REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*

2. REPORT DATE 2/28/98

3. REPORT TYPE AND DATES COVERED Final Tech Rpt 3/6/97 - 2/28/98

4. TITLE AND SUBTITLE

Computational Models of Human Organization Dynamics
Final Technical Report

5. FUNDING NUMBERS

A1 9760400.1320 E495
P6S10 2525 DPAC 6 5219
503733

6. AUTHOR(S)

Dr. Gregg Courand, Dr. Michael Fehling

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Synergia LLC
2400 Broadway, Ste 203
Redwood City, CA 94063-1551

8. PERFORMING ORGANIZATION REPORT NUMBER

COD - 4

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

DARPA / ISO
3701 N. Fairfax Dr.
Arlington, VA 22203-1714

10. SPONSORING / MONITORING AGENCY REPORT NUMBER

19980324 049

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION / AVAILABILITY STATEMENT

Approved for public release; distribution is unlimited.

12b. DISTRIBUTION CODE

A

13. ABSTRACT *(Maximum 200 words)*

We describe a prototype environment for computational modeling of social systems, ACCORD. We also present the theoretical foundation, Generative Practice Theory, and the methodology we have developed, Practice Mapping, to acquire the necessary information about the social system (e.g., an organization), develop its computational realization, validate the models via formal experimentation, and develop implications of the models. We have tested our research in several crisis response domains.

14. SUBJECT TERMS

Organization Dynamics, Computational Modeling, Generative Practice Theory, ACCORD

15. NUMBER OF PAGES

53

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18 298-102

USAPPC V1.00

Synergia LLC

2400 Broadway, Suite 203
Redwood City, CA 94063-1551
Phone: (650) 569-4999
Fax: (650) 569-4990
World Wide Web: www.synergia.com

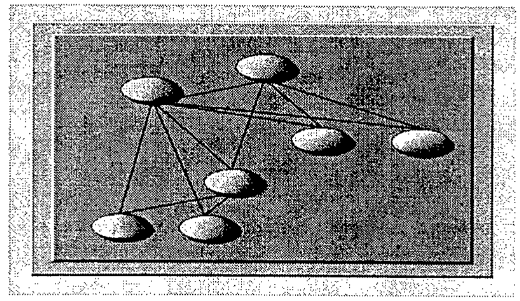

Computational Models of Human Organization Dynamics

Final Technical Report

Sponsored by
Defense Advanced Research Projects Agency
Information Systems Office
Computational Models of Human Organization Dynamics
ARPA Order No. E495
Program Code No. 6S10
Issued by DARPA/CMO under Contract #MDA972-97-C-0001

Objectives

Our work aims to produce a capability to develop and validate models of organizational dynamics, or more broadly, of psycho-social systems. This modeling and analysis activity is in service to planning, decision-making, design, intervention, or other actions that may be taken with respect to managing and improving the psycho-social system in the light of the model.

Four objectives derive from this aim. First, we need to develop principled methods to gather the data, develop models based on the data, validate the models, project implications of the models under varying conditions, and maintain the support for models and conclusions. Second, owing to the complexity of the phenomena and the potentially vast array of conditions that may be of interest, we need to develop computational tools to support data capture, modeling, model validation, and analysis/visualization of model implications. Third, we need to demonstrate how this capability would productively support individuals, groups, and organizations in their work. Fourth, these methodology and technology-development activities should be used to test and improve our theory of social/organizational dynamics, which forms the foundation for our work.

Technical Problems

The technical problems that derive from our objectives apply mainly to the technologies we seek to develop. During the course of this first contract year, we have focused our attention on the creation of an initial set of tools for modeling and model validation. This has lead us to focus on technical problems such as the representational adequacy of the tools we have developed to specify the data, agent practices, agent architectures, social relations, and the interdependence of these. The material in the remaining sections of this report describes our responses to the technical and methodological issues. In the remainder of this section, we present a few of the longer-term technical challenges.

As we proceed computability/tractability problems will need to be formally addressed. Here, we mean both the practical computational requirements for simulations of the kind we envision, as well as the fundamental limits of computation. In the case of the former, we must face processing limitations in the expression of simulation experiments. In the case of the latter, issues such as the computational complexity (e.g., an experiment we might wish to run may contain NP-hard tasks), or nonlinearity in the dynamics, may have critical consequences for what we can or cannot in principle accomplish.

We are particularly concerned that the tools support validation of models. Among other reasons, the complexity of the models can be sufficiently high that they could potentially be tuned to generate almost any reasonable (real-world) data. This issue has theoretical, methodological, and technical features. For the near term, we are placing a primary emphasis on maintaining rationales (evidence and assumption) for data and modeled structures. For the intermediate term, tools that perform sensitivity analysis, and simple tools that verify that at least some of the variables in an experiment are not mentioned already in the data, would provide a significant service. Over the long term, simulation studies should be embedded in rigorous experimental frameworks suitable for intervention-oriented action research.

There are many important human-computer technical problems. There are presentation and complexity-management concerns, driven by the need to make the tools intuitive, robust, and usable by a wide variety of users in various (often difficult) problem-solving situations. And, there are deeper issues having to do with the degree to which the interface both supports users and guides/disciplines them to enter data and build structures and studies that are consistent with the theoretical and methodological assumptions that ground the environment. For example, how might we depict the degree and nature of validity conditions met by different elements of a model?

Research and Development Methodology

In a nutshell, the method of this project has been to articulate/design, construct, seriously critique, and revise a collection of prototype tools aimed at supporting computational modeling and analysis of organizational dynamics. We have worked hard to develop critiques from the point of view of theory and method, as well as relative to various demanding domains, such as are faced by crisis response organizations. We have not simply built some tools and tested them on a toy example.

In more detail, our research and development methodology for this technology derives from our theoretical and methodological commitments, and most importantly, from our concern for validation. We have developed a theory of psycho-social dynamics – Generative Practice Theory, or GPT. Its distinctions define the items that must appear in our technologies (i.e., social structure across agents, agent architecture, practices). Hence part of our R&D method is to implement tools for users that permit entry of data and construction of models along these lines. Validation here consists of assessment of whether the tools are adequate for the sorts of organizations, agents, and practices that appear in the crisis domains of which we are familiar.

Similarly, we have an overall concept of operations that asks users to critique models and evaluate their quality relative to events in the world. This concept of operations, called critical reflection (or, in group interaction, critical dialogue), is based on treating data and models and simulation results as hypotheses that ground in empirical data (observations or testimony), in models, in assumptions, and/or in scenario specifications. R&D method follows directly, in that we maintain rationales for all data and structures, and we propose that users interact by articulating, critiquing, and revising hypotheses. We say "propose that users" since we have not built technology to support model-sharing and critique – though we do have tools to carefully maintain the source for each structure.

Beyond this, we have during the course of this project tested our theory and methods in the light of data from a variety of domains, including: non-combatant evacuation operations, health-care delivery (clinic design, critical and intensive care units), business decision making (critical events, supply chain design), response to hazardous materials spills, and municipal planning (crisis response, general long-term planning). We consider it essential to undertake these tests, owing to the dependence of our technology on our concepts of organization and agency and on our methods of investigating these – the various domains present challenges from 'outside of' our theoretical and methodological work.

Technical Results

We present our technical results for this contract in this section. We begin with a summary of our results relative to the project objectives. Then we describe the theoretic framework, practice mapping methodology, and the technology we have developed based on the theory and methods.

Appendix A at the end of this document presents additional detail on our concepts and project results, and includes an algebraic formulation of the underlying theory that guides this work.

Executive Summary of Results

We have met each of the project objectives described in the first section.

We have significantly refined our methodology, called Practice Mapping, for gathering data, modeling, and validating models of the inter-dependent practices of agents in psycho-social systems.

We have developed a collection of tool suites, in the form of prototype user interfaces that support the activities of psycho-social modeling and analysis. Specifically, we have developed prototypes for: data capture, formalization of data as practices, computational reification of practice descriptions in a form suitable for monte-carlo simulation, design of experiments to validate models and/or project its implications, analysis/visualization of results, and maintenance of rationales (evidence and assumptions) for each item in the system (data, model fragment, result). Our goal, given project resources, was simply to develop look-and-feel interfaces for these user activities. However, we found we were in many cases able to develop real, usable tools.

We have developed an overall concept of operations for the use of these tools. It is a form of structured dialogue among concurrent users, based on articulating and critiquing requirements and models, then selecting and undertaking realignment steps based on the resultant critiques. This is a type of critical method. The following activities occur: users develop models and cast them as hypotheses with attached rationales (supporting data, assumptions); users share these structures; users critique the models and/or justifications of others; users design and run simulation experiments, to validate models, to explore anomalies (i.e., trusted data from the world is inconsistent with conclusions of trusted models), and to develop model implications (e.g., to support planning); and users gather data, revise requirements, and update models accordingly.

As a mechanism to test our overall progress on methods and tools, and as a way to demonstrate how they would be used to support important actions, we invented and implemented in the tools a detailed non-combatant evacuation operation scenario. We demonstrated the way the tools are intended to be used, in a formal project review at our offices, on 12/1/97.

