haider S. and Zaidi A. K. (2005), "On Temporal Analysis of Timed Influence Nets using Point Graphs," in the Proc. of the 18th International FLAIRS Conference, FL.
 - [12] Haider S. and Levis A. H. (June 2005), "Dynamic Influence Nets: An Extension of Timed Influence Nets for Modeling Dynamic Uncertain Situations," in Proc. 10th International Command and Control Research and Technology Symposium, Washington DC.
 - [13] Kushmerick N., Hanks S., and Weld D., (1995), "An algorithm for probabilistic planning," Univ. of Washington, Seattle, Artificial Intelligence Journal, pp. 76(1-2):239-86
 - [14] Likhachev M. and Stentz A., (2006), "PPCP: Efficient Probabilistic Planning with Clear References in Partially Known Environments", Twenty-First AAAI Conference on Artificial Intelligence, Boston, Massachusetts, The AAAI Press/Menlo Park, California, pp 860-868.
 - [15] Maes P., (1990), "Situated agents can have goals," Robot. Auton. Syst., vol. 6, pp. 49-70.
 - [16] Mausam and Weld D.S., "Probabilistic Temporal Planning with Uncertain Durations", Twenty-First AAAI Conference on Artificial Intelligence, Boston, Massachusetts, The AAAI Press/Menlo Park, California, pp 880-888.
 - [17] Mitchell T.M., (1990), "Becoming increasingly reactive," in Proc. 8th Nat. Conf. Artificial Intelligence (AAAI-90). Chicago, IL: AAAI Press/The MIT Press, pp. 1051–1058.
 - [18] Nikolova E. and Karger D.R., "Route Planning under Uncertainty: The Canadian Traveler Problem", Twenty-Third AAAI Conference on Artificial Intelligence, Chicago, Illinois, The AAAI Press/Menlo Park, California, pp 969-975.
 - [19] Tate A., Hendler J., and Drummond M., (1990), "A review of AI planning techniques," in Readings in Planning, J. Allen, J. Hendler, and A. Tate, Eds. San Mateo, CA: Morgan Kaufmann, pp. 26–49.
 - [20] Wagenhals L. W. and Levis A. H. (June 2002), "Modeling Support of Effects-Based Operations in War Games," in Proc. 2002 Command and Control Research and Technology Symposium, Monterey, CA.
 - [21] Wagenhals L.W., Levis A. H., and McCrabb M. (June 2003), "Effects Based Operations: a Historical Perspective for a Way Ahead," in Proc. 8th Int'l Command and Control Research and Technology Symposium, National Defense University, Washington, DC.
 - [22] Wagenhals L. W. and Wentz L. K., (June 2003), "New Effects-Based Operations Models in War Games," in Proceedings of the 2003 International Command and Control Research and Technology Symposium, National Defense University, Washington, DC.
 - [23] Wagenhals L. W. (2000), "Course of Action Development and Evaluation Using Discrete Event System Models of Influence Nets", PhD Dissertation, George Mason University.
 - [24] Wagenhals L. W., Levis A. H. (2000), "Course of Action Development and Evaluation," in Proceedings of the 2000 Command and Control Research and Technology Symposium.
 - [25] Wagenhals L. W. and Levis A. H., "Modeling Effects Based Operations in Support of War Games," in Proc. Of the 15th International Symposium on Aerospace / Defense Sensing, Internal Society for Optical Engineering, Proceedings of SPIE, Vol. # 4367, 2001.
 - [26] Wagenhals L. W., Shin I., and Levis A. H., (1998) "Creating Executable Models of Influence Nets with Coloured Petri Nets," Int. J. STTT, Springer-Verlag, Vol. 1998, No. 2.
 - [27] Wagenhals L. W., Reid T. J., Smillie R. J., & Levis A. H. (June 2001), "Course of Action Analysis for Coalition Operations," Proceedings of 6th Interna-

- tional Command and Control Research and Technology Symposium, Annapolis, Maryland.
- [28] Wellman M. P., (1990), "Formulation of tradeoffs in planning under uncertainty". San Mateo, CA: Morgan Kaufmann.
- [29] Wilkins D.E., (1988), "Practical Planning: Extending the Classical AI Planning Paradigm". San Mateo, CA: Morgan Kaufmann.
- [30] Zaidi A. K., Wagenhals L. W., and Haider S. (2005), "Assessment of Effects Based Operations Using Temporal Logic," in Proc. of the 10th International Command and Control Research and Technology Symposium.
- [31] Zaidi A. K., Mansoor F., and Papantoni-Kazakos P. (October 2007), "Modeling with Influence Networks Using Influence Constants: A New Approach," in Proceedings of IEEE-SMC Conference.

