REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
The public reporting burden for this collection of information sources, gathering and maintaining the data needed, and co information, including suggestions for reducing this burden, Davis Highway, Suite 1204, Arlington, VA 22202-4302. Re collection of information if It does not display a currently vai	is estimated to average 1 hour per res ompleting and reviewing the collection of to Department of Defense, Washingtor spondents should be aware that notwith d OMB control number PLEASE DO I	ponse, including the time f of information. Send common Headquarters Services, I standing any other provisi NOT RETURN YOUR FOR	or reviewing instructions, ser nents regarding this burden Directorate for Information Op on of law, no person shall be RM TO THE ABOVE ADDRE	arching existing data sources, searching existing data estimate or any other aspect of this collection of perations and Reports (0704-0188), 1215 Jefferson a subject to any penalty for failing to comply with a SS
1. REPORT DATE (DD-MM-YY) May 2016	MM-YY) 2. REPORT TYPE 3. DATES C Final 12 May 2		OVERED (From - To) 2005 - 11, May 2016	
4. TITLE AND SUBTITLE Network and Information Scien	ces (NIS) International	Technology All	iance (ITA)	5a. CONTRACT NUMBER 5b. GRANT NUMBER W911NF-06-3-0001 5c. PROGRAM ELEMENT NUMBER 6FE0N0
6. AUTHOR(S) Dinesh Verma, IBM US David Watson, IBM UK				5d. PROJECT NUMBER C201 5e. TASK NUMBER 5f. WORK UNIT NUMBER
<ol> <li>PERFORMING ORGANIZATION NAME IBM T. J. Watson Research Center 1101 Kitchawan Road Yorktown Heights, NY 10598</li> </ol>	(S) AND ADDRESS(ES) IBM United Kingdom Labs Emerging Technology Services Husley Park Winchester, Hants		8. PERFORMING ORGANIZATION REPORT NUMBER ITAFINAL001	
9. SPONSORING/MONITORING AGENCY UNITED STATES ARMY RESEARC LABORATORY (ARL) US ARMY RDECOM ACQ CTR - W9 S. MIAMI BLVDDURHAM NC 2770.	NAME(S) AND ADDRESS(E H UNITED K DEFENCE D11NF4300 3 (Dstl) Porton Dov	S) (INGDOM MINI (MOD) ience and Techno vn, SALISBURY	STRY OF ology Laboratory , SP4 0JQ	10. SPONSORING/MONITORING AGENCY ACRONYM(S) RDECOM/ARL MOD/Dstl         11. SPONSORING/MONITORING AGENCY REPORT NUMBER(S)
<ol> <li>DISTRIBUTION/AVAILABILITY STATE Approved for public release; di</li> <li>SUPPLEMENTARY NOTES</li> </ol>	MENT stribution unlimited.			
<ul> <li>14. ABSTRACT <ul> <li>The final report for the NIS ITA</li> <li>a) The Invoice for full NIS IT</li> <li>b) The ITA Book summarizin</li> <li>c) The ITA Science Library p</li> </ul> </li> <li>15. SUBJECT TERMS <ul> <li>Network, Information, Science,</li> </ul> </li> <li>16. SECURITY CLASSIFICATION OF:</li> </ul>	A is comprised in three p A costs incurred (attach g the program, success rovides a rich user inter International, Technolo 17. LIMITATION	parts: ed) factors, and tech face to ITA pate ogy, Alliance, N 18. NUMBER	inical achievemer ents and publications IS, ITA 19a. NAME OF RE	nts (attached) ons- see http://nis-ita.org SPONSIBLE PERSON (Monitor)
a. REPORT b. ABSTRACT c. THIS F Unclassified Unclassified Unclass	AGE OF ABSTRACT: fied SAR	OF PAGES 165	John Pellegi 19b. TELEPHONE (301) 394-2	ino NUMBER (Include Area Code) 100

Standard I	orm	298	(Rev	. 8-98)
Prescribed	by A	NSI	Std.	Z39-18



# Network and Information Sciences International Technology Alliance

## Network and Information Sciences International Technology Alliance

Network and Information Sciences International Technology Alliance

Research was sponsored by U.S. Army Research Laboratory and the UK Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the UK Ministry of Defence, or the UK Government. The U.S. and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

All rights reserved. This book or any portion thereof may not be reproduced or used in any manner whatsoever without the express written permission of the publisher except for the use of brief quotations in a book review or scholarly journal.

First Printing: 2016

http://nis-ita.org

#### Authors

Graham Bent, IBM UK Dave Braines, IBM UK Cheryl Giammanco, U.S. Army Research Laboratory Thomas La Porta, Pennsylvania State University Kin Leung, Imperial College London Gavin Pearson, UK Defence Science and Technology Laboratory Tien Pham, U.S. Army Research Laboratory R. Srikant, University of Illinois at Urbana-Champaign Paul Smart, University of Southampton Michael Underhill, Underhill Research Ltd Dinesh Verma, IBM US David Watson, IBM UK

#### Editors

Seraphin Calo, IBM US Christopher Gibson, IBM UK John Ibbotson, IBM UK Alun Preece, Cardiff University Ananthram Swami, U.S. Army Research Laboratory Don Towsley, University of Massachusetts Amherst

#### Reviewers

Trevor Benjamin, UK Defence Science and Technology Laboratory Greg Cirincione, U.S. Army Research Laboratory Stuart Farquhar, UK Defence Science and Technology Laboratory Brian Rivera, U.S. Army Research Laboratory Mudhakar Srivatsa, IBM US

## Table of Contents

I.	Foreword
11	Executive Overviewxiii
	NIS ITA Technical Accomplishments
	ITA Collaborations
	Conclusion xx
	References: xx
1	Introduction to the NIS ITA 1
	Operational Background and Coalition Requirements 1
	Strategic Aims
	Achieving the Objectives 6
	Promoting Collaboration
	Technical Scope and its Evolution 10
	Promoting Cross-Disciplinary Collaboration
	How it all Worked
2	Network Infrastructure for Coalition Operations
	Network Tomography 23
	Energy Efficient Networking
	Fundamental Performance Limits of Hybrid Networks
	Robust Routing in Coalition Networks
	Multipath Control of Hybrid Networks

Mobility Modelling in MANETs 38
References
Security for Coalition Operations
Policy-based Security Management
Information Theoretic Security Capacity
Key Management in MANETs 48
Hierarchical Identity Based Encryption
Declarative Infrastructure for Security/Networking
Knowledge-Based Security Policies
Verified Outsourced Computing
References
Distributed Information Processing in Coalition Networks. 63
Distributed Dynamic Processing
Quality of Information
Mission Optimized Sensor Networks
Sensor Mission Assignment
Dynamic Distributed Federated Databases
Fault Isolation in Information Networks
Self-Organization of Services
References
Human and Cognitive Issues in Coalitions
Advances in Argumentation Theory
Shared Understanding
Controlled English
Advances in Cognitive Modelling
Cultural Network Analysis
Analysis of Coalition Differences
Collaborative Planning Model
References

6	Exploiting the Science
	Broad Achievements
	Public Domain and Open-Source Technologies
	Exploitation Narratives
	Final Note 117
	References
7	External Perspective
	The NIS ITA
	Summary of the 2015 Annual Fall Meeting
	Some Reflections on Peer Review 2008–2016
	Critical Success Factors
	Concluding Remarks in September 2015
8	Achieving an Outstanding International Collaborative
	nesearch Alliance
	Challenges of International Collaborative Alliances
	Challenges of International Collaborative Alliances
	Challenges of International Collaborative Alliances       128         Success Factors       131         Summary       138
	Challenges of International Collaborative Alliances       128         Success Factors       131         Summary       138         References       139
A	Challenges of International Collaborative Alliances       128         Success Factors.       131         Summary       138         References       139         NIS ITA Leadership       141
AB	Challenges of International Collaborative Alliances       127         Challenges of International Collaborative Alliances       128         Success Factors       131         Summary       138         References       139         NIS ITA Leadership       141         NIS ITA Metrics and Publications       143
AB	Challenges of International Collaborative Alliances       127         Challenges of International Collaborative Alliances       128         Success Factors       131         Summary       138         References       139         NIS ITA Leadership       141         NIS ITA Metrics and Publications       143         Metrics       143
A B	Challenges of International Collaborative Alliances       127         Challenges of International Collaborative Alliances       128         Success Factors       131         Summary       138         References       139         NIS ITA Leadership       141         NIS ITA Metrics and Publications       143         Metrics       143         Bibliography of Scholarly Publications       143

## Table of Figures

The three primary goals of the programme
The NIS ITA technical management process
TA1: Network Theory 11
TA2: Security across a System of Systems 12
TA3: Sensor Information Processing and Delivery 14
TA4: Distributed Coalition Planning and Decision Making 15
TA5: Coalition Interoperable Secure and Hybrid Networks 17
TA6: Coalition Information Processing for Decision Making 18
Network satisfying criteria for additive link tomography 25
Greedy algorithm performance with number of monitors 26
Reduction in power attainable with cooperative networking , 28
Critical phenomenon in linear networks
Disruption in coalition networks motivating the creation of a new Inter-Domain Routing protocols for MANETs
MPTCP experimental setup: 2-path (solid), and 4-path (solid and dashed)
Multicast on physical network constrained by command and control hierarchy
Graph representation of a cellular/MANET hybrid multicast problem
Policy life cycle model for coalition operations 44
Sample network for analysing security capacity 47
Declarative networking framework
The DDFD concept
Cultural difference in planning between US and UK 87
Collaborative science to research exploitation

Figure 23.	Edgware Fabric logical bus 102
Figure 24.	Expanded Gaian Database node
Figure 25.	Distributed policy management using WPML and the Gaian Database
Figure 26.	Operational context for the Pathfinder transition projects 107
Figure 27.	NIFC deployment of the Gaian Database 111
Figure 28.	Multi-layered conceptual models to support collaborative planning in the CPM
Figure 29.	Sensor Assignment for Missions (SAM) framework 115
Figure 30.	Visualizing NIS ITA paper co-authors

## I Foreword

In May 2006 a landmark collaboration known as the *Network and Information Sciences International Technology Alliance* (NIS ITA) was initiated by the U.S. Army Research Laboratory (ARL) and the UK Ministry of Defence (MOD). This 10-year basic science programme was one of the first examples of a US/UK collaborative research programme under the 2002 Memorandum of Understanding established between the U.S. and UK Governments to facilitate co-operation on defence research. This programme was executed by an international research Alliance comprised of ARL and the UK Defence Science and Technology Laboratory (Dstl), integrated with a consortium of 25 leading academic and industrial organizations from both the US and UK with deep expertise in the fields of network and information sciences. This book has been written at the end of this highly successful ten-year programme, covering a range of perspectives of the work and the results achieved by the integrated technical leadership and wide research collaboration.

A core focus of the research has been on coalition operations—specifically the advancement of fundamental science that enables more effective coalition operations—with an emphasis on the hard problems associated with distributed decision making that enable rapid, secure formation of ad hoc coalition teams. This focus, coupled with the existing research interests and expertise of the twonation, multi-organization, cross-discipline Alliance has brought a sharp focus to the programme, enabling different skills from academia, industry and government to be used together to great effect. The excellent track record of the NIS ITA in terms of basic scientific advancement and impact on current and future defence technologies and capabilities is described in detail throughout this book.

Success like this, on such a large and complex programme, does not come by simply defining the goals at the outset and delivering research over the ten-year period. Instead, there was a continuous process of reinvention and refocusing of research efforts. Key aspects of this success included an environment structured and managed to foster deep collaboration among researchers throughout the Alliance; competitive research plan selection every two years that enabled adaption to feedback, scientific advances, and coalition needs; and an innovative transition model that facilitates the rapid and affordable transition of technologies. The NIS ITA also benefited greatly from rigorous peer-review by a senior team of external academic, industry and government experts with deep experience in relevant fields. The peer reviewers consistently found the programme to be wellaligned with the stated goals, collaborating deeply and creating great scientific results; and they frequently gave constructive advice that enhanced the overall programme.

As the NIS ITA comes to a close we are proud to have been a part of this Alliance team and would like to take this opportunity to thank everyone that has contributed their time, energy and ideas to this significant collaborative endeavour. Many people have contributed in so many ways to the scientific discoveries and technology transition successes of the NIS ITA. It has also been very gratifying to follow the development of individuals throughout the programme, especially students who gained experience and built their networks, in part by working in rotations throughout industry, academia and government, and often ending up employed by Alliance organizations. The *Alliance* may come to an end, but the relationships established between researchers will live on. In addition to the measurable scientific and application impact statistics, there are layers of personal stories, each with intangible professional and personal impact, which attest to the success of the program for the organizations and the scientific community writ large.

It has been a thrilling ride to be engaged in this groundbreaking joint initiative. Possibly the most significant outcome of the NIS ITA has been the strong bonds that have formed between the US and UK leadership and researchers that will endure long after the programme ends.

John Pellegrino US Collaborative Alliance Manager ARL George Vongas OBE UK Collaborative Alliance Manager Dstl

## II Executive Overview

The Network and Information Sciences International Technology Alliance (NIS ITA) was formally launched as a landmark collaborative venture in 2006. At the core of the ITA was the idea of developing a new way of doing business that required deep and enduring collaborations among government, industry and academia in both the UK and the US.

Leaders in the U.S. and UK Governments stated: "The ITA represents a new way of doing business: forming an international alliance of government, industry and academia to jointly conduct fundamental research. The origins of the concept were in similar national programmes, such as the US Army Collaborative Technology Alliances (CTA) and the MoD's Defence Technology Centres (DTC). [...] Experience has shown that the development of new technologies is best achieved when governments work in close co-operation with technology providers in industry and academia from the very earliest stages of research." [1]

The ITA was initiated with the strategic goal of producing fundamental advances in network and information sciences to enhance decision making for coalition operations. The emphasis was on enabling future coalitions that needed rapid and secure formation of ad hoc coalition teams. The domain of "*Network and Information Sciences was chosen ... due to its importance in enhancing our mutual capabilities in support of future coalition military operations*". The initial intent was for "*research focused on enhancing distributed, secure and flexible networks for information delivery and decision-making [...] while providing a better understanding of the theoretical underpinnings of networking*" [1]. The need for such understanding of the behaviour of "*large, complex and interacting networks*" including both "*social and communication networks*" was well recognized [2]. Similarly, the changing nature of warfare and the central role played by information was also recognized (cf. General Sir Rupert Smith, The *Utility of Force: The Art of War in the Modern World* ["War Amongst People"]). The NIS ITA consortium was led by IBM (*IBM Research* in the US and *IBM Emerging Technology* in the UK), and included some of the major defence system integrators in the US—*Applied Research Associates, Honeywell International Inc., Raytheon BBN Technologies, and The Boeing Company*—and in the UK—*Airbus Group Innovations, LogicaCMG UK Ltd., Roke Manor Research Ltd., and Systems Engineering & Assessment Ltd.* The academic partners included some of the most prestigious universities in the US—*Carnegie Mellon University; Columbia University; Pennsylvania State University; Rensselaer Polytechnic Institute; The City University of New York; University of California, Los Angeles; University of Maryland, College Park; and University of Massachusetts, Amherst*—and in the UK—*Cardiff University; Cranfield University; Imperial College London; Royal Holloway, University of Southampton; and University of York.* 

This consortium, along with the U.S. Army Research Laboratory (ARL) and the UK Defence Science and Technology Laboratory (Dstl), formed an international research Alliance to jointly conduct collaborative research. Two of the key aspects of the ITA were: (i) it was structured and managed to foster an open collaborative research environment to support deep collaborations, and (ii) it sought to facilitate the rapid and affordable exploitation of research results through an innovative transition model.