Finally, we have made progress on our theory and methodology, based on the development of tools and the challenges offered by the various domains we have investigated. For example, we have been able to refine our concept of the communication-based inter-dependence of multi-agent practices, based on work we did to elicit information flows and interpretive relations in a local municipality. We learned it is quite productive to elicit problems and opportunities informants see with respect to joint actions, then ask who should be communicating and about what topics as well as whether the specified conversations are occurring. This sharpened our theoretic understanding of the communications infrastructure (one of the three dimensions of social structure stipulated by the theory). It also led us to see the value of building up networks of interpretation flows and blockages, which then suggested requirements for the kind of entry and display facilities we need in the tools to capture and reason about this sort of data.

Foundations, Methods, and Technologies

In this section we summarize the theory, methods, and prototype technologies that have been developed or refined on this contract.

Theory: Generative Practice Theory

We have developed a theory of social activity, called Generative Practice Theory (GPT). This theory of social activity forms the foundation for our methods and technologies; in turn, the methods and technologies offer tests that have helped us revise and refine the theory.

The name indicates our major commitments. First, the theory is generative – social structure is understood as regularities generated and re-generated through the enactment of agent practices (i.e., structure means 'stable features of dynamic social processes'). A primary feature of social structure is that action has two consequences; one is the accomplishment of the aims of the action, and the other is to establish the relevance of that kind of action, enacted that way, in those kinds of situations, by those kinds of enacting agents. So action occurs within and has the effect of reproducing social structure.

Second, the theory is about agents viewed as bundles of practices, understood psychologically (cognitively) and socially. We have developed a cognitive architecture for individual agents. And, we have developed a conception of the primary categories that determine overall social structure. Most importantly, we have developed a concept of practice that is uniform with respect to individuals and collectives, thereby spanning the individual and the social with a single construct. We discuss the architecture, social structure, and practice construct below.

Third, the theory is dynamic, reflecting several important commitments. Dynamics are entailed by any view of structure that consists of regularities in processes, as we claim is the best way to think of social structure. And, human action is dispositional – our knowledge is in the form of competencies that embody tendencies to carry out practices in situations, and to adapt them to the situation and according to internal criteria. Some practices are ongoing (non-terminating), unless intervened upon – e.g., blood sugar maintenance, bureaucracies that persist even when the problems they were formed to address have disappeared. Finally, a dynamic conception is necessary for the uses of the theory – simulations to validate our understanding of social systems, and to develop new designs, and etc.

The GPT is summarized by its commitments to individual and joint practice, to social structure, and to a resource-bounded architecture for individual agency. One preliminary: we use "behavior" to refer to an instance of an enacted practice, and "practice" to refer to the general capacity of the agent to undertake and manage a situated action. So a practice is a generating structure and a behavior is an actualized manifestation of the practice.

Virtually all human behavior adapts during the course of the action's enactment. In particular, behavior is regulated such that criteria are attended to in the face of the ongoing progress of the action within the (usually changing) situation. This in-process adaptation is observed in the lowest-level physical motions, in the highest levels of thought and communication, and etc. In contrast with open-loop systems that run to completion once triggered, closed-loop ones are responsive to information and interrupts during the course of their execution. We therefore represent practices as closed-loop systems.

With this in mind, we depict practices formally in terms of the "Action Proposition" – a template of the form "In Situation S, to achieve/maintain Criteria C, do Action A". Note the contrast with (open-loop) situation-action rules, which have the simpler form "In Situation S do Action A". Our more general form is necessary to capture human practice, with its features of monitoring, interruption, and in-process adaptation (with arbitrary temporal grainsize).

The computational realization of an action proposition is called an "Action Schema". Action schemata implement action propositions via active triggering and data environments (i.e., the situation component), a procedure (i.e., the action component), and a collection of monitors that execute *concurrently* with the procedure and when activated suspend, abort, or redirect the executing procedure (i.e., the criterial component). The practices we implement generally have non-deterministic procedures.

We focus particularly on joint practices. These are practices of multiple agents, where the ongoing adaptation internal to an agent's enactment reflects conditions that have been, are, or are projected to be enacted by other agents. Practices, such as multiple individuals collaborating to move a table, plan an operation, or engage in discussion, are obviously joint. Yet many actions that appear to be individually undertaken are regulated by the social norms and are for this reason joint. (The standard contexts for types of behaviors, the assumptions (often tacit/hidden) that underlie tasks and behaviors, the criteria that are used to measure and adapt results, and features of the action that reflect the way other agents will interpret it, all serve to constrain and guide the action in the light of other agents, even if it seems to be undertaken in isolation.) In essence, if an action can be named in a way that others find meaningful, or if others adapt in any way to the occurrence or products of the action, then the action is joint.

The practice model has been formalized algebraically, in the language of automata theory. We employ Bucci automata, which allow for non-deterministic, non-terminating, interruptible behavior. The agent and social structure models are developed as constraints over these automata. We have proved theorems that characterize the conditions under which an agent must be seen as a collective, structured by resources/communications/role-relations, versus when the agent can be treated as an individual.

The theory of social structure that informs this work is that there are exactly three categories of joint practice, three ways that practices can be seen as inter-dependent. All practice inter-dependence can be captured by some combination of one or more of these ways. The three dimensions of inter-dependence are resource flows, communication flows, and role relations.

All practices adapt to the availability and quantity of resources, both physical and informational, and so the flow of resources in a social system is a strong determinant of its dynamic structure (and performance). Concretely, we inquire into the devices/artifacts that are required and/or may be used or consumed by each practice, as well as the information/data that is required and that may be used.

Communications flows, in this theory, are the flow of interpretations as a consequence of communication. Communication is distinguished from data or information, which we treat as a resource. Communication is the meaning that agents take/assign to the information of which they are aware, which, loosely, is the effect information may have on the ongoing and future interpretations and other actions of the agent. This is clearly a co-adaptive process. Concretely, we inquire into the topics of discussion that occur within the execution of a practice – how they are initiated, and what they might produce in the way of revised commitments on the part of the discussants.

Finally, practices are part of large bodies of inter-related practices called roles. E.g., to be an intelligence officer, or a planning officer, or a baker or a doctor, is to embody a host of practices that are inter-related in terms of the underlying assumptions they make and, especially, the criteria that they attend to. Roles are defined relationally, in terms of the pattern of co-adaptation in consequence of the actions of other agents. For example, we have the role relations given by: commander / intelligence-officer / planning-officer – or, parent / child, or, doctor / patient. Abstractly, role relations distribute criteria – revealed in authority, decisions, and native interpretations – over a collection of socially-interrelated agents. Concretely, we seek to elicit: conditions of practice initiation, expectations agents have for the execution of practices (i.e., about the environment, themselves, others, and about the total system of agents), how behavior is guided by expectations, and how expectations adapt as a function of the evolution of the joint behavior. The criteria served by these roles can then be extracted, so that practices that generate the role relations are centered on the maintenance or satisfaction of the criteria.

Roles are the pre-eminent element of social structure. This means that the information we elicit and the generators we develop must satisfy a condition: it has to span agents and contexts. Simply put, a practice complex is a role if and only if it is played by, and recognizable to, some sizeable collection of agents. This point is true of the resource and communications infrastructures as well, but less so.

The agency model is designed to capture situated and resource-dependent agency (including as a critical resource, the agent's personal cognitive capacity), under conditions of greater or lesser ambiguity about the next action(s) to take. A reactive capacity provides for immediate response if a set of executable practices exist, for which there are sufficient execution resources (e.g., processors). A deliberative capacity recognizes when there are no immediately-executable practices, when there are conflicts among practices (e.g., their criteria), or when there are insufficient execution resources – this triggers practices to resolve the ambiguity, uncertainty, or conflict by producing a non-empty set of actions ready for immediate execution.