Appendix

Proof of Theorem 1

In the derivations below, the following considerations are incorporated:

1. Effects are fully dictated by the actions taken; thus, when probabilities are conditioned on actions and preconditions, the conditioning on preconditions drops.
2. By probability of action success, we mean the probability that the action may succeed, given the preconditions. The final action is selected among those that have positive probability of success. The probability of action taken, given that the action may succeed is the criterion that dictates the final action selection.

$$(a) \quad P(y_1^m | x_1^n) = \sum_{1 \leq j \leq k} P(y_1^m, a_j \text{ taken} | x_1^n) + P(y_1^m, \text{no action taken} | x_1^n) \quad (1.1)$$

where

$$\begin{aligned} P(y_1^m, a_j \text{ taken} | x_1^n) &= \\ &= P(y_1^m | a_j \text{ taken}, x_1^n) P(a_j \text{ taken} | x_1^n) = \\ &= P(y_1^m | a_j \text{ taken}) P(a_j \text{ taken} | x_1^n) = \\ &= P(a_j \text{ taken} | x_1^n) q_j(y_1^m) \\ &= [P(a_j \text{ taken}, \text{succ for action } a_j | x_1^n) + \\ &\quad + P(a_j \text{ taken}, \text{no succ for action } a_j | x_1^n)] \times \\ &\quad \times q_j(y_1^m) \\ &= [P(a_j \text{ taken} | \text{succ for action } a_j, x_1^n) \times \\ &\quad \times P(\text{succ for action } a_j | x_1^n) + \\ &\quad + P(a_j \text{ taken} | \text{no succ for action } a_j, x_1^n) \times \\ &\quad \times P(\text{no succ for action } a_j | x_1^n)] q_j(y_1^m) \end{aligned}$$

$$\begin{aligned} &= [P(a_j \text{ taken} | \text{succ for action } a_j) \times \\ &\quad \times P(\text{succ for action } a_j | x_1^n) + \\ &\quad + P(a_j \text{ taken} | \text{no succ for action } a_j) \times \\ &\quad \times P(\text{no succ for action } a_j | x_1^n)] q_j(y_1^m) \\ &= P(a_j \text{ taken} | \text{succ for action } a_j) \times \\ &\quad \times p_j(x_1^n) q_j(y_1^m) \\ &\quad \{ \text{Using } P(a_j \text{ taken} | \text{no succ for action } a_j) = 0 \} \end{aligned}$$

Equating in (1.1)

$$\begin{aligned} P(y_1^m | x_1^n) &= \\ &= \sum_{1 \leq j \leq k} P(a_j \text{ taken} | \text{succ for action } a_j) \times \\ &\quad \times p_j(x_1^n) q_j(y_1^m) + P(y_1^m, \text{no action taken} | x_1^n) \quad (1.2) \end{aligned}$$

where

$$\begin{aligned} P(y_1^m, \text{no action taken} | x_1^n) &= \\ &= P(y_1^m | \text{no action taken}, x_1^n) P(\text{no action taken} | x_1^n) \\ &= P(y_1^m | \text{no action taken}) P(\text{no action taken} | x_1^n) \\ &= P(\text{no action taken} | x_1^n) q_0(y_1^m) \\ &= P(\text{no action taken}, \text{no action succ} | x_1^n) q_0(y_1^m) \\ &= P(\text{no action taken} | \text{no action succ}) q_0(y_1^m) \\ &\quad \{ \text{Using } P(\text{no action succ} | x_1^n) = 1 \} \end{aligned}$$

Equating in (1.2)

$$\begin{aligned} P(y_1^m | x_1^n) &= \\ &= \sum_{1 \leq j \leq k} P(a_j \text{ taken} | \text{succ for action } a_j) p_j(x_1^n) q_j(y_1^m) + \\ &\quad + P(\text{no action taken} | \text{no action succ}) q_0(y_1^m) \end{aligned}$$

$\Rightarrow \max P(y_1^m | x_1^n)$ attained
 if $P(a_j, \text{taken} | \text{succ for action } a_j) = 1$;
 for $p_{j^*}(x_1^n)q_{j^*}(y_1^m) = \max_{1 \leq j \leq k} p_j(x_1^n)q_j(y_1^m) > q_0(y_1^m)$
 otherwise, $\max P(y_1^m | x_1^n)$ attained
 if $P(\text{no action taken} | \text{no action succ}) = 1$;
 for $\max_{1 \leq j \leq k} p_j(x_1^n)q_j(y_1^m) < q_0(y_1^m)$

(b)

$$\begin{aligned} \bar{U}(x_1^n) &= \sum_{1 \leq j \leq k} \sum_{y_1^m} P(a_j, \text{taken}, y_1^m | x_1^n) U_j(y_1^m) + \\ &+ \sum_{y_1^m} P(\text{no action taken}, y_1^m | x_1^n) U_0(y_1^m) \\ &= \sum_{1 \leq j \leq k} \sum_{y_1^m} P(y_1^m | a_j, \text{taken}) P(a_j, \text{taken} | x_1^n) U_j(y_1^m) + \\ &+ \sum_{y_1^m} P(y_1^m | \text{no action taken}) P(\text{no action taken} | x_1^n) \\ &\times U_0(y_1^m) \\ &= \sum_{1 \leq j \leq k} P(a_j, \text{taken} | \text{succ for action } a_j) p_j(x_1^n) \times \\ &\times q_j(y_1^m) U_j(y_1^m) + P(\text{no action taken} | \text{no action succ}) \times \\ &\times \sum_{y_1^m} q_0(y_1^m) U_0(y_1^m) \end{aligned}$$