## NIS ITA Technical Accomplishments

Over the course of its 10 years of collaborative research, the ITA programme has made many advances in the state-of-the-art of coalition network-and-informationscience, breaking new ground in the areas of network theory, security, information processing, and shared-understanding, that are important to coalition military operations. In this overview we provide a summary of some of the key technical accomplishments. Subsequent chapters of the book describe specific technical areas in detail and provide a more in depth view of these technical achievements.

#### Key Accomplishments

The ITA programme was conceived to deliver new insights into the complex and difficult issues associated with rapid and secure access to trusted information, across dynamic ad hoc coalition teams, to enable understanding and decision making. The programme had three primary goals: (i) to make advances in the fundamental science of coalition socio-technical information systems; (ii) to promote collaboration among US and UK researchers across organizations and

scientific disciplines; and (iii) to transition the fundamental science and enhance capabilities to conduct coalition operations. These three goals are depicted in Figure 1.



Figure 1. The three primary goals of the programme

One of the underlying assumptions behind the programme was that joint collaboration between the two countries would result in outcomes that would not be possible without the synergies gained from robust UK/US collaborations.

ITA research results are evident in the innovations that are described in peerreviewed publications in leading conferences and journals, by scientific leadership at conferences and workshops, and in other ways that ITA work has influenced research in the scientific community. Exploitation of these results is also important for impact in the technical community and to advance the capabilities for coalition operations. The impact of this is felt not only in advanced capabilities but also in shaping the thinking on how technologies can be used by coalitions.

#### Selected Scientific Accomplishments

Selected scientific achievements for the ITA include:

- Network Tomography: ITA researchers developed the scientific principles underlying monitoring of dynamically changing coalition networks with minimum overhead. These insights can be used to instrument and observe a variety of networks with minimum possible probing ("Network Tomography" on page 23).
- · Distributed Dynamic Processing: The ITA programme developed the

concept of bypassing network bottlenecks at the coalition edge by moving processing within the network, and analysed approaches for mapping distributed applications onto hybrid coalition networks. It has created new techniques for distributing streaming and transaction oriented applications, analysing their performance, and improving the effectiveness of distributed applications ("Distributed Dynamic Processing" on page 64).

- Policy-Based Security Management: ITA researchers developed new paradigms for security management using a policy-based approach, creating new frameworks for policy negotiation, policy refinement, and policy analysis. They applied them to create constructs like self-managing cells, and to manage coalition information flows. The team developed techniques for determining security policies that can preserve privacy and sensitive data while allowing partners to make limited queries on that information ("Policy-based Security Management" on page 43).
- Cryptography Applications in Coalition Contexts: The ITA has made fundamental advances in making cryptographic techniques applicable in the context of coalition networks. These include the development of new identity-based encryption paradigms, efficient implementationfriendly reformulation of fully homomorphic encryption algorithms, and outsourcing computation securely to untrusted devices belonging to coalition partners ("Security for Coalition Operations" on page 43).
- Advances in Argumentation Theory: ITA researchers provided the theoretical link to accommodate trust, inconsistency and uncertainty in distributed networked information systems. They have also proposed a principled method for linking provenance data with the evaluation of competing hypotheses to counter cognitive bias inherent in human analysts ("Advances in Argumentation Theory" on page 79).
- Insights into Fundamental Limits and Properties of Mobile Network Structures: ITA researchers developed a variety of models characterizing the scaling properties of mobile ad hoc hybrid networks found in coalitions. These models determine the fundamental communication capacity of disruption tolerant networks, modelled limits on structures with mathematically tractable topologies ("Fundamental Performance Limits of Hybrid Networks" on page 29), identified information theoretic limits on capacity with security constraints ("Information Theoretic Security Capacity" on page 46), and characterized the performance of multi-path

and multi-point communications ("Multipath Control of Hybrid Networks" on page 34). A pragmatic output was a universal mobility-modelling framework ("Mobility Modelling in MANETs" on page 38).

- Energy Efficiency Techniques: The Alliance invented a variety of approaches to improve battery power consumption and energy efficiency in ad hoc networks. The approaches include distributed beam forming using cooperative communications and techniques for improving duty cycling behaviour in networks using self organization ("Energy Efficient Networking" on page 27).
- Coalition Communications Interoperability: The Alliance created new paradigms for inter-domain routing ("Robust Routing in Coalition Networks" on page 32), identified differences in coalitions cultural norms ("Cultural Network Analysis" on page 86), defined a new paradigm for shared understanding ("Shared Understanding" on page 81), and used declarative technologies for networking and security in coalition environments ("Declarative Infrastructure for Security/Networking" on page 52). Another related activity was the creation of a collaborative planning model, including support for rationale, that could be used in the future for the development of effective planning tools ("Collaborative Planning Model" on page 89).
- Quality of Information (QoI): The ITA pioneered the concept of QoI, and created the framework, algorithms, and various use cases surrounding the use of QoI in intelligence, surveillance and reconnaissance (ISR) and sensor networks ("Quality of Information" on page 66). The concept had a significant impact on the scientific community, including starting the International Workshop on Information Quality and Quality of Service for Pervasive Computing (I2QS), and being a major thrust in the Network Science Collaborative Technology Alliance (NS CTA) programme.
- Mission-Aware Information Networking: The ITA developed a variety of techniques to adapt the network to meet the requirements of a mission, including approaches for optimizing networks to meet mission needs, matching assets to missions, and isolating faults in information networks ("Distributed Information Processing in Coalition Networks" on page 63). One of the key transition outputs was the Sensor Assignment for Missions (SAM) tool for matching missions to assets available in the field to perform that task ("Actionable Intelligence Technology Enabled

Capability Demonstration (AI-TECD)" on page 114).

- Dynamic Distributed Federated Databases: ITA researchers created a model to represent sensor information flows as distributed databases, and devised the principles that allowed them to be federated dynamically in a manner that is both self-organizing and scalable ("Dynamic Distributed Federated Databases" on page 69). The work resulted in the Gaian Database technology, which has had multiple transitions to other programmes in MOD and the U.S. Army ("The Gaian Database" on page 103).
- Advances in Cognitive Modelling: The ITA made significant advances in the state of cognitive modelling, including computational modelling of specific cognitive processes, using the Adaptive Control of Thought— Rational (ACT-R) cognitive architecture for understanding collective agent and human interactions, and conducting cognitive social simulations ("Advances in Cognitive Modelling" on page 84).
- Controlled Natural Language/Controlled English: The ITA programme made several advances in using a limited subset of English to improve the usability of computing systems by soldiers in the field in a variety of contexts, including mission planning, asset allocation, and policy specifications ("Controlled English" on page 83). Controlled English (CE) led to several transition activities through the development of the CE Store open source software technology ("CE Store" on page 101).

#### Selected Transitions

Although the ITA was a fundamental science programme, exploitation of research results was critical to its success. The ITA has transitioned significant technologies using an innovative transition model to rapidly exploit research results.

From the beginning the ITA engaged with users and stakeholders in both countries. At the core of this was the continual engagement of the research triad consisting of government, industry, and academic researchers—each bringing their own diverse perspectives—to connect emerging multidisciplinary research results with potential technology needs. These transitions occurred through diverse pathways including joint demonstrations and exercises, capability enhancement to military customers, enhancing commercial products with capabilities relevant to the coalition environment, and open source software. Once transition partners were engaged, a key enabler was the use of the ITA companion transition contracts set

up to facilitate rapid exploitation. The details on these technologies and their use in various transition activities is described in "Exploiting the Science" on page 95.

A brief summary of some of applied research technologies that incorporated ITA scientific output is the following:

- Information Fabric: The Information Fabric is a messaging system that provides the abstraction of a messaging bus spanning different sensors and information sources across different coalition partners. It grew out of the experimentation needs of researchers working on mission aware information networking.
  - Gaian Database: The Gaian Database is a dynamic distributed federated database that implements some of the concepts from distributed relational algebra.
  - Sensor Assignment to Missions Tool: The SAM tool embodies the algorithms and scientific output from mission aware information networking, and provides a semantically driven assignment of resources in a tactical environment to satisfy the needs of a mission, leveraging the Military Missions and Means Framework.
- Controlled English Store: The CE Store is software that realizes the principles underlying Controlled English, enabling it to be created and experimented with for the representation of knowledge and reasoning.
- Watson Policy Management Library: This library implements many of the algorithms developed during the policy-based security management initiative. It was integrated with both the Information Fabric and the Gaian Database for many transition activities. The technology found its way into the hands of the U.S. Army as a component of an IBM network management product that it procured.

## **ITA Collaborations**

From its inception the ITA programme focused heavily on cross-organization, cross-disciplinary and international collaboration throughout the Alliance. Innovative mechanisms to promote this were introduced including: hosting an annual boot camp to establish and develop collaborative working relationships, promote multidisciplinary research, and spark innovation; summer internships for students at industrial partners; and promoting staff rotation among the Alliance members to deepen collaboration and co-invention.

In many cases, the collaborations lasted well beyond the involvement of the researchers in the programme. As students graduated from their universities, they continued collaboration with other colleagues they had befriended during their ITA years. There were several students who graduated to join the industry or government labs that were part of the consortium. For example, one student started out at a university, joined a government laboratory, and then joined an industrial research team completing his round of each type of organization involved in the programme.

By all measures of success, collaboration within the programme has been outstanding. Evidence of this includes significant collaborative peer-reviewed publications, leadership activities, and assessments by external peer reviewers.

## Conclusion

In conclusion, the NIS ITA programme has proven itself to be highly successful as exemplified by the ITA independent external peer-review panel that stated it was "an outstanding example of true, deep and enduring International Research Collaboration" that has "significantly advanced the state-of-the-art in network and information science through multi-disciplinary research". The ITA programme was also cited in the official White House Press Release during the visit to the US by the UK Prime Minister David Cameron on 14 March 2012.

Possibly the most significant outcome of the ITA programme has been the strong bonds that have formed between the Alliance researchers that will endure long after the programme ends.

## References

- T. Killion, P. Sutton, M. Frame and P. Gendason, "A New Paradigm in International Collaboration: The US-UK International Technology Alliance in Network and Information Sciences", *RUSI Defence Systems*, pp. 46-49, June 2007.
- [2] Network Science. The National Academies Press, 2005, p. vii.

## 1 Introduction to the NIS ITA

The NIS ITA programme was a unique international endeavour that brought together a large number of US and UK research institutions from academia, industry and government to work together on a collaborative research agenda for a decade. Given the size, complexity, diversity in technical focus, and cultural differences between the different organizations, challenges in managing the programme over the course of ten years have naturally arisen. The ITA has successfully addressed all of these, resulting in a highly successful collaborative programme.

In addition to addressing a set of key problems posed by the two governments, the ITA challenged an alliance of government, industry and academic partners to adopt new ways of working to break down barriers, build relationships, develop mutual understanding and work in partnership to develop advanced technologies for the US and UK military. Collaboration between partners on both sides of the Atlantic would be a key aspect of this new way of working and the success of the alliance would be measured by how well they performed this aspect of the work. In addition to the research outcomes, the new way of collaborative working would create a community of researchers that understood the problems emerging from the adoption of network centric warfare and could contribute solutions to the US and UK military. This community and its shared knowledge would exist and be of benefit long after the formal ITA research programme had completed.

#### **Operational Background and Coalition Requirements**

Information is key to addressing the issues raised by the landscape of current military operations, with enhanced situational understanding in all its forms being seen as the key technology driver. To address this, the ITA programme sought to develop the fundamental underpinnings to enable secure, dynamic, decision-enabling information flow for tactical military operations in a coalition environment. This led to the identification of two strategic goals:

- 1. Enable the rapid and secure formation and maintenance of ad hoc teams
- Enhance distributed, secure, and flexible decision making for coalition operations

The emphasis of the research would be on dynamic coalitions that bring together a number of different partners, both governmental and non-governmental, into a single operation or mission. This implies a degree of transience of existence and membership, distinct from persistent alliances with permanent infrastructures such as the *North Atlantic Treaty Organization* (NATO). Each partner would bring different cultural and social perspectives, different policies and procedures, different systems and networks, all supported by a variety of technologies and conventions. These must be brought into harmony to achieve a common goal.

The research programme would focus on technologies for Brigade-level and below, down to the lowest level, with an operational context evolving from major combat operations to counter insurgency, and military assistance to stabilization and development, including how to operate in transitional epochs. It is expected that future conflicts will likely occur in a congested, cluttered, contested, connected and constrained battlespace, and adversaries will aim to use a mix of high-end and low-end asymmetric techniques to exploit weaknesses. UK and US forces will be operating in complex situations that involve coalitions of cooperating organizations with varying degrees of mutual trust, trying to resolve complex problems with many dimensions (military, economic, political, social, legal, etc.) using a mix of military force and concerted influence.

The operational requirements arising from these complex situations will continue to evolve as forces adapt to complex coalition operations. This gives rise to a set of emerging trends that will impact future coalition operations.

*Counterinsurgency operations are complex:* Counterinsurgency (COIN) operations are information driven, and have highlighted the importance of tactical networks by placing increasing demands on their performance. COIN is an extremely complex form of warfare that places significant burden on the people and technologies that are employed. This implies that there is a need to increase the emphasis on information and analysis at the lowest operational levels, shortening decision-making timescales, and the exploitation of a wider array of information sources given increased problem complexity.