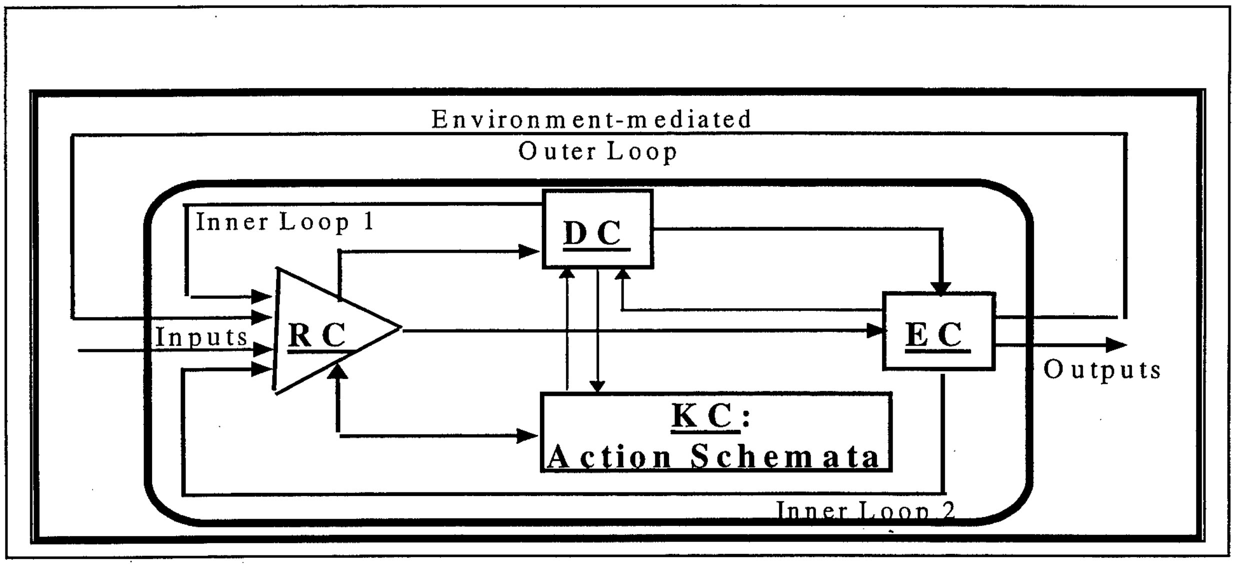

Figure 1: An Agent Model, Dynamic Schema Instantiation

Figure 1 depicts this agent architecture in terms of four components: a reactive component (RC), a knowledge component (KC), a deliberative component (DC), and an execution component (EC). The reactive component senses and captures internal/cognitive and external/environmental states. The knowledge component stores the practices of the agent that are suitable for immediate execution, should the right triggers exist in the situation and providing resources are adequate. The deliberative component resolves conflicts, manages underdetermination (choices) or ambiguities, adapts practices to available or accessible resources; its products are practices that are suitable for immediate execution. The execution component dispatches execution resources (cognitive, physical) to practices that are ready for immediate execution.

Methodology: Practice Mapping

A method is needed to enter into the rich phenomenon that is an organization and efficiently make sense of it in service to self-understanding and improvement – after all, there is no limit to the phenomena that may be selected for study and possible improvement. The ideal method would embody our detailed theoretic positions on agency, social structure, and individual practices/competencies. And, the method must support the discipline imposed by our aim of developing computational realizations of what we learn. Further, and unlike many social scientists, we believe it is necessary to have a basis for validating statements/models of the social system – our work is not just about creating hypotheses. For all these reasons, a rigorous methodology is required to guide observation, interviews, and experimentation aimed at gaining an understanding of the system, and this method must support a suitable concept of validation.

We have found that such a methodology does not exist – e.g., participant-observation methods in anthropology and ethnography are inadequate – and so have created "Practice Mapping". Practice Mapping is the collection of methods used to elicit or observe organizational phenomena, develop a computational realization of the phenomena in terms of a generative system of practices, and validate that computational model.

The main features of the Practice Mapping method are summarized in Figure 2. The figure depicts Practice Mapping, or simply mapping, as cycling between articulation/focusing methods, and critiquing methods. Each small box represents a number of specific methods we have developed, some of which are described very briefly here. In fact, both of articulation and critique can be undertaken concurrently, so the figure depicts a logical, not necessarily sequential, relationship among sub-methods.

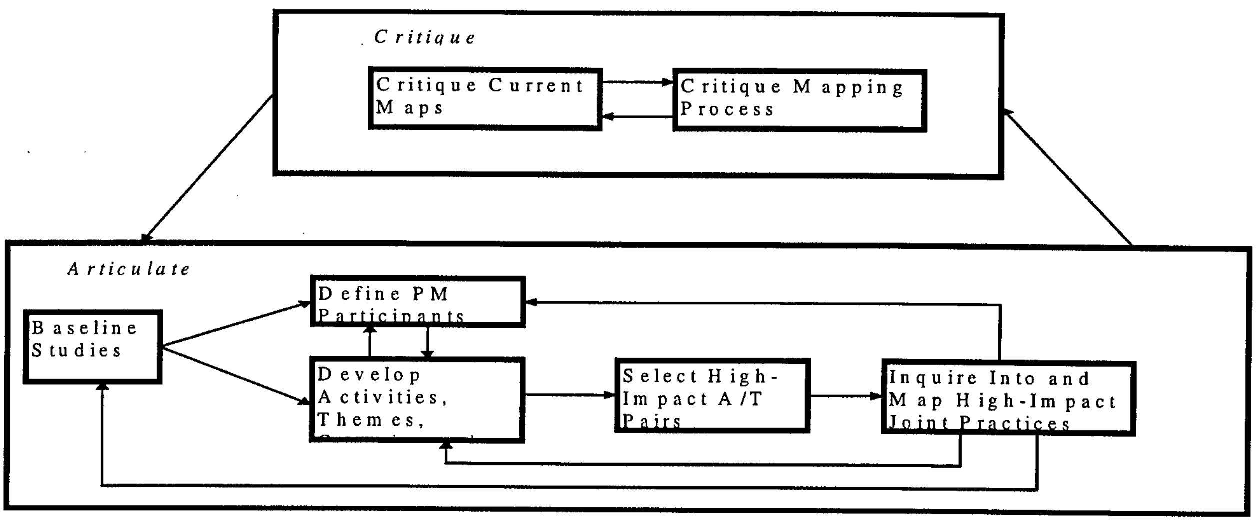

Figure 2: The Structure of Practice Mapping

We now present the essentials of the Practice Mapping (PM) method. Baseline articulation studies establish the general features of a class of social systems of which the subject one is an instance. E.g., a review of response directives and general crisis response materials (e.g., the EPA's "Incident Command System" for responding to hazardous materials spills) precedes our mapping of any particular organization's responses. In addition to learning the 'lingo' of an organizational field, baseline studies can produce a preliminary analysis of the reliance of the subject type of organization's reliance on its physical plant and other resources, the typical tasks the organization faces, the types of workflow or decision-making styles that are used to coordinate task response, the types of participants for different tasks, and broadly-shared norms on actions. Background analysis can contribute directly to the organization model. But in most cases its role is to inform hypotheses that practice mappers (i.e., the individuals we train to use practice mapping methods) use to focus the general elicitation strategies of practice mapping to the particular characteristics of the organization types with which they are about to begin work.

Practice Mapping participants are those within and external to the system that are interviewed and/or observed during mapping. Our methods allow us to discover through practice inter-dependencies the agents or types of agents that need to be consulted, and so the participants change over time in a way that responds directly to the quality of the modeling.

We employ three concepts to determine entry points (phenomena in the form of action patterns) that we should investigate, and we revise these as part of the process. Behaviors are the enacted practices, as described previously. Constituencies are collections of agents that relate to a practice or action pattern either as participants or as stakeholders – so practice structures are primary determinants of which individuals and groups need to be explicitly modeled as agents. Themes are patterns of action that are judged by the PM participants as sufficiently important and/or sufficiently recurrent to warrant analysis (e.g., "we fail at the end of a project to reflect on what we might learn that would help us do better next time"). Themes are a generalization of objectives that are sensitive to the practice structure of the social system being mapped. We have developed a process for eliciting candidate themes from participants, an instrument for eliciting judgements about the significance and recurrence levels on the candidates, and a simple but robust statistical analysis method that identifies themes and constituencies of those themes. Hence themes enable us to determine, at any point in the PM process, what to investigate next.

Once themes are identified and high-impact ones selected (i.e., according to the statistics on importance, recurrence, and scope of the affected constituency), we inquire into and map the practices that generate the theme. We seek the ways behaviors are situated, the main flow of the action, and the way action is adapted and interrupted. We are also guided by the social structure model, so we seek to understand how the practice is situated in a larger social setting, and how that setting both arises from and constrains the practices of agents. We seek to understand the role relations among agents, the constituent practices of the roles, the co-adaptation (initiatives and responses) as the roles are enacted, and then the way communication and resource manipulations occur under the guiding influence of the controlling roles. Further, we are guided by our agent architecture to understand and depict the resource-dependent nature of action, the conditions that determine how critical recurrent uncertainties are resolved, and the skills agents embody for resolving conflicts and ambiguities.

A large number of specialized methods have been developed and tested to acquire specific features of the practices, agents, or social structures. For example, we help PM participants articulate the criteria that inform practices by varying the resources available to a practice from non-existent to minimal to opulent. We then study how agents adapt their action – the actions/commitments that are neglected, kept, substituted, enhanced, and etc. to learn what is criterial, or valuable, under that practice. Similarly, we ask for examples of terrible, poor, adequate, good, and great outcomes, and how they might arise, to learn about the range of internal and environmental contingencies, and associated responses, embodied by a practice.

All PM inquiry is guided by the fact that human action is generally defensive – we humans are reluctant to open ourselves to the risk failure or error, and we are reluctant to admit to them as well. Yet these defenses have profound impacts on what we can and do learn, individually and collectively. We employ special communications methods to surface and explore defenses.