$\Rightarrow \max \bar{U}(x_1^n)$ attained for :
 $P(a_j, \text{taken} | \text{succ for action } a_j) = 1$;
 if $p_{j^*}(x_1^n) \sum_{y_1^m} q_{j^*}(y_1^m) U_{j^*}(y_1^m) =$
 $= \max_{1 \leq j \leq k} p_j(x_1^n) \sum_{y_1^m} q_j(y_1^m) U_j(y_1^m) > \sum_{y_1^m} q_0(y_1^m) U_0(y_1^m)$
 $P(\text{no action taken} | \text{no action succ}) = 1$;
 if $\sum_{y_1^m} q_0(y_1^m) U_0(y_1^m) > \max_{1 \leq j \leq k} p_j(x_1^n) \sum_{y_1^m} q_j(y_1^m) U_j(y_1^m)$

Proof of the Network Propagation – Problem 1

Using the notation in Section 4, Problem 1, and via the Theorem of Total Probability and the Bayes Rule, we obtain:

$$\begin{aligned} r(y_1^m(N)) &\stackrel{\Delta}{=} \max_{\text{sequence of actions}} P(y_1^m(N) | x_1^n(1)) = \\ &= \max_{\text{sequence of actions}} \sum_{y_1^{m(N-1)}} P(y_1^m(N), y_1^{m(N-1)} | x_1^n(1)) = \\ &= \max_{\text{sequence of actions}} \sum_{y_1^{m(N-1)}} P(y_1^m(N) | y_1^{m(N-1)}) \times \\ &\times P(y_1^{m(N-1)} | x_1^n(1)) \leq \max_{y_1^{m(N-1)}} \left[\max_{\text{action}} P(y_1^m(N) | \right. \end{aligned}$$

$$\begin{aligned} &\left. y_1^{m(N-1)} \right) \max_{\text{sequence of actions}} P(y_1^m(N-1) | x_1^n(1)) \Big] \\ &= \max_{y_1^{m(N-1)}} \left[\left\{ \max_{\text{action}} P(y_1^m(N) | y_1^{m(N-1)}) \right\} r(y_1^{m(N-1)}) \right] \\ &\max_{\text{action}} P(y_1^m(N) | y_1^{m(N-1)}) = \\ &= \max_j \left[\max_j p_j(y_1^m(N-1)) q_j(y_1^m(N)), q_0(y_1^m(N)) \right] \end{aligned}$$

Thus, via substitution principle, we obtain:

$$\begin{aligned} r(y_1^m(N)) &\leq \max_{y_1^{m(N-1)}} \left[r(y_1^{m(N-1)}) \times \right. \\ &\times \max \left\{ \max_j p_j(y_1^m(N-1)) \times q_j(y_1^m(N)), q_0(y_1^m(N)) \right\} \Big] \end{aligned}$$

with equality iff the $y_1^{m(N-1)}$ value that attains $r(y_1^{m(N-1)})$ is selected. The above proves the general step in the network propagation of Problem 1.

Proof of the Network Propagation – Problem 2

Using the notation in Section 4, Problem 2, and via the use of the Theorem of Total Probability and the Bayes Rule, we obtain:

$$\begin{aligned} &\max_{\text{sequence of actions } y_1^{m(N)}} \sum_{y_1^{m(N)}} U(y_1^m(N)) P(y_1^m(N) | x_1^n(1)) = \\ &= \max_{\text{sequence of actions } y_1^{m(N)}} \sum_{y_1^{m(N)}} U(y_1^m(N)) \sum_{x_1^i(N-1)} P(y_1^m(N), \\ &x_1^i(N-1) | x_1^n(1)) = \\ &= \max_{\text{sequence of actions } y_1^{m(N)}} \sum_{y_1^{m(N)}} U(y_1^m(N)) \sum_{x_1^i(N-1)} P(y_1^m(N) | \\ &x_1^i(N-1)) P(x_1^i(N-1) | x_1^n(1)) = \\ &= \max_{\text{sequence of actions } x_1^i(N-1)} \sum_{x_1^i(N-1)} P(x_1^i(N-1) | x_1^n(1)) \times \\ &\times \sum_{y_1^{m(N)}} U(y_1^m(N)) P(y_1^m(N) | x_1^i(N-1)) = \\ &= \max_{\substack{x_1^i(k); \\ k=1, \dots, N-2}} \sum_{x_1^i(N-1)} P(x_1^i(N-1) | x_1^n(1)) \times \\ &\times A_{r(x_1^i(N-1))} x_1^i(N-1) \end{aligned}$$

The latter expression proves the back propagation property and the steps in the algorithm.