2 Network and Information Sciences International Technology Alliance

Data volumes will continue to grow, especially informal information with limited structure: The continued rapid growth in the volume of data and information in the network—from sensors and from users, particularly informal information usually with limited structure—places increasing demands on decision-making. It has been shown that most data (80%) used for decision-making is unstructured, making it hard to extract meaning, patterns, and trends. Therefore in order to support decision-making, disparate information types from multiple sources must be transformed into relevant, actionable knowledge.

*Processing power and storage capacity is increasing faster than communications capacity:* The processing power and storage capacity of computer systems is increasing (and expected to continue to increase) at a faster rate than the capacity of the links (bandwidth) between systems. Therefore research is required to understand how to efficiently use available bandwidth resources. Data and services must be smart and smartly positioned within networks to minimize their impact on any limited network resources.

There is an increased use of commercial cellular networks: The initial emphasis of the ITA programme was to understand the underpinning models of existing *Mobile Ad hoc Networks* (MANETs). However, the increasing use of cellular networks by coalitions has provided them with more communications capabilities, but the lack of security and interoperability with military networks limits their effectiveness. While the successful outcome of coalition operations continues to depend on the right information being available to the right soldier, at the right time, and in the right form, this continues to be dependent on the provision of infrastructures that maintain appropriate protection for the information and data they handle. Military infrastructures are designed to have the characteristics and functionality needed to support these requirements, but may not always be appropriate or adequate to form the sole provider of services to the military. Thus civil and commercial infrastructure services are being brought into use to operate in conjunction with the military systems. This implies that hybrid networks that exploit and interoperate with commercial wireless networks are key.

Enhancing coalition decision-making depends on secure communications and information networks: To support coalition decision making, the communication network must be closely linked with the information infrastructure, and security mechanisms must support both seamlessly. Technologies are needed that can manage and control secure hybrid networks and the information content/exchange to support decision-making. An increased emphasis on securing information infrastructures, including information exchange among multiple security domains

3

and security for distributed information services, is necessary. Any research must address the end-to-end problem of data-to-decision (coalition) across multiple security domains.

## Strategic Aims

The NIS ITA was therefore specifically designed to encourage and facilitate collaborative multi-disciplinary research between government, industry and academia across the US and UK institutions that formed the consortium. Central to the programme were three major themes:

- 1. A new, more collaborative way of working between all members of the alliance
- 2. Delivering significant, innovative and disruptive science
- 3. Achieving rapid and broad exploitation of the science (in both defence and civil domains)

In addition it was recognized that to achieve the understanding required of *"large, complex and interacting networks"* a challenge-led, multi-disciplinary approach would be required.

The need to conduct both fundamental research and rapid exploitation meant the ITA consisted of two parts: a fundamental research component and a technology transition component. The fundamental research component provided for collaborative research, the results of which would be published in the public domain, whilst the technology transition component provided for the application of the fundamental research results to defence and security applications (and included the ability to address any security or export issues that might arise).

### **Great Science**

*Focused and Innovative Science:* The primary objective of the NIS ITA was the creation of fundamental underpinning knowledge. It was recognized that this would require intellectually-hard science across multiple disciplines, where traction would be achieved through the application of a cross-community, cross-disciplinary ITA team to the problem.

*Disruptive Science:* It was anticipated that the NIS ITA would succeed in identifying a selection of nascent concepts, which might offer novel solutions to significant military problems in network or information science. It was therefore

4 Network and Information Sciences International Technology Alliance

necessary to ensure that the research programme structure encouraged work on new ideas that were potentially disruptive, rather than incremental.

*Quality Assured Science:* To ensure that the scientific quality of the research was maintained, an external peer-review panel comprised of independent US and UK industry and academic domain experts was tasked with formally reviewing the programme biennially, with an informal review in the intermediate years. The peer reviewers were also encouraged to offer guidance to improve the evolutionary direction of the research portfolio.

#### **Critical Science**

New Directions for Science and Technology (S&T): The NIS ITA research programme was created to investigate new, blue-skies concepts for network enabled capability to support coalition operations, and to set new directions for research and technology development within the wider military S&T community. It was therefore essential that the researchers were aware of the specific challenges that the military community faced in relation to coalition operations and work being undertaken in other research programmes.

Driving Applied Research: Whilst the NIS ITA is a fundamental science programme it was recognized that experimental validation of the concepts that were being developed was equally important, and that such validation would enable the derisking and development of innovative and disruptive solutions within the applied research activities of both the US and UK. In some case this would be through the validation of theories, which set the design-bounding conditions and limits for applied research solutions.

#### **Technical Impact**

*Fundamental and Applied Research:* One of the most challenging aspects of any fundamental science programme such as the NIS ITA is to provide tangible evidence of its impact. The demarcation between fundamental research and applied research is never easy to define and in reality they are inextricably inter-related. In recognition of this, from the outset the NIS ITA programme made provision for any of the basic research that indicated new, or disruptive, approaches for future coalition operations to be taken to higher technology readiness levels. The programme therefore included a transition contract mechanism that enabled the US or UK to unilaterally take and exploit any of the fundamental research results for specific national programmes. In addition joint exploitation of the research was made possible through multinational agreements such as the *Coalition* 

#### Warfare Program (CWP).

*Transitioning Research:* The availability of the transition contract mechanism associated with the NIS ITA fundamental research programme was considered to be an essential element that would enable the easy transition, development and testing of the underpinning science within more realistic settings. The creation of shared science within the Alliance, and the availability of transition contracts in both the UK and US, enabled the Alliance to exploit synergies between UK and US government funded applied research, and private venture funding, in the testing, maturation and de-risking of technologies.

## Achieving the Objectives

To ensure that its scientific objectives were met, the NIS ITA programme was specifically designed to create an environment in which researchers from industry, academia and government could be encouraged to work together on specific research projects that required scientific breakthroughs and new understanding to address specific coalition challenges. The technical management process is shown in Figure 2.



Figure 2. The NIS ITA technical management process

It was recognized that as the research programme progressed the coalition challenges would change and that, consequently, new research opportunities

would emerge. The NIS ITA was therefore designed to implement a rolling two-year programme of research. Each two-year cycle would solicit new research proposals that fostered collaborative, interdisciplinary research aimed at answering fundamental scientific questions. The proposals were then assessed and a set of projects selected to create a *Biennial Program Plan* (BPP), which was then implemented and continuously monitored against clearly defined criteria.

The solicitation process required Dstl and ARL to issue a scoping document that defined, in broad terms, the specific coalition challenges that were being faced (from a military perspective) and the areas of research that were considered to be important. The NIS ITA consortium was then requested to create project proposals that would identify the corresponding scientific challenges and how they would be addressed. The process also encouraged proposals that might suggest new or disruptive areas of research.

The following criteria for assessment were specified:

- Technical merit and innovation (including experimentation plans for validating research)
- Synergistic value of collaboration
- · Military relevance
- · Transition opportunity
- · Past performance

Technical merit was judged from the perspective of scientific excellence and the degree of risk associated with the proposal. High scoring research would exhibit scientific originality and disruptive potential, whilst low-risk incremental science would produce a low score in this category. Plans for experimental validation of the basic science were also encouraged where this was appropriate. In the later stages of the NIS ITA specific plans for using the *ITA Experimentation Framework* were also requested so that the value of the fundamental research could be judged in the context of other research activities.

Synergistic value of collaboration was a measure of the diversity of the proposed project team, and was a measure of the shared-science theme. Projects scored highly if they showed strong collaboration across multiple partner institutions and that included academic, industry and government partners from both the UK and US. Single institution proposals resulted in a low score in this category. The

proposals had to demonstrate that the project team had the resources to address the scientific challenges and how they would work together to achieve the goals. The process also requested evidence to suggest how the project might link to projects in other technical areas. This required the prospective project teams to interact during the solicitation process and aided the process of developing a BPP that exhibited an interdisciplinary, cross-sector culture and programme of sufficient mass.

*Military relevance* was included to ensure that the research was focused on the critical-science theme. The proposals were assessed against their relevance to the specific military coalition issues identified in the scoping document. If the proposal identified a new or disruptive science then this category was judged in relation to its potential to transform future military coalition operations.

*Transition opportunity* was a measure of the maturity of the proposed research project and was included to meet the requirements of the exploited-science theme.

*Past performance* was used to enable the proposed project teams to demonstrate how they had been successful in delivering the required impact in previous BPP cycles.

Proposed projects were selected using these criteria and then grouped into the different ITA technical areas (TAs) to create a coherent BPP. To achieve this within the available budget constraints often required some adjustments to the proposed project tasks, and in some cases individual tasks were removed and occasionally replaced with tasks from other project proposals. However, the selection process was transparent with appropriate feedback being given to both the successful and unsuccessful project teams.

## Promoting Collaboration

A key message from the Program Bulletin issued by ARL and MOD was that collaboration between consortium partners must be the foundation of the programme. Without this the benefits of multinational research would be lost and the programme would revert to siloed research within small groups of participating institutions. The aim was to create a coherent research community in the US and UK that understood and could address the future research demands of the two countries.

Encouraging collaboration within the ITA took several forms. Firstly, all projects submitted under a particular BPP had to be staffed by a mixture of UK and US

industrial and academic partners. This ensured that collaboration would take place not only internationally but also between the academic and industrial participants of the consortium. Government participation was via the activities of *Technical Area Leaders* (TALs) and project participants within the programme. This combination ensured that the military benefits of the research (identified by the government participants) and potential for transition (identified by the industrial participants) could influence the academic direction of the research leading to successful outcomes for the programme.

To further support this aspect of collaboration, all papers, posters and demonstrations submitted to the annual *ITA Fall Meeting* had to be by authors from multiple institutions. Although this might seem a minor constraint, it reinforced the interaction between participants within ITA research projects. Whilst this policy was not enforced for other non-ITA conferences and journals, the collaborative nature of the ITA has also resulted in the vast majority of external publications being jointly authored by participants from multiple institutions (academic, industry and government).

ARL and MOD also saw regular face-to-face meetings as being necessary to engender a collaborative atmosphere between ITA participants. To this end, annual *Boot Camps* took place that allowed researchers to work together to progress their projects. The Boot Camps were held each year in June and alternated between sites in the US and UK. Further working sessions were also held as part of the September Fall Meeting that, like the Boot Camp, was held alternately in the UK and US. A condition of ITA participation was that all researchers should attend the Boot Camp and Fall Meeting each year as part of their project obligations. In addition to these two meetings, further face-to-face meetings were held in January with other working meetings being used to augment the regular phone conferences that took place amongst project members.

Finally, ARL and MOD highlighted the importance of staff rotations within the Project Bulletin. Their objective was "to foster and facilitate collaborative research where face-to-face interaction is advantageous, to enable a researcher to utilize unique facilities, and to facilitate the exchange of research results". Staff rotations were particularly encouraged into and out of the government establishments so that civilian researchers would gain insight into MOD and ARL specific requirements. Rotation of government staff into ITA participant organizations gave them the insights into commercial practices and provided them with the opportunity to work alongside leading academic researchers. This aspect of the ITA programme has been particularly successful. An interesting example of staff rotations has been the annual visit of cadets from the *U.S. Military Academy at West Point* to IBM UK's Hursley Laboratory site. During their visits of two to three weeks, the cadets worked with IBM researchers and applied their military expertise to develop demonstrations based on ITA research. Examples of their work included identifying aspects of cyber defence of which a Commander should be aware in order to assess mission readiness, and using controlled natural language to develop a flexible knowledge database for intelligence applications.

The aspirations for collaboration expressed by MOD and ARL in the initial Program Bulletin have been achieved during the lifetime of the programme. Over time, the participants in the ITA have worked closely together and have developed strong international working relationships that have continued outside of the ITA programme; this has been a major contribution to its overall success.

## Technical Scope and its Evolution

To address the requirements identified in the previous section, the ITA research programme consisted of four TAs (TAs 1–4) during the first five years of the programme. Following the decision to extend the programme for a further five years this was reduced to two TAs (5 and 6) that combined those from the earlier period. Each TA typically consisted of three projects to address different aspects of the research.

The redefinition of the technical areas marked the shift in the requirements of the fundamental research required for improving coalition operations over the span of ten years. A major shift in emphasis between the first five years and the second five years of the programme was to leverage hybrid networks.<sup>1</sup>

### TA1: Network Theory

An adaptive self-organizing network that adjusts automatically and rapidly to ever-changing tactical situations is fundamental to the success of the future vision of the U.S. Army and UK Armed Forces. Wireless communication, a prerequisite for the agile operations of coalition forces, is susceptible to detection, identification, location, hostile jamming, abrupt loss of nodes, mobility of elements, and interference from a variety of sources. Furthermore, the information infrastructure deployed in the field must shield its users from the complexities of the underlying

<sup>1</sup> Hybrid networks use a combination of fixed infrastructure and mobile infrastructure to deliver connectivity, as opposed to a purely mobile infrastructure.

network infrastructure while allowing them to access information required to undertake the mission at hand.



Figure 3. TA1: Network Theory

The three projects for TA1 were:

- Theoretical Foundations for Analysis and Design of Wireless and Sensor Networks: This project investigated the theoretical limits of wireless and sensor networks in the military context to establish limits on capacity, scalability, reliability, detection, energy efficiency, and lifetime of networks. In addition, it developed a mathematical framework within which coalition forces could develop robust, high performance network protocols for military wireless networks.
- Interoperability of Wireless Networks and Systems: Insufficient network interoperability between coalition nations and even different military units of an individual nation's armed forces is a common barrier that dramatically inhibits the formation of agile mission groups. This project modelled and analysed the interoperability of different wireless networks

and systems. It then provided PHY (*physical*)/MAC (*media access control*)/ network/application network layer solutions for seamless interoperation, and developed cross-layer adaptation methodologies to achieve optimal performance.

• Biologically Inspired Self-Organization in Networks: Self-configuring and highly adaptive networks can significantly enhance the survivability of the infrastructure critical to a military operation. Biological systems provide extensive examples of survivable self-organization. This project investigated models, theory, and algorithms for creating self-organizing wireless and sensor networks inspired by biological systems.

## TA2: Security Across a System of Systems



Figure 4. TA2: Security across a System of Systems

Future military coalitions will consist of partners with heterogeneous technology, skills, interests, and trustworthiness. These partners will come together in *communities of interest* (CoIs) with common goals, perhaps only for a short period. This imposes new requirements, for example the ability to negotiate

interoperation between groups with different security policies, and the ability to make security policy decisions on-line and in real-time, rather than in advance. At the beginning of the research programme, existing security mechanisms did not scale, or were considered ineffective for future systems-of-systems; the primary aim of this research strand was to challenge the orthodoxy and to propose radical new approaches to security that were appropriate for the new era of coalition operations.