As recollections of behaviors are elicited and generalized as practices, their basis in participant testimony and other observations is captured/recorded with the data and models. Note that as part of practice mapping practice, we audiotape and/or videotape all interviews; therefore, a data record is available for later review and for critique of mapping practices.

Once practices are elicited and formalized as action propositions, they become available for computational expression as action schemata. We discuss this in the section on technology.

In addition to articulating practices, agents, and social structures, we are compelled to critique the ongoing models. A number of techniques have been developed and tested to critique both the contents developed by mapping thusfar and the mapping process itself. For example, the data underwriting a model must be evaluated for its scope and validity, and the process must be continually reviewed for its attention to the right themes and its assessment by the right participants. The records we keep on the basis for model elements are used here.

Practice mapping concludes when the objectives that stimulate mapping have been satisfied. Usually, this means that a generative model linking themes (in terms of dynamic execution) has been developed, inter-subjectively validated by means of experiments and assessments of participants, and used to support interventions to improve the social system – one or more of training, redesign of resource flows or communications or role relations, technology insertion.

This brings us to the requirement of validating social models. The overall validation method is to seek ongoing critique of a model and its implications, and to revise them accordingly. Psycho-social models can never be proved true – the next piece of data might invalidate what has been elicited and induced thusfar (e.g., reveal that a generalization must be made contingent) – but they can be shown to withstand certain criticisms, subject to what we now know. In other words, a model is explicated in terms of its supporting and refuting evidence as well as modeling assumptions, and it is then justified in terms of the challenges that it is known to withstand.

We have tested and employ the following three specific validation methods. We develop computational experiments, attempts to generate simulated results that correspond (and ideally, predict) observable organizational phenomena. We also seek direct review of models by participants and stakeholders mentioned in or affected by the elements of the model. Experiments, and inter-subjective agreement on the model, validate that model up to the point that critiques are successfully met. A third validation method is to intervene to change the social system on the basis of what has been learned through simulation. Arguably, if one can on the basis of a model cause predictable change/improvement, then it seems fair to say that a valid model of the social system has been developed. Usually, these validation methods are combined – agents critique models, then mature models are subjected to experimentation, and interventions are designed, analyzed, enacted, and evaluated. Each of these methods is implemented within a cycle of articulation, critique, and realignment.

Technology: ACCORD

We have begun to implement a collection of technologies, called ACCORD, to assist us with practice mapping, modeling, and analyzing social systems. ACCORD stands for Agency, Cognition, and Coordination in Organizational Dynamics. ACCORD is composed of four suites of tools:

1) Mapping Suite
   a) collect and record data on agents, practices, resources, and the environment – the latter two being represented as static and dynamic objects
2) Modeling Suite
   a) construct a computational realization of agents, using the agent architecture, and of action schemata, linked to supporting mapping data
3) Experimentation/Analysis Suite
   a) define execution parameters for monte-carlo simulations of the computational realization of the system, under various environmental or practice assumptions,
   b) define measurements (probes) to be made to capture data during simulation runs
   c) develop experiments to validate the model of the system, or to explore what appear to be anomalous relations between system observations and model implications
   d) develop studies that derive implications of the system's performance and dynamics
   e) analyze and visualize the results of validation experiments and simulation studies
4) Simulation Suite
   a) The distributed discrete-even simulation engine and distributed processing network that executes action schemata (within agent architectures) concurrently.

We have begun to prototype the first three suites of tools. We have not yet built the simulation suite, but have funds to do so and will be developing this engine over the next 6 to 8 months.

The three suites of tools are, in essence, interfaces. They support capturing data, designing and capturing models, and designing and capturing analyses. These interfaces have been developed, in JAVA, on a powerful PC laptop (Toshiba Tecra 740CDT). We selected the laptop since this makes it a portable device that can be carried to the subject organization and used as a "mapper's notebook" during practice mapping.

The simulation engine will run on a distributed network of workstations (DEC Alpha's). Presently, we plan to develop the host for the engine (i.e., the controlling manager for the simulation processes) on a DEC 600au, a very powerful workstation. This manager will assign execution resources to simulation processes; these resources reside on other DEC workstations, one of which is a powerful dual-processor workstation. The use of JAVA allows the laptops and the host workstation to share data and control information.

The remainder of this section illustrates the use of the tool suite. We include screen captures of one element within each of the mapping and modeling and analysis tools, and describe the process associated with their use. The interfaces have been instantiated according to a demonstration scenario we have invented, of a NonCombatant Evacuation Operation (NEO) enacted under conditions of high threat and deception on the part of key actors. In this scenario a collection of rebels has formed and assembled west of the capital city and is marching west to east through the city to the government palaces, with the intention of taking control. It has been determined that American citizens should be evacuated. They assemble at the embassy and must travel north to south, crossing the rebel path, to get to the evacuation docks. The interactions of government troops, rebels, evacuees, and evacuation forces are studied and used in real time to inform evacuation planning decisions. Deception and hidden actors figure into the scenario, as a way to motivate why and how models are validated.

The screen captures that follow illustrate the fundamentals of each of mapping, modeling, and analysis. However, we stress that mapping, modeling, and analysis are not ends in themselves but rather are in service of action in the world, e.g., operations review, design, planning, decision-making, management, and etc. When we demonstrate the tools we show how mapping, modeling, and analysis activities are combined (in crisis situations) to develop, critique, revise and explore the implications of psycho-social models for operations. Hence psycho-social modeling can serve processes of interpretation, planning, decision-making, communication/coordination, and execution monitoring, e.g., for real-time support for the management of complex in situ responses to crises.

We now present graphics from each of the prototype interfaces.

Figure 3, from the mapping interface, illustrates dual methods for developing an understanding of agent practices. On the left is displayed the 'bottom-up' method, in which agents are defined as compositions of other agents, which are composed of complex practices, which are composed ultimately of primitive practices. On the right is the social structure view; mappers use it to establish the role relations within large-scale practices, e.g., the various roles taken with respect to the civil uprising. Roles devolve into practices – ultimately, the same practices that appear as constituents in the bottom-up view. Hence, both approaches lead to the overall set of practices, with the bottom-up method disciplining the modeler to think about how practices are aggregated within agents, and with the top-down or social discipline disciplining the modeler to think about how practices are structured across agents (that is co-adapt according to resource, interpretive, or role/authority relations).

Once these practices are named, the user can enter (via additional tools) information about the practice in the form of behavioral data, and then build action propositions based on the behavioral data. In particular, mappers are provided with tools to encode their elicited and observed data about practices, and then to formally describe that data in terms of the situations to which the practice responds, the criteria that determine what can interrupt and/or redirect the enactment of the practice, and the general procedure of response. The description of the practice in terms of situations, criteria, and procedures forms the basis for the development of the computational expression of the practice.

Over the long term these tools will provide direct support for specific mapping methods we have developed. For example, one of our methods, aimed at developing the detailed structure of a practice, is as follows: vary resources relevant to the practice from absent or minimal to opulent; elicit the way the practice adapts; determine what is dropped, preserved, and added in consequence; and from this induce the criteria being served and the contingencies that are recognized. This method suggests that the tool to enter behavioral data might be required to support the structured maintenance of a family of behavioral observations, indexed by informant, situation, and features of resource availability.

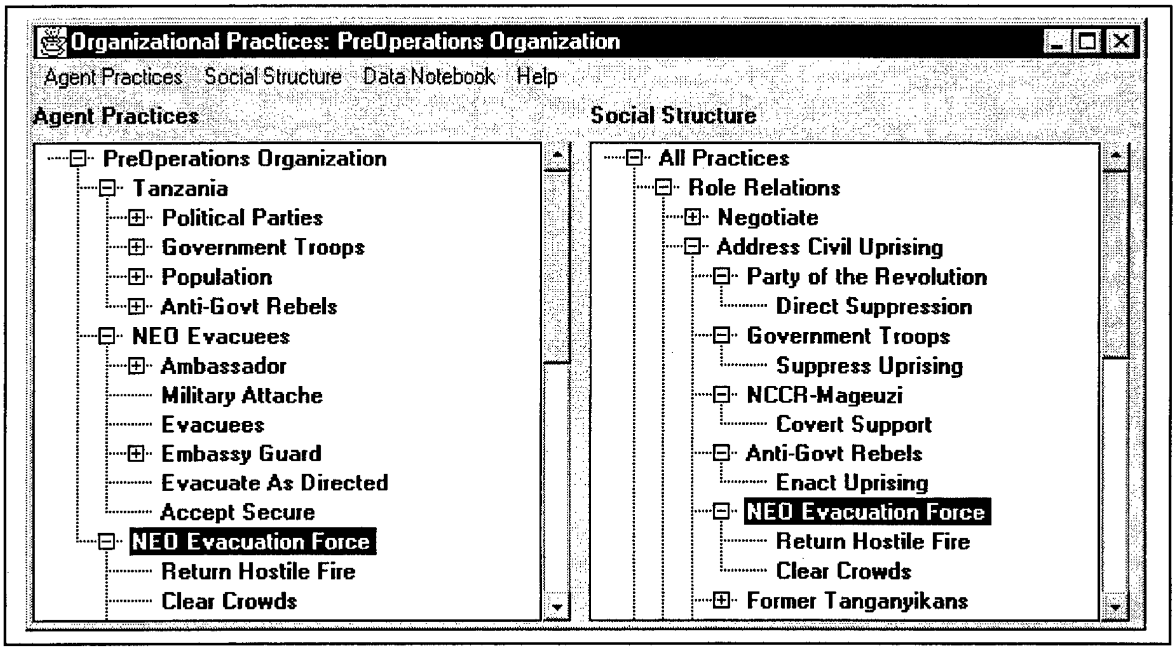

Figure 3: A Sample from the Mapper's Interface

Figure 4 illustrates one modeled practice. Data has been collected and structured in the form of an action proposition prior to this step. Here, we see the components of an action schema have been filled in for a practice of the rebels in the scenario – the small battle that defines their behavior when they reach the goal location (i.e., the government palace they wish to control). The trigger and binding environment capture the situation information, the procedure the basic action (which is itself non-terminating) internal to the practice, and the monitor list the criteria that are used to regulate the action – manage its execution, terminate it if successful, suspend or abort it under failure conditions.

The more general process is as follows. Once practices are elicited and formalized as action propositions, they become available for computational expression as action schemata. For this task the components of the action proposition are converted to the associated forms in the action schema. Situation information, usually derived from the contexts identified in behavior observations, is partitioned into triggering information and in-process information (depicted as local data access/storage structure). Process information is converted into a procedure that not only captures the basic plan of action but also contains restart loci that provide for contingent and variable response to interrupts. Criteria are converted to monitors – special procedures that execute concurrently with the main procedure. When the conditions a monitor tests for are met (e.g., a criterion is violated), the monitor a) suspends the procedure pending some conditions that warrant restart, b) restarts the procedure (at a new control point) based on new conditions (e.g., steer/adapt behavior to restore satisfaction of the criterion), and/or c) aborts the procedure (e.g., as a result of measuring success, deciding to accept failure, or losing interest in the action).

Action schemata are linked to their basis in modeling assumptions and action propositions, and so ground in the data gathered from interviews, observations, and surveys.

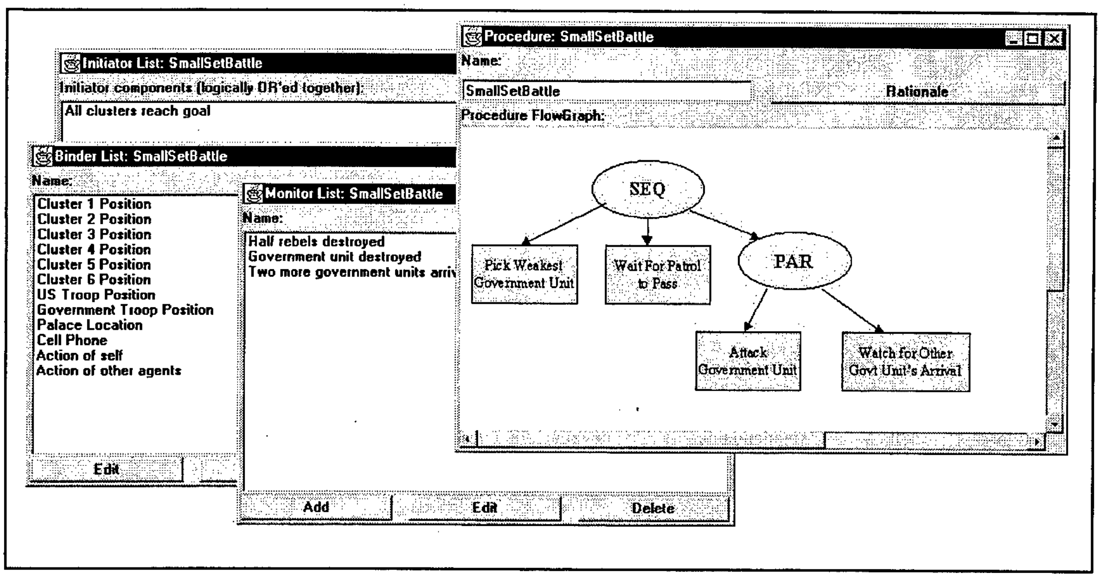

Figure 4: A Sample of the Modeling Interface

Figure 5 illustrates capabilities of the analysis environment. This interface supports the design, execution, and analysis/visualization of experiments and studies. By experiments, we mean simulations directed toward validating or refuting elements of an organization model (e.g., to test whether it can generate statistics on practices that are not part of the model basis and yet empirically verifiable). By studies, we mean the projective implications of practice response dynamics, for various situations, of a validated model.

The method of conducting a study or an experiment is as follows. Users select a variation structure for a model of interest; anything from a practice component to a practice to an agent to an entire organization can be selected. This will be varied for the purposes of the study or experiment. On the basis of structures marked for variation, scenarios are developed that define elements of a space of alternatives over which the varied structure ranges. Probes are additionally designed to take measurements of the varying structures during simulation runs. Results (probe products) are stored and made available to analysis and visualization applications.

In the case of the screen capture shown in Figure 5, an anomaly is being studied. The analysis of anomalies is central to our model validation methodology, and in fact our tools are designed to maintain careful records of anomalies and analyses based on them. Anomalies are puzzles: seemingly trustworthy real-world data is incompatible with implications of a currently trusted model. We study anomalies by developing an experiment that enlarges the range of variation of a given model structure, often beyond what would seem reasonable, and then testing whether under simulation any of the model variants is capable of generating outcomes that match real-world data. If a successful variant is reasonable it may be taken as the new model; if no reasonable variant exists then the model is called into question.

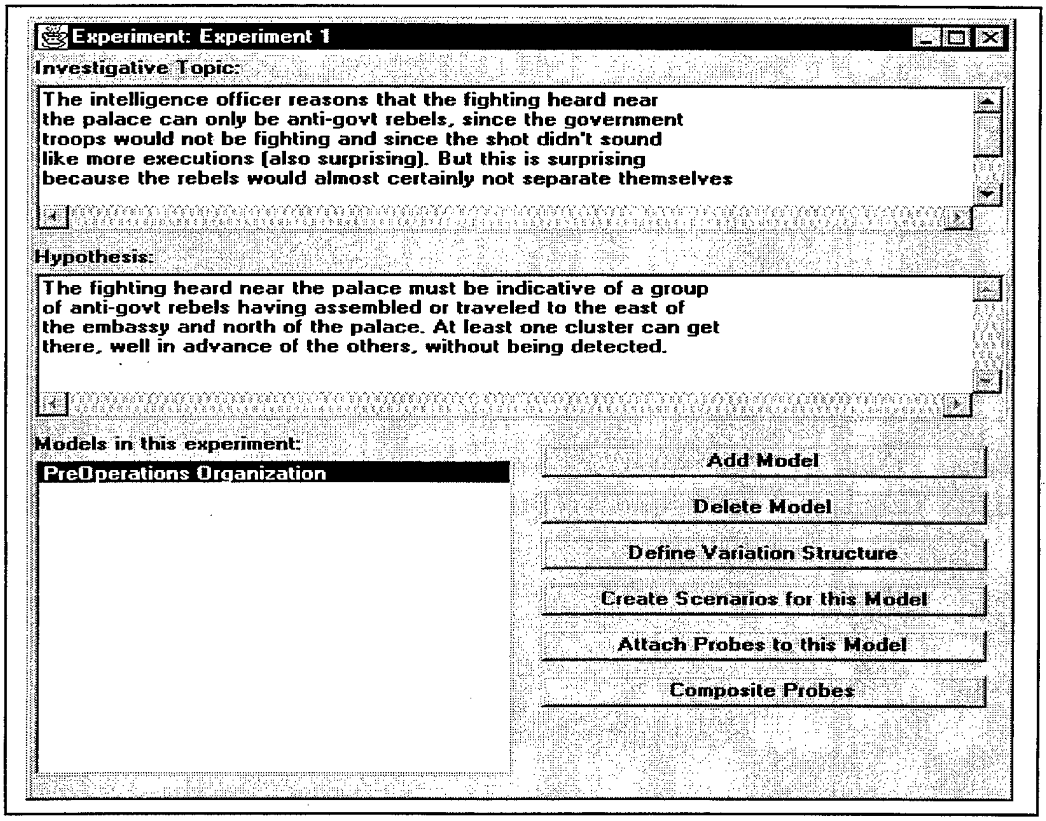

Figure 5: A Sample of the Analysis and Execution Interface

Important Findings and Conclusions

At the start of this contract we were in possession of a detailed theory of psycho-social dynamics, as well as an initial version of the practice mapping methodology. Hence following the order of development of the statement of work, at the project's inception we believed our first and primary task to be the development of a discrete-event simulation engine suitable psycho-social studies, followed by the other tasks, e.g., refine Practice Mapping methods and develop a validation concept.