The three projects for TA2 were:

- Policy Based Security Management: This project investigated computing
  platform-independent policy frameworks to specify and analyse security
  and networking policies. The aim was to provide easy-to-use mechanisms
  for refining high-level user-specified goals and decisions into low-level
  controls, such as networking firewalls. The project developed algorithms
  to detect policy conflicts and investigated strategies for conflict resolution
  in Cols.
- Efficient Security Architectures and Infrastructures: This project developed and analysed lightweight and adaptive security architectures and infrastructures to facilitate the formation of, and operations by, secure, flexible Cols. A focus area of this project was dynamic trust establishment among various members of Cols by taking into account both positive and negative evidence. In addition, it explored alternatives to traditional public key infrastructures that are inherently more energy and bandwidth efficient, and that promise to provide natural support for coalition operations.
- Trust and Risk Management in Dynamic Coalition Environments: This
  project developed a trust and risk management framework that could be
  used to define and manage the concepts of trust, risk, and operational
  benefits in dynamic coalition environments. It also investigated, developed
  and validated mechanisms for assessing risk and benefit during operations.

#### TA3: Sensor Information Processing and Delivery

Situation awareness superiority can provide tremendous tactical and strategic advantage to coalition forces over their adversaries, especially in cases of asymmetric warfare in urban contexts. Within the context of this technical area, sensor networks and the data they generate were viewed as powerful tools that aid in achieving situational awareness, and supporting context-aware decisionmaking and other high level military operations.



Figure 5. TA3: Sensor Information Processing and Delivery

The three projects for TA3 were:

- Quality of Information of Sensor Data: This research project studied formalisms to describe, analyse and estimate the quality of information delivered by a sensor network, i.e. knowledge of the quality of information, derived from different data sources, and expressed through a rich meta-data set. This could include representations of the raw sensor data and quantified knowledge (e.g. unreliable, sufficient, superior), designed to allow decision making entities to appropriately weigh this enriched information and so make better decisions.
- Task-Oriented Deployment of Sensor Data Infrastructures: This project investigated algorithms, architectures and procedures to aid in building data source management support, taking into account troop and data source mobility across a geographical area. It explored techniques for data sources and fusion elements to be proactively deployed and operated to automatically augment the information gathering experience, for improved situational awareness during the execution of a task.
- Complexity Management of Sensor Data Infrastructures: This project investigated techniques to reduce the complexity of managing sensor data infrastructures, including changing the operating point of a sensor

network in real-time in response to ever changing and fluid mission goals. It investigated a simplified control interface for managing the multitude of disparate sensing and processing elements, finding their optimal operating points without disrupting the network operation.



## TA4: Distributed Coalition Planning and Decision Making

Figure 6. TA4: Distributed Coalition Planning and Decision Making

Coalition operations are very different from unilateral operations. They include a number of organizational, operational, political, and cultural challenges that are unique to the operational environment, and to the specific composition of the coalition.

Organizationally, coalitions can be very complex. They often include various branches of the native and foreign military, as well as non-military organizations such as *non-governmental organizations* (NGOs) and *private voluntary organizations* (PVOs), each bringing their own unique cognitive and collaborative styles (culture), doctrinal languages, technological capabilities and core competencies. This blending of capabilities makes possible certain operations that a single coalition member could not, or would not, conduct unilaterally.
The three projects for TA4 were:

- Mission Adaptive Collaborations: The essence of this project was to develop tools, methods and techniques to analyse and synthesize coalitions of agents (human and synthetic), to discover how teams can understand one another, and how they can most effectively adapt and redirect themselves.
- *Cultural Analysis:* This project focused on developing an understanding of the command processes within coalitions as well as the processes at work determining their external interactions. It analysed the cognitive and sociocultural factors that facilitate or impair communication and understanding. The aim was to provide the means to monitor and, if necessary, transform those processes.
- Shared Situation Awareness and the Semantic Battlespace Infosphere: This project researched tools and methods for understanding situations unfolding in a distributed environment. The effort to enhance situation awareness also integrated a range of planning and decision making services.

#### TA5: Coalition Interoperable Secure and Hybrid Networks

Defined in the final five years of the programme, this TA aimed to develop the fundamental underpinnings for secure hybrid wireless networking that enables adaptable and interoperable communication and information services for military coalition operations. Key challenges included efficient and adaptive secure networking that adapts rapidly to dynamic missions and ad hoc team formation, operates without reliance on centralized network or security services, and the composability of disparate security and networking systems. An underlying challenge was to treat networking and security together rather than as separate technologies.

The three projects for TA5 were:

- *Wireless Networking Performance, Metrics, and Estimation.* Developed the mathematical abstractions, models, and frameworks to enable adaptive control of the behaviour of hybrid wireless networks.
- Agile and automated Security/Network Management and Control. Developed integrated security and network management and control techniques that reduce the human expertise required to maintain dynamic hybrid networks.

 Security for Dynamic, Distributed Coalition Data and Network Services.
Developed the fundamental underpinnings for adaptable security that seamlessly and quickly changes to support the rapid assembly of CoIs that have different technologies, policies, levels of trust, and roles.





#### TA6: Coalition Information Processing for Decision Making

This TA—also defined for the final five years of the programme—aimed to develop the fundamental underpinnings for exploiting and managing an agile network of data and information sources, for effective understanding and decision-making across a coalition for dynamic complex problems. The overall challenge was to develop the fundamental science to underpin a 2-way end-to-end socio-technical chain reaching from data-to-decision (and from decision-to-data) resolving complex problems while operating in a coalition environment. The requirements included provisioning of information in a format that is understandable, augments cognitive performance, and avoids data deluge.



Figure 8. TA6: Coalition Information Processing for Decision Making

The three projects for TA6 were:

- *Human Information Interaction:* Developed fundamental mathematical representations of human-machine collective cognition and representational frameworks to enable rich human-machine interaction.
- Distributed Coalition Services: Developed fundamental principles for dynamic composition of services across a coalition hybrid network with constrained network resources to meet user needs.
- *Collective Sensemaking Under Uncertainty:* Develop integrated techniques to enhance collective reasoning and understanding in complex situations with uncertain information.

## Promoting Cross-Disciplinary Collaboration

Although the programme was structured along fairly well delineated technical areas (TA1–4 for the first half, and TA5–6 for the second half), there was a strong desire to promote cross-disciplinary research that would enable synergistic exploration by cutting across traditional discipline boundaries. The challenge for a fundamental science programme was to determine the right way to attain this without impacting the depth and excellence of scientific exploration in a single research discipline.

The approach used within the ITA programme was to define broad cross-disciplinary

research *themes* and *grand challenges* that fostered such collaboration. By defining a shared research challenge that all areas needed to address jointly, and thinking about ways to address that challenge in boot camps and collaboration meetings, the researchers were motivated to align their work in a cross-disciplinary manner. The themes also helped to promote dialogue, but had the additional task of showing relevance to the needs of coalition operations. Each of the cross-cutting themes addressed a challenge that coalition operations faced on the ground, and researchers were encouraged to think about how their research addressed those.

Another approach used to bring researchers from different disciplines together was the definition of a common scenario to motivate the research and resulting application of that research. The definition of common scenarios enabled each researcher to discuss how their scientific advances would contribute to its specific context.

In this section we provide a brief overview of the grand challenges, themes and scenarios that were used for this purpose.

#### **ITA Grand Challenges**

The grand challenges provided a means to promote collaboration and synergy across multiple TAs. Teams with expertise in different areas came together to discuss how they could join forces to meet one of the grand challenges defined. As a basic research programme, the grand challenges did not force different projects to co-develop any system or prototypes. However, they were useful in their role of encouraging conversation across multiple areas.

The ITA programme defined three grand challenges as follows:

*Challenge 1:* To ensure that coalition warfighters at the tactical level get the right information at the right time even when they do not know they need it, and even if this requires the network to mediate access to information based on resources, risk, and tactical context.

*Challenge 2:* To enable coalition warfighters at the tactical level to work collaboratively to share knowledge, build trust, and solve problems across space and across cultural boundaries and to enable this to happen rapidly.

*Challenge 3:* To enable the design of systems that are secure, dependable and flexible; parsimonious in the use of all resources, whether this is the number of nodes required to provide a capability or electrical power, processor power and spectrum bandwidth; and allow the user to make full use of the system capability

without any need to understand the complexity of the system or the detail of security/information management policies and their implementation.

### The ITA Themes

The purpose of the ITA themes was to demonstrate the value of the research programme to military stakeholders, and to ensure that the ITA was addressing real problems faced by the US and UK defence communities. To do this three themes were defined that crossed all technical areas. Each theme identified a specific issue that was of relevance to the military, and could be used to demonstrate how research projects addressed that issue.

Theme 1, Information Flows: Information sharing is a key component of networkcentric operations (NCO), providing both the right information to the right users at the right time (even though they may not know that they need it) and also the means for users to share and use information. For the soldier on the ground this is a hard problem since he/she is often dependent upon ad hoc mobile communications networks and will most likely be part of a coalition force with national differences in equipment, security policies and culture.

The key aspects of information flows covered by this theme included specifying user information needs; characterizing information; identifying the information sources best able to meet user needs; efficiently and securely establishing and sharing network flows, from sources to users; and dynamically adjusting information flows in response to changes in the user's information needs and network dynamics.

Theme 2, Dynamic Mission-Focused Communities of Interest: A CoI is a smaller subgroup of a coalition force that is formed to undertake a mission. To be effective, such subgroups require networks that provide robust operation, small world properties, security properties, and efficient use of bandwidth. Several issues were identified, such as enabling multiple secure communities to coexist, with each engaged in different missions; enabling effective community formation in a decentralized manner; and enabling collaboration among members and relationships to evolve locally.

Theme 3, Trust and Risk-Enabling Risk-Based Decisions: The visions for *network-enabled capability* (NEC) and NCO embrace the concept of providing more flexible access to information in response to operational needs. To do this requires a shift in security policy to risk-based decisions. Trust is an important factor in evaluating risk. There were a range of issues to be addressed from

conceptual concerns to system design and usability, including technologies for evaluating trust and risk; defining meta-data to enable dynamic risk evaluation; and expressing operational intent and enforcing risk based access control.

#### **ITA Scenarios**

In order to bring researchers working on different tasks together, motivating scenarios were defined during the course of the ITA research programme. The goal of these was to provide a context in which teams could discuss how a specific research activity would lead to a military capability that could help with the challenges emerging from that scenario.

The scenarios used by the ITA programme envisioned a hypothetical country in which coalition operations were being performed. In addition to defining the country profile, the scenarios defined various vignettes, such as a dirty bomb falling into the hands of insurgents, a water borne disease outbreak in the countryside, or a hostage in need of rescue. In each of the vignettes there was a discussion on the challenges coalition forces would face in dealing with its critical requirements, and how the different technologies could play a role.

These discussions—which happened in both formal meeting venues like the Boot Camps and informally over telephone conferences—helped researchers consider the relevance of their work to coalition operations, and seek opportunities to bring in technologies from other researchers to help in the effort.

All three activities—grand challenges, themes and scenarios—thus played a contributory role in bringing researchers from different areas together in cross-disciplinary interactions.

# How it all Worked

As the ITA programme progressed through its first five-year phase this model began to bear fruit and build confidence between the research teams, and thus engendered a positive attitude to an interdisciplinary approach to research. The second five-year phase saw a merging of technical topics and a strengthening of cross-topic working. Additionally, encouraging the use of a common experimentation framework (that allowed the solutions developed to be made available between teams) assisted in establishing an environment in which both results sharing and, more importantly, shared thinking about problem solving could be facilitated. Out of these changes the number of activities involving cross-disciplinary work has grown and its immense value has been recognized, becoming a success factor of the programme. It is important to note that the amount of time and effort that needs to be invested by all parties in bringing this about should not be underestimated, and this cross-discipline collaboration should be cherished as a lasting legacy of the ITA.

The success of the programme from a fundamental research perspective can be judged from the number of technical publications (in excess of 1200) on a wide variety of topics (see "NIS ITA Metrics and Publications" on page 143). The external peer review process consistently recognized the high quality of the basic research being undertaken, and many of the conference publications were recognized with best paper awards.

The technical chapters of this book highlight some of the areas of fundamental science that have been developed during the ten years of the NIS ITA programme. In the context of research-to-capability, the question was how fundamental discoveries that result in new disruptive technical capabilities can transform the way in which future operations might be performed. Whilst such discoveries often take longer to mature, fundamental breakthroughs in areas such as computing on encrypted data and verifiable outsourced computation are beginning to show the potential to provide capabilities that, until recently, were thought to be impossible. Other examples include the integration of a coalition policy management enforcement point into a fully decentralized database, and the use of that database to provide a secure distributed registry for an embryonic, self-managing tactical service bus (the *ITA Information Fabric*, now available as the open source *Edgware Fabric*).

The chapter "Exploiting the Science" describes a number of successful transition activities that have been undertaken under the auspices of the ITA programme. They range from technical demonstrators that seek to apply the results of the fundamental science research in practice (e.g. Controlled English and Network Tomography) through to full-blown operational systems into which these emerging technologies have been embedded. This type of transition is exemplified with respect to the science of *Dynamic Distributed Federated Databases* (DDFDs), where a UK investigation into the use of DDFDs for counter-IED led to the development of a novel tool that has been deployed in an operational system. This in turn enabled affordable development and demonstration of fine-grained policy controlled access operating on the *Battlefield Information Collection and Exploitation* (BICES) network at the NATO Intelligence Fusion Centre (NIFC) in a joint UK/US applied research programme.

# 2 Network Infrastructure for Coalition Operations

Over the ten years of its existence, the NIS ITA programme made several significant advances in understanding the fundamental properties of coalition tactical networks. These include: creation of mathematical models to identify the performance and scalability limits of network infrastructure, new algorithms and approaches to address communications challenges in coalition networks, as well as insights into new approaches for routing and information-dissemination in coalition environments.

During the initial phase of the research programme, significant attention was paid to the network infrastructure for MANETs since they were viewed as the primary structure for tactical units in a coalition environment. In the latter phase of the programme, the focus shifted toward hybrid coalition networks—which included MANETs as well as fixed infrastructure networks—to obtain the best possible network connectivity for coalition operations.

Some of the key scientific advances in understanding the properties of coalition networks are enumerated in this chapter.