We came to see that the simulation engine places few hard constraints on the specification of models, and offers relatively little guidance on the method that might determine the best use of simulation – and the opposite is true in reverse. Also, the technical issues of the engine are nearly silent on the critical issues of validation, which are central to the utility of the overall tool suite. Hence we discovered that our research and development effort would be better spent on creating a suitable concept of operations, and then working from this to specify interfaces to express and manipulate data, models, and analyses – in essence, reversing the order of tasks in the SOW. Thus we developed the concepts and methods and then on this basis a collection of interfaces that support and discipline users in the light of the method. We are now much better prepared to develop a suitable discrete-event simulation engine for psycho-social dynamics.

Psycho-social modeling, and the way of working that attends it, are new and offer significant challenges. We believe the work we have done thusfar, while not a final confirmation, is at least strong support for the claims that modeling and analysis of psycho-social dynamics is feasible and that it can be successfully integrated into the work of teams. Further, while much work remains, we believe this type of work can be so integrated not just in support of 'off-line' studies, but even as real-time in situ support for teams facing complex and significant challenges (e.g., interpretation, decision-making, and planning for crisis response).

We have tested our concepts and methods, and to a lesser extent our technologies, in a variety of domains. In particular, we have assessed our progress not only with respect to the NEO scenario we invented, but also with respect to what we are learning about hazardous materials response, health-care delivery, business planning, and municipal planning. In each case we are focusing on the need to respond effectively to critical change (e.g., to crisis). We find that the methods, concepts, and tools are informed by yet robust in the face of each domain. This gives us confidence about the research program, the tools, and the ultimate viability of this research for commercialization.

Special Comments

No special comments.

Implications for Further Research

Our initial prototyping efforts have revealed the need for a collection of specific tools, as well as more powerful representations for some of the central constructs. We summarize the central tools we believe are needed below. We present both the foundational requirement the tool addresses and the specific technologies that we think might be brought to bear.

It is clear that the deepest tests of our work, and the ultimate utility of it for tasks such as crisis response, require a running simulation engine. We are ready to build a discrete-event simulation engine, optimized for the monte-carlo simulation of inter-dependent agents behaving according to their concurrent, inter-dependent, architecturally-managed practice dynamics. Such an engine does not exist, since it needs to respond to all of our other commitments; we will design and build one ourselves.

We need a uniform representation for agents, their constituent agents, agent architectures, agent complex and/or aggregate and then primitive practices, and observations/testimony (recollected or hypothesized behavior) used as the basis for designing a practice. At present we use special representations for each; this turns out to increase the difficulty of browsing and validation. We believe a meta-graph (i.e., graphs whose nodes can be entire graphs) representation might be the right general representational abstraction over agents, practices, and evidence. Such a facility would replace the form-based input we have now. By virtue of its support for abstraction, it would provide a quick-entry capability for mappers, as they gather initial/partial information on practices.

We have said that data, models, and results each have the status of hypotheses in our method. Each structure is the product of elicitation, observation, abstraction, assumptions, and design ingredients, and each is potentially criticizable at any time (recall the articulation-critique-realignment framework of critical dialogue, employed by users to assess requirements and models). To support this, we maintain links to the evidence and assumptions of each item. However, the current facilities are non-uniform, and do not in every case represent the item in the form of a hypothesis with more or less justification. We need a template for hypotheses and their supporting and refuting bases that can be uniformly applied both to the structures in the system and as a format for critical dialogue among users. Argumentation templates, such as may be found in Gregg Courand's PhD dissertation ("Cooperation Via Justification-Based Consensus Formation Processes," Stanford University, 1991) or as proposed by Stephen Toulmin ("The Uses of Argument," Cambridge Univ. Press, 1958), offer candidates.

Validation of psycho-social models can be an infinite task. At what point can we claim that our model of some agent is of sufficient scope, offered according to useful distinctions, internally consistent, and properly predictive of or congruent with real-world events and structures (i.e., complete, efficacious, sound, and true)? We need a principled technique to guide our selection of what to validate, and to help us assess the extent to which we have or have not so validated a structure. At present, we employ the concept of an anomaly as an entry into these issues: anomalies are instances in which a trusted datum from the world either cannot be expressed or is inconsistent with implications of an otherwise-trusted model. Critical dialogue raises and investigates and handles anomalies. We seek to formalize this using a rigorous calculus, and are exploring the adequacy of Bayesian decision-making for this purpose. If adequate, the concepts of sensitivity analysis and of value-of-information could provide useful tests for validation purposes.

Distribution of this Report

The following individuals/organizations comprise the distribution list for this report.

DARPA/ISO
Attn: Dr. Steven Flank
3701 N. Fairfax Dr.
Arlington, VA 22203-1714

DARPA/OASB
Attn: Library
3701 N. Fairfax Dr.
Arlington, VA 22203-1714

Defense Technical Information Center
Attn: OCC
8725 John J. Kingman Rd., Suite 0944
Ft. Belvoir, VA 22060-6218

DARPA / OASB
Attn: Ms. Connie Jacobs, SBIR PM
3701 N. Fairfax Dr.
Arlington, VA 22203-1714

Appendix A

The slides that follow, taken from our final project review for this first/base phase of the contract, present a careful summary of the concepts and methods on which our work has been based, as well as the results of this effort. They also include material on the formal algebraic foundation for the theory of generative practices.

COD: Computational Organizational Dynamics

Project Review — Dec. 1, 1997
Presented to Dr. Steven Flank, DARPA ISO

Synergia's Project Team:

Dr. M. Fehling (PI); Dr. G. Courand (Co-PI)

E. Stroh (Proj. Mgr.)

Research Staff: J. Chalidabhongse, A. Guo, C. Wheat, Maj. R. Steinrauf

Agenda

Project Overview

Work on Underlying Concepts / Theory

Demo of Practice Mapping Technology

Research / Development Issues

- Explicate *psycho-social dynamics* of important types of organizational problem-solving
  - Descriptive *models*
  - Prescriptive *standards*
  - *Focus effort on* *crisis planning and response*
- Explore methods/technology to enhance effectiveness of organizational problem-solving
- Design prototype of crisis-response support tools based on simulation/analysis (ACCORD)

Progress in Project Phase II-A

- Extensive domain analyses underway
  - Military crisis-response: NEO
  - Hazmat control and other disaster-relief activities
  - Critical and emergency health care
- Formal theoretical framework (GPSD)
- Refinement of *Practice-Mapping* methodology
- Design prototype of ACCORD: Suite of practice-mapping tools

Collaborators

- Key military analysts (e.g. Army Staff College Major; Middle East, SpecOps, NEO expert)
- HazMat and Disaster Relief Organizations
  - EPA
  - Various California and Bay Area agencies
- Monterey Institute for Int'l and Strategic Studies
- Naval Postgraduate School
- Stanford and Mt. Zion Hospitals

Agenda

Project Overview

Work on Underlying Concepts / Theory

Demo of Practice Mapping Technology

Work on Underlying Concepts and Theory

Domain and Task Conceptions

Theoretical Framework

Methods and Technology

Defining Features of Crisis

- **Threat:** Must act to preserve key *values/criteria*
- **Uncertainty:** Prior plans obviated at outset
  - Risk (defined by current *frame* of understanding)
  - *Frame error* (e.g., invalid content, incompleteness)
- **Urgency**: Resource-limited rational response
  - Deliberation versus
  - Reaction versus
  - Adaptation (cf., Newell–"preparation vs deliberation")

Crisis Response

- Crisis = Problem solving under *threat*, *uncertainty*, and *urgency*
- Divide effort among
  - — Acting on current understanding
    - Deliberation
    - Prepared reaction
  - — Learning (frame adaptation)
- *Manage* ***organizational*** *as well as physical dynamics*

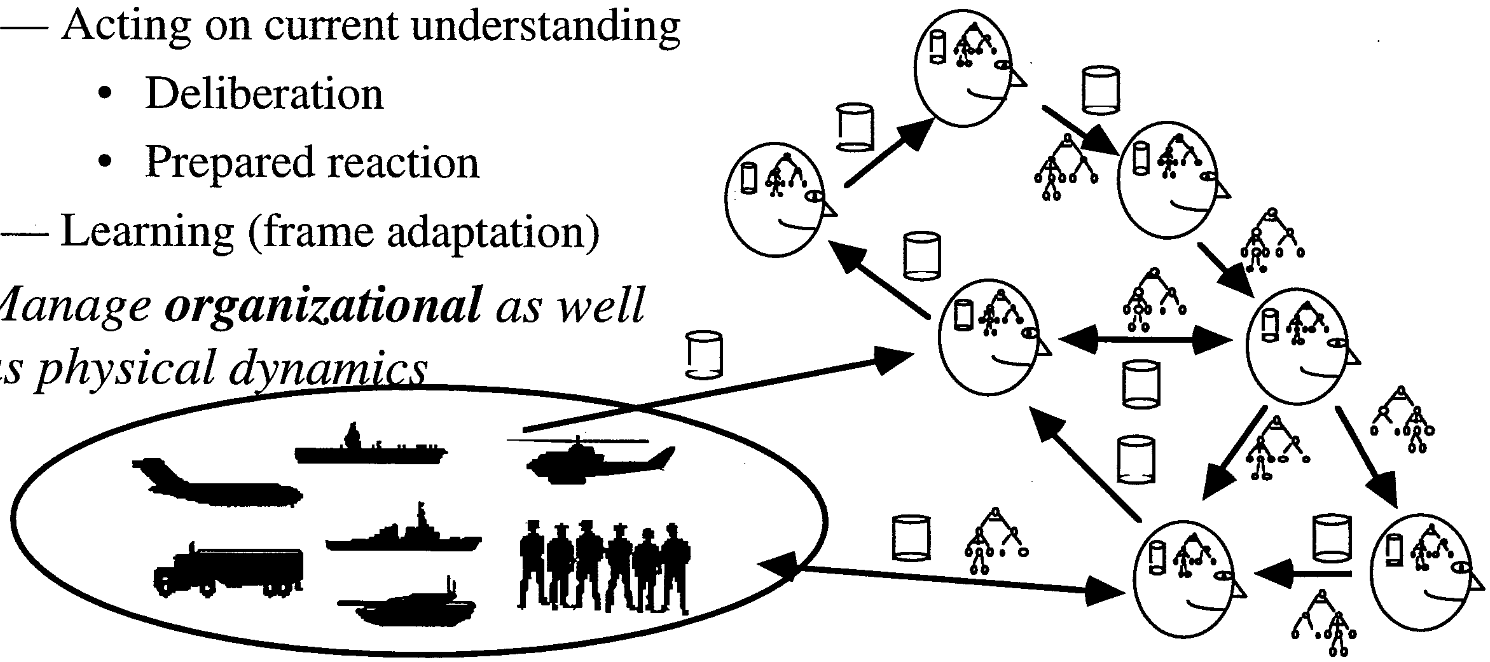

Crisis Response as an Organizational Practice

- Each organizational actor embodies a subjective view of a crisis situation & possible responses
- These subjective perspectives
  - Typically tacit, i.e., unarticulated
  - Embody potentially conflicting ideas and divergent practices
- Crisis response entails more than complex choice
  - Need to articulate, coordinate, and *synthesize* divergent frames (views)
  - *Anomalies* and *conflicts* must be identified and constructively managed

Work on Underlying Concepts and Theory

Domain and Task Conceptions

Theoretical Framework (with ODC)

Methods and Technology

GPSD — Generative Psycho-Social Dynamics

- Social organizations and actors (agents)
  - Actors: Bundles of *practices* embodying beliefs, values, commitments as skills (i.e., *competences*)
  - Social Structure = Dimensions of actor interdependency:
    - Resource flows
    - Communication
    - Role/authority relations (roles as generalized, tacit task structure)
- Explanatory focus: *Tacit, self-sustaining structure-in-practice* (cf., A. Giddens, J. Habermas)
  - Recursivity: Organization -agent
  - Agents/orgs as *self-organizing computational systems*

Formal Structure of Practice

- *Action proposition* — In situation **S**, to meet criterion **C**, perform action **A**

Action Schema: Computational Structure

Initiator: (trigger pattern)

Binder: (data access links)

Procedure: (specification of non-deterministic, structured courses of action)

Monitors: (event-triggered, asynchronous interrupt handlers)

An Actor Model: Dynamic Schema Instantiation

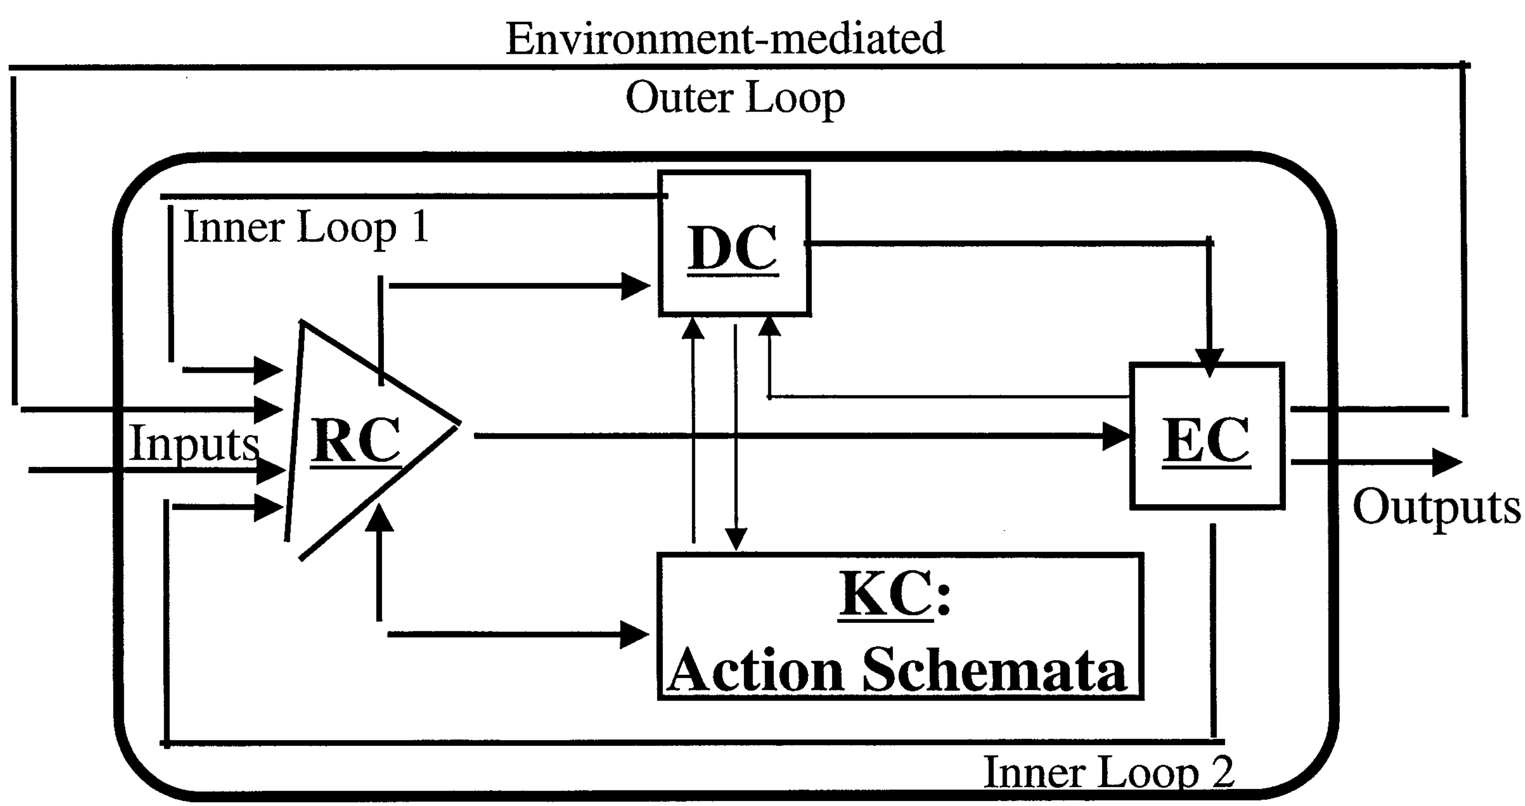

Analytic Challenge: Formal Foundation

- Indeterminacy in interpreting and validating complex simulation models
  - Validation-by-demonstration (e.g., AI) fine for system design purposes but not for empirical science
  - Interpreting and validating simulation-models
    - Credit-assignment problems challenge interpretation
    - Too many degrees of freedom: How much of the implementation matters?
- *Need a rigorous, mathematical version of theory*

Mathematical Formulation of GPSD (cf. Appendix)

- Practices = Non-deterministic Σ-*automata*
  - Bücci (non-halting) automata: Event-driven continuous processes
  - State splitting: Resource/capacity constrained process dynamics
- Actors = Bundle of practices coupled to environment via input translation and effector actions (viz., env. as automaton; cf., Genesereth)
- Organization = Net of coupled practices
  - Strength of coupling determines status of net as actor/agent
  - Captures adaptive feedback and learning processes of orgs/actors

Work on Underlying Concepts and Theory

Domain and Task Conceptions

Theoretical Framework

Methods and Technology

Organizational Crisis Response

Crisis =*Threat*, *uncertainty*, and *urgency*

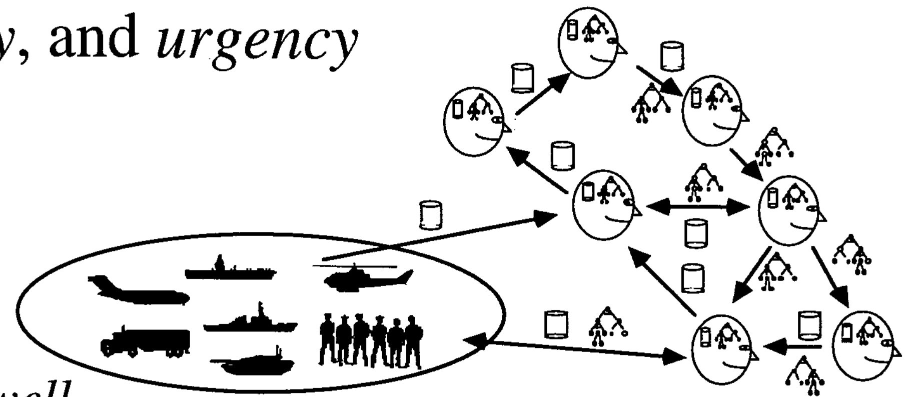

- Balance effort among
  - Reaction
  - Deliberation
  - Learning
- *Manage **organizational** as well as physical dynamics*
  - Articulation, test, and coordination of divergent frames
  - Constructively test/manage conflicting views and commitments to action

••• Support *anomaly-driven adaptation* •••

••• *via **critical dialogue*** (CD) •••

What is *Critical Dialogue*?