# Network Tomography

One of the key challenges in coalition operations is that any one partner only has limited visibility into the structure and performance of the networks belonging to other partners. Since networks form a critical component for any military operation in network-centric warfare, and communication in coalition networks frequently requires traversing coalition partners' infrastructure, mission effectiveness can be impacted by unpredictable network performance. Network tomography provides the principles that enable insight into the performance of a coalition partner's network infrastructure. The word tomography comes from the field of medical imaging, where it is narrowly used to describe procedures that create images by identifying different sections, and more broadly as inferring the internal structure of the body by means of external measurements. In network tomography, an analogous technique is used to understand the attributes of the nodes and links that make up the coalition network, and different sections understood in this manner are composed to create the overall structure of a large part of the network. Network tomography thus consists of understanding the attributes of the components (nodes, links) that make up a network by means of end-to-end measurements.

The ITA work in network tomography was conducted collaboratively by researchers at ARL, Dstl, Imperial College London, University of Massachusetts at Amherst (UMass) and IBM US. While the basic concept of network tomography was known in the academic community, a fundamental theory to relate the number of linearly independent paths (and thus link identifiability) to externally observable parameters such as network topology, number of monitoring nodes, and routing restrictions was not available. The ITA team established necessary and sufficient conditions for identifying all the link metrics using a given number of monitors in a network [1]. They showed that these conditions not only allow efficient identifiability tests, but also enabled an efficient algorithm to place the minimum number of monitors in order to identify all link metrics. To the best of our knowledge, this is the first work providing fundamental constraints on network topology for identifying additive link metrics using end-to-end measurements on cycle-free paths.

The team began network tomography research by investigating the problem of identifying individual link metrics in a communication network by measuring accumulated end-to-end metrics over selected paths, under the assumption that link metrics are additive (e.g. delay) or multiplicative (e.g. packet delivery rate) and constant for the duration of the measurement.

Using principles from the field of linear algebra, the team showed that all the link metrics could be uniquely identified when the number of linearly independent paths is equal to the number of links in the network. Figure 9 illustrates a graph in which this condition is satisfied. The network has 3 monitoring points located at the nodes in red, has 9 links and 9 independent end-to-end paths. The first main result was that it is impossible to identify all the link metrics in any network with a non-trivial topology (having more than one link) using only two monitoring nodes. Nevertheless, the interior links not incident with any monitoring node might be identifiable. The second main result was a set of necessary and sufficient

conditions for identifying all the interior links using two monitoring nodes. Furthermore, it was shown that these conditions have a natural extension to identifying the entire network using three or more monitoring nodes [1]. These conditions were expressed in terms of the 3-vertex connectivity of an extended graph. The proposed algorithm for optimal monitor placement and path construction has surprisingly low complexity, O(n+m), n being the number of nodes, and m the number of links.



Figure 9. Network satisfying criteria for additive link tomography

Another important problem in tomography is that of placing a given number of monitors in a communication network to identify the maximum number of link metrics from end-to-end measurements between monitors, assuming that link metrics are additive, and measurement paths cannot contain cycles. Previous results had shown that complete identification of all link metrics could require a large number of monitors. To address this challenge, the team developed an efficient algorithm for determining all identifiable links for a given placement of monitors [2]. This allowed the creation of a polynomial-time greedy algorithm to incrementally place monitors such that each newly placed monitor maximizes the number of additional identifiable links. A snapshot of links that can be identified in a network using this algorithm with different numbers of monitors is shown in Figure 10. The team proved that the proposed algorithm was optimal for 2-vertexconnected networks, and demonstrated that it is near optimal for several real internet service provider (ISP) topologies that are not 2-vertex-connected. The research enabled a quantifiable trade-off between the level of identifiability and available monitor resources.

A further contribution was to identify the location of failed nodes when the only measurements available were from end-to-end probes. Assuming that a path

behaves normally if and only if it does not contain any failed nodes, the measured paths must show different symptoms under different failure events. In other words, for any two distinct sets of failed nodes, there must be a measurable path traversing one and only one of them. This condition is, however, impractical to test for large networks due to the combinatorial numbers of paths and failure sets.



Figure 10. Greedy algorithm performance with number of monitors

The cross-organizational ITA team characterized this condition in terms of easily verifiable conditions of the network topology and the placement of monitors under three families of probing mechanisms. These differ in whether measurement paths are: (i) arbitrarily controllable, (ii) controllable but cyclefree, or (iii) uncontrollable (i.e. determined by the default routing protocol). They also characterized the maximum identifiability of node failures, measured by the maximum number of simultaneous failures that can always be uniquely localized. They identified both upper and lower bounds on the maximum identifiability and showed that algorithms can be constructed with polynomial running time to find bounds such that the upper and the lower bounds differ by at most one. They also quantified the impact of the probing mechanism on the capability of node failure localization by numerically comparing the maximum identifiability under different probing mechanisms on a variety of random and real network topologies. They found that despite a higher implementation cost, probing along controllable paths could significantly improve a network's capability to localize simultaneous node failures [3].

The team proposed a framework to design probing experiments with a focus on probe allocation when the link metrics are stochastic, and used it to apply tomography to two types of performance metrics—packet loss and packet delay variation. They showed that when the number of probing paths equals the number of links, closed-form solutions for the optimal design exist. They also developed efficient heuristics to optimally distribute probes in networks where a closed form solution is not possible [4].

These results have a profound implication on the management and operation of coalition networks, since they enable better insights into the performance of partner networks that cannot be directly monitored. Network tomography provides tools for automated management of coalition network services.

### **Energy Efficient Networking**

Battery power is a significant challenge in tactical networks, along with the fact that parts of a coalition's network infrastructure can disappear abruptly. Cooperative networking, including use of network coding, has been proposed as a mechanism to significantly increase the robustness of message delivery and network performance. A team of ITA researchers from UMass and Imperial College London pursued research on the issues of cooperative networking. They analysed its properties, and proposed various mechanisms to use cooperative networking to improve power consumption in coalition networks.

After analysing approaches to cross-layer optimization to cooperatively forward information in MANETs, the team developed cross-layer algorithms incorporating network coding and showed that cooperative diversity can significantly improve the power-performance curve [5]. In particular, Figure 11 shows the significant reduction of power consumption using cooperative transmission to achieve a given transmission reliability (right-hand figure), when compared to the case without the use of cooperative transmission (left-hand figure). The power consumption can be used to determine a soldier's power load as a function of mission lifetime. The team also augmented these studies with an analysis of robustness of network coding schemes.

The team considered the exploitation of the available multi-user diversity in wireless networks to greatly improve the performance and overcome problems of algorithms typically viewed as belonging to the physical layer. They considered the recent promising technique of *physical-layer network coding*. By looking up one layer and performing relay selection with cognizance of this underlying technique, they were able to both improve the performance and overcome the onerous distributed phase synchronization requirements. Furthermore, a similar line of thought was applied to a multiple access channel with available relays. By

carefully scheduling the multiple source and relay transmissions in cooperative networks, the full diversity gain is achievable even though only a fraction of the relays is made available to each source. An achievable diversity-multiplexing trade-off was established that approximates the optimal multiple-input singleoutput upper bound to such [6].



Figure 11. Reduction in power attainable with cooperative networking

An alternative efficient technique to save energy consumption in sensor networks is the intelligent use of a duty cycling mechanism—turning sensors on and off for communications. An ITA research team from University of Cambridge and Raytheon BBN Technologies established a tight lower bound on the transmission latency in a duty-cycling network and proved the latency optimality of a new protocol, the *green wave sleep scheduling* (GWSS) protocol on selected network topologies (line, grid, and tree) with low load. This work proved the NP-completeness of the delay-efficient sleep scheduling problem using a polynomial time reduction of the 3-SAT problem, constructed a non-interfering version of GWSS with slanted *wave-fronts* that achieves optimal latency and throughput performance on a square-grid network with multiple simultaneous transmissions, and suggested an extension of GWSS to random extended networks using the construction of an underlying grid of percolating paths across the network area [7].

In networks that deploy duty cycling techniques to conserve energy, gathering global routing topologies can be cumbersome. In such networks, stateless opportunistic forwarding provides an alternative approach that can be more practical. The research team performed an analytical investigation of latency properties of stateless opportunistic forwarding in finite networks. Modelling the network behaviour as a random walk on a weighted graph with link-specific sojourn times, the team investigated the latency incurred by such a random walk. They were able to derive an exact formula for end-to-end latency. This enabled them to derive simple and reasonably accurate scaling laws for latency in certain types of regular graphs, as well as for several random geometric graphs [8].

### Fundamental Performance Limits of Hybrid Networks

In order to design coalition communication networks, an insight into their performance limits and scaling properties is essential. A team of ITA researchers from University of Cambridge, UMass and ARL determined various performance limits on the rate at which information can be exchanged between nodes in a hybrid network.

The team modelled a network as being comprised of n ad hoc nodes scattered randomly and b base stations placed regularly, where the ad hoc nodes can share information through the wired infrastructure in addition to using wireless ad hoc communication. The team studied per-node throughput scaling achievable in this hybrid network as n (and b) grow for both one-dimensional and two-dimensional networks that are fully connected [9]. The work was further extended to analyse more realistic networks with finite bandwidth links and limited interconnections. Upper bounds on the throughput were found, which match the lower bounds for a wide range of values of b. The team also improved previous lower bounds in the extreme case where b is almost of the same order as n, and also showed a matching upper bound in the case of one-dimensional networks [10].

Since tactical networks are dynamic, and often allow packets to be stored and forwarded to the destination, a cross-organizational ITA team from UMass, ARL and Raytheon BBN Technologies analytically studied the probability of source routing success in a class of dynamic networks, where a time-varying stochastic process governs the link (up-down) dynamics. Their key finding was that this exhibited critical phase-transition phenomena (also known as percolation) as a function of the end-to-end message latency per unit path length (or inverse velocity). They evaluated the probability of routing success on dynamic network (1D and 2D) lattices with links going up and down as per an arbitrary binary-valued stationary random process (such as a Markov process), in a source-routing framework. They determined percolation thresholds on the time deadline for high-probability of routing success in terms of the first and second order moments of the link state process, and determined percolation thresholds on the parameters characterizing the link process for a fixed time deadline for a 1D Markov network.

That work generalized results reported in two articles from the 80s that appeared in the *Physical Review Letters* on directed percolation theory. The team analysed the performance of a stateless single-copy opportunistic forwarding algorithm on a 2D probabilistic grid. They discovered that adding a time dimension, i.e. letting the network evolve as per potentially time-correlated link dynamics, the opportunistic store-and-forward routing algorithm exhibits a critical threshold behaviour that can be determined. In this 2D grid network case (as in the 1D network case), the normalized messaging latency (ratio of routing latency to path length) exhibits a critical phase transition when one evaluates the critical latency to path-length ratio as a function of the moments of the link up-down process [11]. This critical phase transition is illustrated in Figure 12, where percolation of normalized message latency happens at a precise critical threshold q/(p(p+q)) for a linear network with a two-state Markov link-dynamics model parameterized by probabilities p (off to on) and q (on to off). These insights have been used in various protocols to set the time to live TTL parameter of data packets in order to meet a timeliness/storage trade-off.



Figure 12. Critical phenomenon in linear networks

An ITA research team comprised of researchers from IBM US and University of California, Los Angeles (UCLA), studied the scaling properties of disruption tolerant networks (DTNs) [12]. DTNs have been proposed as one of the fundamental building blocks of tactical networks. They effectively overcome partial connectivity by letting the nodes carry-and-forward data; and, understanding the scalability of DTN protocols is needed for protocol design and evaluation. The team developed a unified framework that formalizes DTN performance as a function of node mobility for the first time. In this work, they represented DTNs as a class of wireless mobile networks with intermittent connectivity, where the intercontact behaviour of an arbitrary pair of nodes can be described by a generalized two-phase distribution (i.e. a power-law head with an exponential tail). Various experiments and analytical results have shown that this two-phase distribution represents the most realistic model for vehicular and pedestrian mobility. Using this DTN model, they extended the throughput and delay scaling results of Grossglauser and Tse (derived for exponential inter-contact time distributions) to the general case with motion correlation. They also analysed the impact of finite buffers on the capacity scaling properties of DTNs under the two-phase mobility model, and verified the analytical results with a simulation study on a realistic wireless network model.

A team from Imperial College London and IBM US also studied the placement of static gateways in mobile DTNs [13]. Given a limited gateway budget, the problem is to deploy gateways at selected locations such that certain performance metrics are optimized. Different domains may possess heterogeneous properties (such as node mobility patterns, message forwarding strategy, network size), which have to be considered jointly to optimize the end-to-end performance. The team came up with a unified gateway deployment framework, which separates the domain-specific utility computation from the high-level gateway placement. To obtain an efficient solution, they developed practical algorithms with quadratic complexities along with analytical utility decomposition and calculation methods. They derived closed-form analytical expressions to calculate the utilities as functions of several key characteristics of the constituent domains, which were shown to be accurate and robust against the mobility patterns. Evaluation based on the contact traces of a real DTN shows that the calculated utility is very close to the actual value after a constant scaling, and the final deployment performs as well as that from the optimal deployment strategy with a much higher complexity. Compared with deployments that are agnostic to utilities, the proposed strategies improve the performance by up to 30%.

# **Robust Routing in Coalition Networks**

Tactical coalition networks face a variety of challenges when it comes to routing data between various nodes. Network topology and link characteristics vary over time. Ensuring end-to-end packet delivery at low latency is a task for which traditional stateful routing algorithms that maintain a globally consistent single path between each source-destination pair cannot perform well.

In order to address this challenge, ITA researchers have come up with various approaches to improve routing in coalition networks. One approach that a team of researchers from Roke Manor Research Ltd. in the UK and UMass explored was *braided routing* [14].

A braid of links and nodes, initially constructed around a set of links connecting a source/destination pair, is used for forwarding data between source and destination. Braided routing operates at two timescales. At the longer timescale, a braid is constructed with links and nodes that may change over time, and no effort is made to repair the braid in response to such changes in between braid construction epochs. Instead, the diversity of forwarding links within the braid serves to provide locally determined alternative forwarding paths. At the shorter timescale, local forwarding decisions are made to select the *best* next hop out of all possible next hops. With this limited-scope, multipath forwarding in braided routing can explicitly and efficiently control the trade-off between state/signalling overhead and performance in light of topology changes. The team analysed braided routing from several different viewpoints in order to fully explore and understand its properties [15].