Collaborative organizational practices

(problem solving and other actions)

grounded in reflective, critical communication

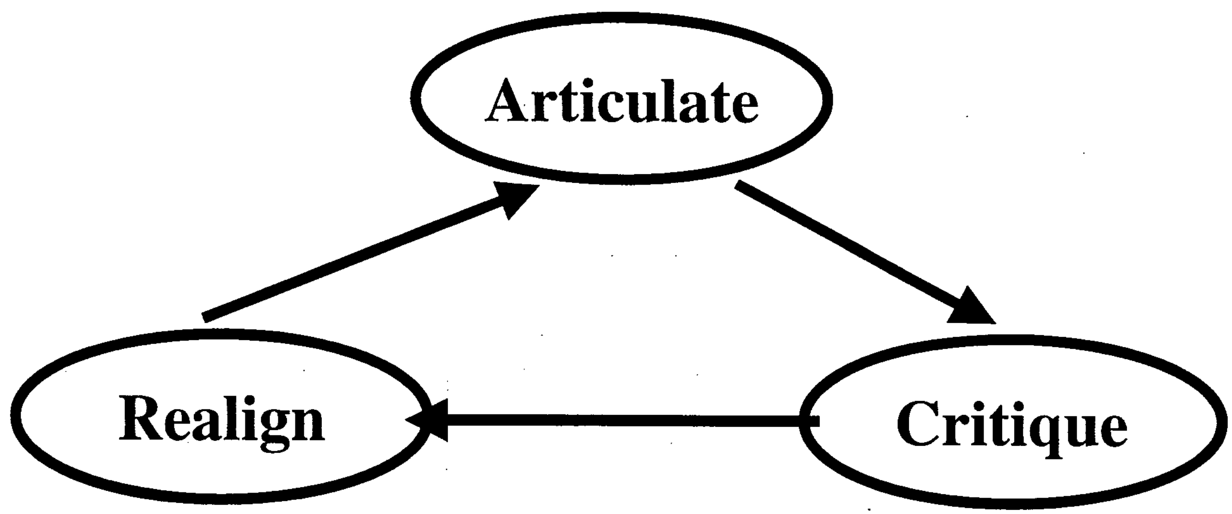

Critical Dialogue: Impact on Social Organization

- Impact on organizational dynamics
  - The substance/structure of organizational practices impact CD via constraints on resource and communication flow
  - CD impacts organizational structure as actors critique commitments to action embodied in views of facts & problems, roles, resource policies, etc.
  - Contribution to organizational adaptivity
  - Critical, coordinated deliberation (analysis w'in consensus frame)
  - Critical, coordinated reaction (within consensus frame)
  - *Critical, coordinated learning (breaking an existing frame)*

Critical Dialogue: Impact on Crisis Response

- Key features of Critical Dialogue
  - Plans, criteria, and other action commitment articulated and revised
  - Addresses problem *framing* as well as problem *solving*
  - Enhances convergence to consensus
- Exploiting organizational conflict/diversity
  - Explore diverse perspectives to critique problem frame
  - Exploit critiques within *and of* the consensus frame
  - Coordinate diverse responses within consensus frame

Agenda

Project Overview

Work on Underlying Concepts / Theory

Demo of Practice Mapping Technology

Overview of Practice Mapping

- Descriptive and prescriptive modeling discipline
  - Articulate and validate model of organizations' dynamic practices
  - Critical appraisal of practice dynamics/structure
  - Aid in revising practices and their interdependencies
- Organizational planning/design *via analysis*
  - Habitat mapping
  - Mapping practice interdependencies as coupled action schemata
  - Projective analysis to appraise organizations' entailed capacities
  - *Model revision to capture planning/redesign commitments*
  - Projective analysis to appraise implications of plans/redesigns

ACCORD Tool Suite

ACCORD Simulation and Analysis Suite

— Mapping (Data collection) tools —

— Model implementation tools —

— Experimentation/validation tools —

(for planning as well as analytic problem-solving)

Agenda

Summary and Conclusions

Research / Development Issues

- Explicate *psycho-social dynamics* of important types of organizational problem-solving
  - Descriptive *models*
  - Prescriptive *standards*
  - *Focus effort on crisis planning and response*
- Explore methods/technology to enhance effectiveness of organizational problem-solving
- Design prototype of crisis-response support tools based on simulation/analysis (ACCORD)

Progress in Project Phase II-A

- Extensive domain analyses underway
  - Military crisis-response: NEO
  - Hazmat control and other disaster-relief activities
  - Critical and emergency health care
- Formal theoretical framework (GPSD)
- Refinement of *Practice-Mapping* methodology
- Design prototype of ACCORD: Suite of practice-mapping tools

Next Steps: Project Phase II-B

- Complete domain analyses
- Develop alternative scenarios from other domains to test generality of approach
  - Hazmat (EPA)
  - Corps level planning, command & control (Lunceford)
- Complete design, refine, and test prototype
- Implement (or *import*) suitable dev. platform
- Formulate phase III implementation plan

Appendix: Algebraic Formulation of GPSD Theory

Formalized Notion of Practice

Actors as Practice Complexes

Organizations as Loosely Coupled Networks of Actors

Toward a Mathematical Theory of Organization

- Algebraic description of a *practice*

  Let Σ be a set. Then define an*abstract practice* on Σ,

  $\Re_{\Sigma} = [Q, I, T, E]$, where Q is a finite set of states,

  I, T $\subseteq$ Q initial & test (terminal) states,resp., and

  $E \subseteq Q \times \Sigma \times Q$ is the set of transitions in $\Re_{\Sigma}$.

- So, an abstract practice is a *non-deterministic* Σ-*automaton* as in mathematical automata theory (cf., Eilenberg)

Formalizing Agents and Organizations

- *Embedded practice (systems)* represent agents and organizations coupled to their environment.

  If S and A are sets, then $\Re_{\Sigma,S,A} = [\Re_{\Sigma}, \rho, \lambda]$, where $\rho : S \to \Sigma$, and $\lambda : Q \times \Sigma \to A$, is an *embedded practice*.

- Essential refinements
  - <u>Resource / capacity constraints:</u> $Q = H \times N$; where is N is a strict order, and least N "causes suspension"
  - <u>Non-halting processes</u> (Bücci); now T represents *test criterion* states (entered infinitely often)

Formalizing Social *Structure*

- Structure: Interdependencies among constituents
  - Sequential
  - Parallel
  - Cyclic (*feedback*)

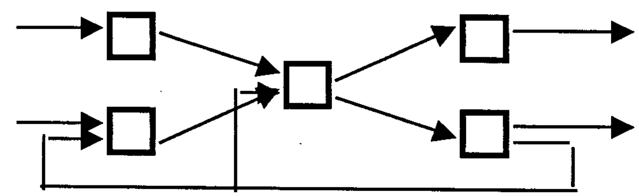

- Theorem A: A finite set of interdependent agents can be represented in a single embedded practice.

  Proof: By induction show that sequential and parallel dependencies among constituents define equivalent embedded practice system.

Formalizing Agent *Structure*

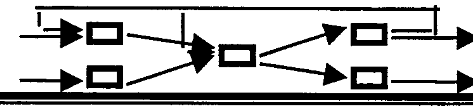

- Agents as *networks of state-output processes:*
  - Let $\Re_{\Sigma,S,A} = [\Re_{\Sigma}, \rho, \lambda]$ be an embedded practice system. If, $\forall q \in Q, \sigma, \sigma' \in \Sigma, \lambda(q,\sigma) = \lambda(q,\sigma')$, then $\Re_{\Sigma,S,A}$ is a *state-output* (*s.o.*) *practice*. We often write $\Re_{\Sigma,S}$.
  - Let $\aleph = [\{\Re_{\Sigma_i,\Sigma_i}\}, I, A, \{\rho_i\}, \lambda]$, where I, A are inputs and actions (outputs), resp., and $\rho_i : (\times Q_i) \times I \rightarrow \Sigma_i$, for $1 \le i \le n$, and $\lambda : (\times Q_i) \times I \rightarrow A$. Then we call $\aleph$ a *practice complex* or, alternatively, an *agent – actor.*