Another challenge in coalition routing is the frequent loss of inter-domain connectivity due to MANET topology changes. The challenge is illustrated in Figure 13 where the link to the coalition partner network at the top has been severed, disrupting communication. A collaboration between UCLA, IBM US, University of Cambridge, and Honeywell International Inc. explored techniques to address this problem.

The team proposed an *Inter-Domain Routing protocol for MANETs* (IDRM). IDRM was the first protocol to address routing across multiple independent MANET domains [16]. In subsequent work [17,18], they came up with a distributed algorithm to dynamically activate gateways to support inter-domain communication in heterogeneous MANET environments with dynamic network topology. The team designed a protocol to support opaque interoperation among heterogeneous MANET domains. They demonstrated that under various

conditions, the new protocol performed a fully distributed dynamic gateway assignment that significantly improved performance (up to 200% compared to a static case where the role of gateways have been pre-assigned depending on the scenario) by adaptively assigning gateways in response to topology changes. They analysed the overhead of the protocol, explicitly considering node mobility, both by modelling and simulations.



Figure 13. Disruption in coalition networks motivating the creation of a new Inter-Domain Routing protocols for MANETs

Routing in networks requires collecting link state, which is a challenge in timevarying networks. Researchers from Raytheon BBN Technologies, Roke Manor Research Ltd. and UMass performed an analytical study of strategies for gathering link state in wireless networks whose structure varies over time. In such timevarying networks, optimal routing computations need to be performed frequently, but it is often infeasible or expensive to gather the current state of links for the entire network all the time.

The team performed an analytical characterization [19] of the effect on the performance of the minimum-latency routing policy of various link-state sampling strategies operating under a limited sampling budget in a special class of dynamic networks. They modelled link dynamics using a two-state Markov link-dynamics model parameterized by probabilities p (off to on) and q (on to off). If links are more likely to turn on than off at each time instant (i.e. p>q), a *depth-first* sampling strategy is optimal. If links are more likely to turn off than on (p < q), a *breadth-first* sampling strategy is negligible compared to the time scale of the link dynamics. The team also presented numerical simulation results comparing

optimal schedules under different latency assumptions, e.g. when the packet forwarding latency is comparable in time-scale to the link dynamics.

# Multipath Control of Hybrid Networks

With the ubiquity of cellular networks in most areas of the world, a hybrid network consisting of cellular infrastructure as well as MANETs will be a reality for many coalition operations which operate in an urban context, e.g. peacekeeping operations in towns where an insurgency has erupted. The ability to use multiple paths spanning cellular infrastructure and MANET hops provides for resiliency using path diversity. Cellular networks can provide broad, long-range coverage and potentially high quality mobile connectivity for the warfighter, using relatively inexpensive standardized commercial components. Cellular technology not only enables mobility, but also allows redundant connectivity using multiple wireless paths to improve availability, reliability, and performance.

Multiple paths enable the potential to shift traffic from broken or congested paths to higher quality ones as traffic characteristics dynamically change, particularly during movement. A team of researchers from University of Cambridge, IBM US and UMass characterized the behaviour of cellular networks, and their utility as a hybrid platform, by examining 3G, 4G, and WiFi networks for both single-path and multipath data transport [20]. They first characterized two major US 4G/3G networks in terms of throughput, round trip time, and loss rate. Then, they conducted experiments of single-path and *multipath TCP* (MPTCP) connections over these cellular data networks, and showed that data transport using MPTCP is a promising solution for a more reliable and efficient data transfer scheme in dynamic environments. In addition, they identified cases where enabling MPTCP might not provide high throughput gain. These situations arise when file size is small and a particular path has significantly higher quality.

MPTCP enables mobile devices to use several physical paths simultaneously through multiple network interfaces, such as WiFi and cellular. However, wireless path characteristics change frequently in mobile environments, causing challenges for MPTCP: for example, WiFi associated paths often become unavailable as devices move, since WiFi has intermittent connectivity caused by the short signal range and susceptibility to interference. The researchers improved MPTCP to manage path usage based on the associated link status [21]. This variant, called MPTCP-MA, uses MAC-Layer information to locally estimate path quality and connectivity. By suspending/releasing paths based on their quality, MPTCP-MA can more effectively utilize restored paths.

Experiments conducted using the setup shown in Figure 14 showed that MPTCP-MA can efficiently utilize an intermittently available path, with WiFi throughput improvements of up to 72%. In a related effort, the team also designed and implemented an improved energy-efficient MPTCP that reduces power consumption while preserving the availability and robustness benefits of MPTCP.



Figure 14. MPTCP experimental setup: 2-path (solid), and 4-path (solid and dashed)

#### Multicast in Hybrid Coalition Networks

Multicast has been well studied in the networking literature in both wired and wireless network contexts in the last several decades. However, multicast routing under physical or logical communication constraints, as is typical in coalition environments, has not received much attention. A team from Raytheon BBN Technologies, University of Cambridge, and The City University of New York (CUNY) considered the multicast routing problem under operational communication constraints found in coalition environments [22, 23].

The team initially considered the problem of minimum-cost multicast routing on a multi-domain network by constructing a node-weighted Steiner tree (for mesh networks) and a Steiner connected dominating set (for wireless broadcast networks) that is subject to a non-additive cost constraint. This is because the multi-domain multicast cost is not just the sum of node costs of a Steiner tree, but it instead depends on the domains of the connected neighbours. They found an efficient algorithm that provides a tree with bounds on how far it is from the optimal solution, which is computationally hard to determine, and showed that taking multi-domain cost constraints into account can help reduce the cost of a multicast tree by up to 40%.

The team also considered creation of such trees when constraints are imposed due to a hierarchy that needs to be maintained among various nodes. One such scenario is shown in Figure 15. Such constraints may arise in a subset of nodes belonging to different coalition partners. They showed that the overall multicast problem can be decomposed into several smaller multicasts, each of which can be solved using the previous approach. They found that necessary hierarchical constraints could cause a significant increase in the total cost of multicast – up to 25%, via simulations on realistic military networks.



Figure 15. Multicast on physical network constrained by command and control hierarchy

In other work, the team envisioned large-scale group communication (or multicast) applications including real-time voice call groups in disaster relief and military hybrid networks, video conferencing (e.g. smartphone video conferencing), P2P video and file sharing. On one hand, multi-hop wireless networking over WiFi can help extend the range of cellular networks in low *signal-to-interference-plus-noise ratio* (SINR) regions as well as alleviate network congestion. On the other hand, equipping a few nodes in a MANET with cellular radios can help to heal wireless network partitions and, thus, to improve the overall network connectivity. The team considered the problem of resource-efficient multicast in hybrid wireless networks that include both point-to-point (cellular) and broadcast (MANET) links, as shown in Figure 16.

The underlying optimization problem is a hybrid of two well-known NP-hard

graph optimization problems—the *Minimum Steiner Tree* problem (for point-topoint links) and the *Minimum Steiner Connected Dominating Set* problem (for broadcast links). They considered both edge- and node-weighted versions of this problem and used distinctly different methodologies to give three algorithms with guaranteed approximation factors. Under different assumptions about the underlying network, one of the three algorithms can provide a practical solution to the optimization problem.



Figure 16. Graph representation of a cellular/MANET hybrid multicast problem

The team also studied the problem of maximizing the multicast throughput in a dense multi-channel multi-radio (MC-MR) wireless network with multiple multicast sessions. Specifically, they considered a fully connected network topology where all nodes are within transmission range of each other. In spite of its simplicity, this topology is practically important since it is encountered in several real-world settings. Further, a solution to this network can serve as a building block for more general scenarios that are otherwise intractable. For this network, they showed that the problem of maximizing the uniform multicast throughput across multiple sessions is NP-hard. However, its special structure allows derivation of useful upper bounds on the achievable uniform multicast throughput  $\mu$ . In particular, they showed that  $\mu \leq \min\left(\frac{T}{d_{max}}, \frac{c}{\kappa}\right)$ , where T is the number of transceivers,  $d_{max}$ is the maximum number of multicast sessions in which a node can participate. C is the number of channels, and K is the number of concurrent multicast sessions. They showed that an intuitive class of algorithms that maximally exploit the wireless broadcast feature can result in very poor worst-case performance. Using a novel group splitting idea, they designed two polynomial-time approximation

algorithms that are guaranteed to achieve a constant factor of the throughput bound under arbitrary multicast group memberships. These algorithms are simple to implement and provide interesting trade-offs between the achievable throughput and the total number of transmissions used for multicast in a coalition environment.

## Mobility Modelling in MANETs

A team drawn from Raytheon BBN Technologies, University of Cambridge, and UMass developed an approach for *Universal Mobility Modelling* [24] that covered two fundamental aspects: (i) generating traces from models, and (ii) extracting models from traces. To generate traces from the model, the team proposed developing mobility models using a set of composable building blocks. The team observed that any mobility pattern can be captured in terms of three fundamental blocks: (i) target selection, (ii) steering behaviours, and (iii) path planning/ navigation. Based on this observation, the team studied the decomposition of mobility scenarios into a limited set of primitive mobility building blocks, from which succinct mobility models can be defined.

A team from University of Cambridge and UMass [25] analysed a novel set of mobility traces collected from a military experiment carried out in Lakehurst, New Jersey, U.S. In this dataset, a number of soldiers navigate through a highly hostile terrain where they were susceptible to adversities such as ambushes and traps. In the hope of maximizing their security, a common practice is to traverse the area along the same path, one squad after another. These squads depart along the same path only when the frontal squad arrives at a confirmed secure point. The frontal squads may also have to rest and wait for the following squads for supply and back up. During the movement, if a squad were in danger, squads nearby would immediately move to (or, of course, near if it is unsafe) the scene for assistance. Such mobility patterns are common in military settings.

The team divided the data set into two sets: a training set and a test set. They analysed the mobility characteristics of the training set, and used that to generate synthetic data according to the values of the mobility characteristics. They compared the synthetic set to the test set, and found that the mobility model worked well and accurately captured aspects of the spatial and temporal characteristics of the test set. This mobility model can be used to generate synthetic traces with different travel schedules.

### References

- [1] L. Ma, T. He, K. K. Leung, A. Swami and D. Towsley, "Inferring link metrics from end-to-end path measurements: Identifiability and monitor placement", *IEEE/ACM Transactions on Networking*, vol. 22, no. 4, August 2014.
- [2] L. Ma, T. He, K. Leung, A. Swami and D. Towsley, "Monitor placement for maximal identifiability in network tomography", in Proc. *IEEE INFOCOM*, 2014.
- [3] L. Ma, T. He, A. Swami, D. Towsley, K. K. Leung, and J. Lowe, "Node failure localization via network tomography", in Proc. ACM Internet Measurement Conference, 2014 (IMC '14).
- [4] T. He, C. Liu, A. Swami, D. Towsley, T. Salonidis, A. Bejan and P. Yu, "Fisher information-based experiment design for network tomography", Proc. ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems, 2015 (SIGMETRICS 2015).
- [5] Z. Ding, K. Leung, D. Goeckel, and D. Towsley, "On the study of network coding with diversity", *IEEE Transactions on Wireless Communications*, vol. 8, no. 3, pp. 1247–1259, March 2009.
- [6] Z. Ding, K. Leung, D. Goeckel, and D. Towsley, "A novel relay assisted cooperative transmission protocol for wireless multiple access systems", *IEEE Transactions on Communications*, vol. 58, no. 8, pp. 2425-2435, August 2010.
- [7] S. Guha, P. Basu, C.-K. Chau, and R. Gibbens, "Green Wave: Latency and Capacity-Efficient Sleep Scheduling for Wireless Networks", *IEEE Journal on Selected Areas in Communications*, vol. 29, issue 8, pp. 1595–1604, September 2011.
- [8] C.-K. Chau and P. Basu, "Analysis of Latency of Stateless Opportunistic Forwarding in Intermittently Connected Networks", *IEEE/ACM Transactions on Networking*, vol. 19, no. 4, August 2011.
- [9] N. Banerjee, M. Corner, D. Towsley, and B. N. Levine, "Relays, Base Stations, and Meshes: Enhancing Mobile Networks with Infrastructure," in Proc. ACM MOBICOM, 2008.
- [10] C. Capar, D. Goeckel, D. Towsley, R. Gibbens, and A. Swami, "Capacity of Hybrid Networks," in Proc. Annual Conference of the ITA, 2011.

- [11] P. Basu, S. Guha, A. Swami, and D. Towsley, "Percolation Phenomena in Networks under Random Dynamics", in Proc. COMSNETS, 2012.
- [12] U. Lee, S. Y. Oh, K.-W. Lee, and M. Gerla, "Scaling Properties of Delay Tolerant Networks in Correlated Motion Patterns", in Proc. ACM CHANTS, 2009.
- [13] T. He, N. Sofra, K.-W. Lee, and K. K. Leung, "Utility-Based Gateway Deployment for Supporting Multi-Domain DTNs", in Proc. *IEEE SECON* 2010, June 2010.
- [14] V. Manfredi, R. Hancock, and J. Kurose, "Robust Routing in Dynamic MANETS", in Proc. Annual Conference of the ITA, 2008.
- [15] C.-K. Chau, R. J. Gibbens, R. E. Hancock, and D. Towsley, "Robust Multipath Routing in Large Wireless Networks", in Proc. *IEEE INFOCOM 2011* mini-track.
- [16] C.-K. Chau, J. Crowcroft, K.-W. Lee, and S. H. Wong, "Inter-Domain Routing for Mobile Ad Hoc Networks", in Proc. ACM MobiArch, 2008.
- [17] S.-H. Lee, S. H. Wong, C.-K. Chau, S. Varadarjan, K.-W. Lee, and J. Crowcroft, "Self-organizing Inter-Domain Routing for Heterogeneous MANETs", in Proc. Annual Conference of the ITA, 2009.
- [18] Y. Song, S. Wong and K. Lee, "Optimal gateway selection in multidomain wireless networks: a potential game perspective", in Proc. ACM MobiCom, September 2011.
- [19] S. Guha, D. Towsley, P. Basu, H. Tripp, T. Freemany, D. Katz-Rogozhnikov, R. Hancock, and J. Kurose, "Optimal sampling strategies for minimum latency routing with imperfect link state", in Proc. *WiOpt*, 2012.
- [20] Y.-C. Chen, Y.-S. Lim, R. J Gibbens, E. M. Nahum, R. Khalili, and D. Towsley, "A Measurement-based Study of Multipath TCP Performance over Wireless Networks", in Proc. ACM Internet Measurement Conference (IMC), 2013.
- [21] Y.-S. Lim, Y.-C. Chen, E. M. Nahum, D. Towsley, and K.-W. Lee, "Cross-layer path management in multi-path transport protocol for mobile devices", in Proc. *IEEE INFOCOM*, 2014.
- [22] P. Basu, C.-K. Chau, R. Gibbens, S. Guha, and R. Irwin, "Multicasting under Multi-domain and Hierarchical Constraints", in Proc. WiOpt, 2013.

- [23] P. Basu, C.-K. Chau, A. Lu. Bejan, R. Gibbens, S. Guha, and M. P. Johnson, "Efficient Multicast in Hybrid Wireless Networks", in Proc. *IEEE MILCOM*, 2015.
- [24] A. Medina, G. Gursun, P. Basu, and I. Matta, "On the Universal Generation of Mobility Models", in Proc. IEEE/ACM MASCOTS, 2010.
- [25] X. Lu, Y.-C. Chen, I. Leung, Z. Xiong, and P. Liò, "A Novel Mobility Model from a Heterogeneous Military MANET Trace", in Proc. *International Conference on Ad Hoc Networks & Wireless* (LNCS), pp. 463–475, 2008.

# 3 Security for Coalition Operations

In modern network centric operations, politics and pragmatics often require military operations to be conducted in cooperation with partners that cannot be completely trusted. Notwithstanding the need to address the standard problems of secure operation in a single partner network, the complexities of dealing with multiple partners makes security much harder in coalition operations.

Because of the critical role that security has in coalition operations, a significant effort in the NIS ITA programme was devoted to address challenges that can arise. The international team of academic, industrial and government researchers addressed several of these challenges, and have made significant progress in understanding the principles that can help in improving security.

The scientific advances made by ITA researchers include: information theoretic analysis, new security algorithms, new security architectures and frameworks, disproving some commonly held myths in the scientific community, as well as optimizations and performance improvement techniques to make secure computation faster. Some of the most significant results are summarized in this chapter.

### Policy-based Security Management

In a coalition environment there is a constant need to interoperate with multiple administrative domains. A policy-based approach allows multiple coalition partners to work towards common mission objectives, while protecting sensitive information from unwarranted access. Policy based technologies were proven to be effective in a single administrative domain at the start of the programme. The primary contribution of ITA researchers, drawn from Imperial College London, IBM US, Honeywell International Inc. and ARL, was in addressing challenges in coalition settings. In order to address coalition policy challenges in a holistic manner, the research team developed an abstract model for a policy life cycle for applying security policies in a coalition environment. This life cycle is shown in Figure 17. It provides a context for different policy related activities that need to be performed at different stages of a coalition operation to support dynamic communities of interest.



Figure 17. Policy life cycle model for coalition operations

The research team developed a policy refinement framework [1] to address the challenges associated with the steps of this life cycle. They also developed algorithms [2,3] to refine policies defined in constrained natural language into enforceable security policies through a series of transformations that preserve their semantic intent, with each refinement step being: *correct*, in that the set of refined policies correctly implements the higher level policy; *consistent*, in that the refinement must not lead to conflicts between the derived policies or the other policies existing in the system; *valid*, in that the policies must be enforceable in the system context to which they will be applied; and, *minimal*, in that all policies in the derived policy set must be required for the correctness of the refinement. Often, high-level security policies describe decisions (i.e. access-control decisions or obligations) based on system-state conditions that occur prior to the state at which the decision is made. Most existing policy systems do not natively support past temporal conditions. The research team developed a methodology for transforming policies with references to historical information into history-free policies in an automated fashion. For this form of policy refinement, the main idea is not to keep the whole history but to identify, based on the syntax of the policies, what part of the history should be retained and for how long. By syntactically analysing the given policies, the refinement approach identifies a subset of the system history sufficient to evaluate the conditions. The system monitors this relevant information at each state and maintains a store of auxiliary facts. Such monitoring is done by rules generated from the syntactic analysis of the given policies, which are then transformed into stateless policies where a simple store query replaces the original history conditions. Each system state is extended with a store, and the evaluation of history conditions (hence policies) is reduced to an equivalent evaluation of current conditions. The research team has proved the correctness of the transformation, i.e., the history-free policies plus the event monitoring system enforce exactly the same policies as the original history-based policies. In this way, they have enabled execution of a general class of historybased policies [2] on end-user devices that only support history-free policies (e.g. those implementing XACML or Ponder2 policies).

The approach borrows ideas from Chomicki's prior work on efficient checking of *Past Linear Temporal Logic* (Past LTL) integrity constraints in the context of database management systems. This provided the foundation for the approach that the team used to provide complexity results on the performance of policy analysis algorithms [4].

A case study [5] was conducted to apply this theory to firewall configuration management. In particular, using argumentation logic, mechanisms were derived to automatically generate firewall policies from higher-level requirements, demonstrating the refinement process for these particular sets of policies. While the previous references dealt with a generic method of transforming temporal references for all types of policies, this reference deals with the refinement process for policies that are used by a specific security enforcement mechanism, viz., firewalls. Firewalls remain the main perimeter security protection for networks. However, network size and complexity make firewall configuration and maintenance notoriously difficult. Tools are needed to analyse firewall configurations for errors, to verify that they correctly implement security requirements and to generate configurations from higher-level requirements. This work extends previous work on the use of formal argumentation and preference reasoning for firewall policy analysis, and develops means to automatically generate firewall policies from higher-level requirements. The research team validated the approach by applying it to both examples from the literature and real firewall configurations of moderate size (150 rules).

One salient characteristic of generating firewall policy rules is that it entails generating an ordering (i.e. priorities). Rule priorities are not restricted to firewall rules. Other standard policy languages such as XACML and CIM-SPL need ordering of the rules; and, as with firewall rules, this is not a simple process. The basic algorithm proposed in this work, that combines argumentation logic with abduction, can also be used to automatically generate rule order in other policy languages.

In other advances for improving policy management for coalition operations, the research team developed algorithms and approaches that could ease the task of policy negotiation for coalition commanders. They investigated different approaches by which negotiation can be brought to a quick closure by successively reducing the amount of conflict existing between the set of policies that were being negotiated [6].

## Information Theoretic Security Capacity

ITA researchers from UMass and Imperial College London have done work to determine fundamental bounds on the capacity of information theoretic security [7,8].

The wireless secrecy capacity scaling problem asks how much information can be shared among n randomly located nodes such that the throughput is kept information-theoretically secure from m eavesdroppers also present in the network. The team found bounds on this sharing capacity for both one-dimensional and two-dimensional networks. They showed that in a 1-D network, n nodes can share a per-node throughput that scales as 1/n, which can be kept secure from mrandomly located eavesdroppers of unknown location as long as m grows more slowly than  $n/\log n$ .

Figure 18 illustrates a simplified version of this approach. Network coding is used along four paths between the source (S) and the destination (D). The paths have the same minimum spacing throughout the route; hence no eavesdropper can be close enough to all four paths at once. This allows us to determine the per node throughput regardless of monitor position. For a 2-D network, the per-node secure throughput scales as  $\frac{1}{\sqrt{n \log n}}$  for any number of eavesdroppers of unknown location, which could be arbitrarily located inside this network.



Figure 18. Sample network for analysing security capacity

These results provide a significant improvement over previous work, which assumed either that eavesdropper locations were known or the number of eavesdroppers that could be tolerated was very limited. The key technique realizing these improvements is the application of simple network coding methods, which were known to help secrecy in a network but their extension to wireless physicallayer secrecy had been limited. The researchers also derived a square root limit on low probability of detection (LPD) communication over additive white Gaussian noise (AWGN) channels. Specifically, if an eavesdropper (E) has an AWGN channel to the transmitter with non-zero noise power, they proved that  $O(\sqrt{n})$  bits can be sent from the transmitter to the receiver in n AWGN channel uses with probability of detection by the eavesdropper less than  $\epsilon$  for any  $\epsilon > 0$ , and, if a lower bound on the noise power on the eavesdropper's channel is known, then  $O(\sqrt{n})$  bits can be covertly sent in *n* channel uses. Conversely, trying to transmit more than  $O(\sqrt{n})$  bits either results in detection by the warden with probability one or a non-zero probability of decoding error as  $n \to \infty$ . Further, they showed that LPD communication on the AWGN channel allows one to send a non-zero

symbol on every channel used, in contrast to what might be expected from the square root law in image-based steganography.

They also studied the connectivity properties of coalition networks under a simple security constraint. The problem was to study connectivity in a coalition network, which includes nodes from two different teams (red and blue); here, two nodes from different teams can only communicate through gateway (purple) nodes also present in the network. The constraint that some nodes are not allowed to communicate even when they are within transmission range makes the connectivity problem more interesting but also harder to address. Their results suggest that a relatively small number of gateway nodes are enough to significantly improve overall connectivity, which is an encouraging result as gateway nodes may be costly to deploy in real networks.

The reason that the gateway nodes have such an impact is that they serve three purposes: (i) they help connect two nodes from different teams, (ii) they help improve the connectivity of the red (blue) networks individually, and perhaps less obviously (iii) they enable nodes from one team to connect nodes from the other team, e.g. two purple nodes with many blue nodes in between can form a bridge between two remote red nodes. This research has significant implications on secure communication in coalition networks.

## Key Management in MANETs

Often the first step to secure communication in MANETs is the key distribution phase that allows any (or most) pairs of nodes to establish a secure communication channel. Our work helped identify serious drawbacks in prior proposals and led to the development of new key pre-distribution solutions for coalition networks.

A major research focus due to its importance in the 2006–2010 time frame was the development of lightweight key pre-distribution schemes with applications to MANETs. Briefly, the goal of such schemes is to (a) allow any pair of nodes in the network to compute a common key, based only on their respective node identities, while (b) being resilient to the compromise of any t nodes in the system. This can be achieved by having each node store a key for every other node in the system; this results in per-node storage that is linear in the size of the network, which is unacceptable. It is known that if information-theoretic security is desired, then storage O(t) is optimal. In settings where t may be large, one can ask whether better efficiency can be obtained by relying on computational security instead. A series of prior papers [9,10,11] suggested using perturbation polynomials to construct key-pre-distribution schemes. Roughly speaking, these schemes all use *perturbation polynomials* to add noise to polynomial-based systems that offer information-theoretic security, in an attempt to increase the resilience threshold while maintaining efficiency. These schemes had so far been resistant to any known attacks and have been highly regarded in the community appearing in leading conferences. ITA researchers attempted to attack various pre-distribution protocols proposed in the literature.

This was the first attack on these protocols and the team demonstrated [12] that they can all be completely broken once one allows even a slight extension of the parameters beyond those achieved by the underlying information-theoretic schemes. They further showed that the general approach was not sound, and unlikely to ever yield secure schemes. This research was stimulated by the discussions of the state of the art in cryptography for MANETs and sensor networks, and has led to a fundamental result that contributes to the research community's greater understanding of what hard problems can and cannot be used as the basis for constructing secure identity-based encryption schemes.

Another important aspect of key management in MANETs was explored by a joint team of Royal Holloway, University of London (RHUL) and IBM US [13]. Key management is perhaps the most complex and most vulnerable part of any cryptographic implementation. While key generation and activation had been extensively studied in the context of mobile ad hoc and wireless sensor networks, there was a dearth of good research in designing techniques for key deactivation (revocation) and even more so for key reactivation. The researchers designed a new trust-based revocation scheme with the following three characteristics to support security services in a mobile ad hoc network: (i) *distributed* – the scheme did not require a permanently available central trusted authority, (ii) *active* – the scheme statistically guaranteed an incentive structure in which rational nodes are always encouraged to revoke malicious nodes while malicious nodes are discouraged from revoking other nodes in the network, and (iii) *robust* – the scheme was resilient against large numbers of colluding malicious nodes (e.g. it can tolerate collusions of size 30% of the network for an IDS error rate of 15%).

Several schemes in the literature have two of the above three characteristics, but none had all three. The prior best revocation scheme failed to support rational nodes (assumes good nodes are unselfish – in their scheme a rational node would defer revocation responsibilities to others, thereby, indefinitely postponing a revocation), and assumed an error free IDS system (which is infeasible in most practical settings). This was the first revocation scheme that satisfies all three properties making it well suited for environments, such as ad hoc networks, which are bandwidth constrained and dynamic and where nodes must operate under the threat of compromise.

### Hierarchical Identity Based Encryption

Network-centric coalition operations are faced with the challenges of resource constraints (e.g. due to the use of hand-held devices) and the necessity to interoperate across multiple administrative domains. Traditional cryptographic solutions (e.g. identity-based encryption, functional encryption) often assume a single trusted authority (e.g. during the key initialization and key set up phase) and thus needed to be revisited to accommodate multiple roots of trust.

Prior work on *identity-based encryption* (IBE) allows the use of a node's identity as its public key, without requiring a separate public key that is bound to the node's identity using certificates (issued by trusted authorities or by a chain of trusted authorities). In military networks, nodes are often associated with a rank and chain of command, thereby necessitating the need for supporting hierarchical relationships between nodes. The ITA research presented the first *hierarchical identity-based encryption* (HIBE) system that has full security for more than a constant number of levels.

A HIBE scheme is an extension of identity-based encryption with identities organized in a tree with a single root master authority. An identity in nonhierarchical IBE is a unique single identifier, e.g. "Bob Smith serial number 123456789". An identity in a HIBE scheme is an ordered collection of identifiers that can mirror the organizational structure of the Trusted Authority (TA), e.g. ("U.S. Army", "2-16 Infantry", "Bob Smith"). For concreteness, consider a two-level TA called "U.S. Army" with a hierarchical structure in which the first level of the hierarchy will identify a particular grouping of personnel, with the secret key for a grouping-e.g. ("U.S. Army", "2-16 Infantry")-being held by its commanding officer. In our example, the second level will identify a specific individual or role within the grouping; our examples use individuals' names for this level. In a hierarchical IBE, as in regular IBE, encryption is accomplished by using the identity of a recipient in place of using the recipient's certified public key. The hierarchical structure is meant to mirror the trust or authority structure of the organization. Any entity in a HIBE can use its own secret key to generate a secret key for any subordinate. If the head of the 2-16 Infantry Battalion has the secret key for identity ("U.S. Army", "2-16 Infantry") then he is able to generate secret keys of the form ("U.S. Army", "2-16 Infantry", "Name"). Any *ancestor* capable of generating a secret key for the addressee may then decrypt messages, but a subordinate (or *child*) is not capable of generating a secret key for or decrypting messages addressed to any ancestor. The root authority holds the master key for the HIBE scheme, making it capable of generating a secret key for any entity in the hierarchy and decrypting any message sent using the scheme.

An ITA research team from RHUL and CUNY explored the use of IBE to make it applicable to coalition networks. They proposed a lightweight, generic and broadly applicable framework enabling the refreshing of private keys in coalition networks [14]. This algorithm is an improvement upon the approach of simply distributing new private keys by encrypting them using the old public keys.

The approach used a refresh algorithm to generate a new private key for a given identity, sending it in a refresh message via a secure mechanism, which is coupled with a recovery algorithm for a node to recover its new private key. The scheme is secure and the framework is applicable to enable secure interoperation between entities with different trusted authorities.

Subsequent work [15] extended the notion of HIBE from the domain of a single TA to a setting with multiple, independent TAs each with their own HIBE. In this multi-TA HIBE environment, a group of TAs may form coalitions, enabling secure communication across domains. These coalitions may be temporary in the sense that a grouping may be formed and/or dissolved without compromising the (future) security of its member TAs. Similarly, the coalitions may be dynamic in the sense that individual coalition TA members may be added or removed at any time. Each individual TA produces its parameters and keys independent of any other TA. Coalition formation is achieved by way of two stages in which member TAs first exchange (public) messages and then broadcast (public) update messages to members of their respective hierarchies, allowing any system user to form a coalition key which allows the decryption of messages sent to members of the coalition. The work extended the functionality from coalition-based HIBEs to coalition-based WIBEs (here W indicating the capability of wildcarding in the hierarchical namespaces). This allows broadcast encryption from any user to large subsets of users within a coalition. Among the various uses of this approach, one can imagine a collection of national forces with independently generated WIBE schemes based on a common naming convention. This architecture would allow the forces to form ad hoc coalitions and provide any entity the ability to send messages to a large range of users within the coalition. The team developed a full syntax and security model for multiple TA WIBE schemes, and gave example

constructions with full security proofs against passive and active attackers.

Another effort by a separate group of ITA researchers presented the first HIBE system that has full security for more than a constant number of levels [16]. In all prior HIBE systems in the literature, the security reductions suffered from exponential degradation in the depth of the hierarchy, so these systems were only proven fully secure for identity hierarchies of constant depth. (For deep hierarchies, previous work could only prove the weaker notion of selective-ID security.) In contrast, the new system offered a tight proof of security, regardless of the number of levels; hence the proposed system is secure for polynomially many levels. This work represents an important theoretical advance in our understanding of HIBE and its security.

# Declarative Infrastructure for Security/Networking

Hybrid networks provide many technical benefits but are challenging to manage and secure, because the security and performance properties of hybrid technologies are harder to reason about. To understand the issues involved with hybrid network management, ITA researchers from Imperial College London and IBM developed a declarative infrastructure for agile and automated security/network management and control, constituting a unifying framework for the specification, analysis, implementation and evaluation of a wide range of network protocols and security constructs [17].

The declarative framework is depicted in Figure 19, and permits the specification, execution and analysis of network/security protocols. The first component of the framework is the *protocol model* that allows the specification of a given protocol, eased by the presence of a novel non-deterministic *choice* construct. The model's operational semantics interpret networks as collections of communicating state transition systems. For development and testing purposes, specified protocol models can be executed in a realistic setting. The model is general; it has been used to express path vector, link state, and MANET routing protocols, along security-relevant constructs. For example, it was used to implement a distributed network anomaly detection system, which can detect abnormal traffic without any prior knowledge of attack signatures.

The framework included an analysis query language with which an analyst can express security (and other) properties of interest. Queries were answered using an *Answer Set Programming* (ASP) solver, which performs surprisingly well on NP-hard problems. The query language could express properties such as loop
freedom and routing path convergence, which from a security point-of-view are important for avoiding denial of service.



Figure 19. Declarative networking framework

The team has applied the framework to several situations. As one compelling example, the team analysed a newly proposed MANET routing protocol [18] that aims to hide network topological information from routers, during the route discovery phase, due to security concerns. By expressing the correctness property (i.e. "paths computed by the protocol are disjoint") and the security requirements (i.e. "the protocol can always discover disjoint paths, if they exist") as queries, the analysis framework showed that the protocol, contrary to the claims of its authors, fails to meet its security requirements. It also generated counter-example traces to show the source of the flaw.

The team of researchers applied declarative networking principles to other aspects of secure network management. As one example, they considered the security of MPTCP, in particular considering the problem of multipath attack signature detection, which attempts to identify threats whose components have been split and sent across multiple paths, where traffic on each path seems innocuous, but in combination it is dangerous.

The team developed the first multipath signature-based intrusion detection algorithm that is able to detect signature-based attacks distributed across multiple paths via MPTCP protocols [19]. Inspired by the distributed and asynchronous features of the declarative networking framework, the researchers formalized a distributed signature-based intrusion detection problem as an asynchronous online exact string-matching problem and proposed a new algorithm that was first tested using their declarative networking simulation environment. Monitors located on different paths are modelled as automata, running for partially observed input strings. Asynchronous communication among the automata allows them to share a global state of the string matching. Different monitors may receive different sub-flows, with split signatures. Each of them scans each received packet locally and broadcasts its automaton state to all the other monitors. The team proved that if every packet in a multipath connection is captured by at least one monitor and if there exists a malicious pattern in the connection then the algorithm will detect it. Moreover, they proved that the time complexity of the algorithm is linear with respect to the size of the input, but detection time grows linearly with throughput. This result matches the typical single-path approach to detection. The researchers have reproduced this result experimentally (using an implementation written in C), under different network configurations (e.g. WiFi-WiFi, WiFi-LAN, and LAN-LAN).

During the security analysis of multipath TCP, ITA researchers found out a new vulnerability [20] in which an adversary can infer properties of a flow split across its own and a hidden network: Adjusting the properties of the flow (e.g. its bandwidth), and observing performance of the sub-flow on its own network, the adversary can infer properties of the network transmitting the other sub-flow.

### Knowledge-Based Security Policies

One of the key challenges in coalition operations is to support flexible collaboration amongst coalition members while protecting sensitive data. For example, a coalition member may be willing to reveal that it can provide sufficient resources to staff a mission, but is not willing to reveal all of the resources that it owns. Or, a coalition member may reveal that it has identified enemy activity in a general area of operations but does not want to reveal the exact location, for fear of giving away locations of its sensors. In short, coalition members should be able to collaboratively filter, fuse, and query information from multiple sources to meet mission objectives while, nevertheless, retaining some level of secrecy.

Abstractly, the technical problem can be framed as follows: given a piece of sensitive information X and a function (or *query*) F, how much can be learned about X based on the computed result F(X)? What additional information might

another function G leak about X via its result G(X)? How should a data owner specify a limit on the amount of knowledge a querier might be able to obtain?

ITA researchers at University of Maryland and IBM developed a technique called *knowledge-based security policies* (KBSP) to address these questions [21,22]. The technique works by estimating the recipient's belief about the likelihood of possible values of X as a probability distribution D, and then using a novel technique called *probabilistic abstract interpretation* to compute revised belief D' = D | F(X). This technique is noteworthy in that it works for any function F implemented in a fairly general programming language. If D' would assign too much probability mass to a possible secret X', then function F is too revealing, and should be refused. Otherwise, F can be answered and D' is retained as the updated belief used to assess future queries.

In related work [23], they considered the application of queries over a location trace with anonymized identities, but showed that the trace can be easily deanonymized when social network information about participants is available. They also generalized the work from quantifying flows from static secrets, which do not change over time, to those from dynamic secrets, which do change (with location being a prime example of a dynamic secret) [24]. They developed a novel formal model, and metrics, for quantifying information flows from dynamic secrets. This model allowed for adaptive adversaries, whose interactions with the system can influence it, and be influenced by it. Using this model experimentally they found that, counter intuitively, frequent change of a secret can actually increase leakage; this occurs when the pattern of change is readily inferred. They proved that waitadaptive and input-adaptive adversaries, who can choose particular inputs to, and when to exploit, a target system, can derive significantly more information, and in a monotonic fashion, than adversaries that cannot make such choices. Thus, ignoring the adversary's adaptivity might lead one to conclude secrets are safe when they really are not.

The team also generalized their KBSP work to secure multi-party computation (SMC), which is a cryptographic technique that computes a function  $F(X_p,...,X_N)$  among N parties, each with a secret  $X_p$  where each party receives the output of the function but nothing else (i.e. as if F were computed by a separate, trusted party). Their generalized technique [25] can determine what each party will learn about the others' inputs based on the output. KBSPs for SMCs can be used for mission planning: N partners attempt to compute a mission plan that will only succeed if the resulting plan does not reveal too much about each partner's (sensitive) resources, according to their own policies. They also developed a

means to deconstruct an SMC function F into several smaller SMC functions  $F_{I}$ ,  $F_{2}$ , etc. which can be executed sequentially, at much lower cost than F, but which (according to knowledge-based reasoning) reveal no more information about the input secrets, in aggregate, than does F. For some functions, like joint median, the benefit can be orders of magnitude performance improvements.

SMCs are useful, but hard to use because only low-level mechanisms (e.g. circuit descriptors) have been available to specify them. Therefore, the team developed a new, high-level programming language for SMCs called *Wysteria* [26]. Wysteria supports generic, N-party, mixed-mode programs, which involve any number of participants (determinable dynamically) and combine local, private computations with synchronous SMCs.

They also developed SCVM as complementary work to Wysteria, which employs *RAM-model secure computation*. This is an SMC approach that addresses the inherent limitations of circuit-model secure computation considered in almost all previous work, including Wysteria, by using *Oblivious RAM* (ORAM) to avoid side channels due to programs with random accesses [27]. SCVM constituted the first automated approach for RAM-model secure computation. Leveraging compile-time optimizations, this approach speeded up computation by about two orders of magnitude compared to both circuit-model secure computation and the state-of-art (but by-hand) RAM model secure computation.

SCVM leveraged a concept called *memory-trace oblivious program execution* (MTO), which is a security property for secure computations executing on untrusted platforms, in which memory accesses (e.g. to fetch program instructions or encrypted memory) can be observed. They showed how a relatively standard information flow type system could be combined with ORAM to ensure MTO (which is a stronger property than that provided by ORAM on its own) but at relatively low cost (orders of magnitude faster than using only ORAM)[28].

### Verified Outsourced Computing

In coalition operations, it is quite common that adequate processor capacity for performing a computation task is only available from a coalition partner that cannot be completely trusted. It would be desirable to outsource the computation to the partner's systems, but verify that information has not leaked and that the computation has been performed properly. ITA researchers have explored the use cryptographic security technologies to allow data from one security domain to be processed in a different domain (or in many domains at once) in a way that respects the security requirements of the data's original security policy [28]. One of the benefits from this exploration has been an improvement in the practicality of *fully homomorphic encryption* (FHE), which supports arbitrary computations on encrypted data without revealing the data. Research in this area was conducted at CUNY, University of Maryland and IBM.

Among several advances along this frontier, ITA researchers developed a *verifiable database* that allows a weak client to outsource a large database table to an untrusted server and makes retrieval and (unbounded) update queries [29]. This approach was a first-of-a-kind that simultaneously satisfied compactness (and thus incurred little verification overhead on the client) and yet offered full security. Follow-on work [30] supports public verification, i.e. the result of the computation is verifiable by any third party, and requires no secret key, including optimizing for simpler classes of computations.

To prove the practicality of their research, ITA researchers developed and released a C++ library of the verifiable database scheme, and later created a version that can be efficiently implemented on devices such as *field-programmable gate arrays* (FPGAs). Rather than program the FPGAs directly, they developed an implementation in the Lime high-level language [31], thereby allowing for cross compilation onto an FPGA to obtain the necessary performance improvements. To illustrate how FHE can be used in practice, they developed a demonstration of secure two party oblivious transfers. The approach has many applications in coalition environments [32].

This work led to development of a practical protocol for verifiable keyword search, which is the problem of a client storing a document D with an untrusted server and later wanting to verify whether a given keyword w appears in D. Using our verifiable database scheme would create a significant workload at the server—to respond to a query or update, the server must read or reread the entire file, performing exponentiations along the way. In response, the researchers developed a basic indexing technique that dramatically improves the server's efficiency without compromising the client's efficiency. Keyword searches in this modified protocol ran in constant time without substantially changing storage requirements or the cost of the initial upload phase. Moreover the same indexing technique yielded an efficient way to handle updates to the file D. Implementation experiments with a mobile client app and a server operating on large files showed that the approach is practically feasible.

The above two problems are specific instances of the general problem of building

an authenticated data structure (ADS) [33], which is a data structure whose operations can be carried out by an untrusted server (or prover) which provides a proof to the client (the verifier) that the operation was done correctly (was authentic). Past work on ADSs had focused on particular data structures (or limited classes of data structures), one at a time, often with support only for specific operations (such as reads, not updates). To ease the construction of ADSs, the researchers developed a simple extension to a functional programming language, in which a programmer simply designs the data structure and its operations as usual, and then adds a few annotations and coercions to identify the key places to ensure authenticity. Pleasantly, the programmer gets security (i.e. verifiability) by construction-all data structures and operations that type check are sure to be authentic. The annotations only affect performance. The approach has been implemented as an extension to the OCaml compiler, and used to produce authenticated versions of many interesting data structures including binary search trees, red-black trees, skip lists, and more. Performance experiments show that this approach is efficient, giving up little compared to the hand-optimized data structures developed previously.

ITA researchers have also considered variations and generalizations of the secure outsourcing problem. In particular, they have explored schemes that ensure *privacy*, so that the outsourced function and/or data are kept hidden from the computation agent. In a specific instance, they formalized and implemented a scheme for *Verifiable Oblivious Storage* (VOS), in which a client outsources the storage of data to a server while ensuring privacy of the data and verifiability and obliviousness of access to that data [34]. VOS generalized the notion of ORAM in that it allowed the server to perform computation, and also explicitly considered the issue of data integrity and freshness. They showed that allowing server-side computation makes it possible to circumvent the known lower bound on the bandwidth required for ORAM constructions. They also proved theoretical results about the limits of *adaptively secure* FHE schemes, suitable for use in interactive protocols.

In additional research [35], ITA researchers have considered a generalization of a *single-client* solution to the secure outsourcing problem to a *multi-user* setting (that is common place in coalition environments). In particular, suppose there are N clients who wish to outsource the computation of a function F over their joint inputs  $x_p$ , ...,  $x_N$  to an untrusted server. They proposed a protocol for noninteractive computation verifiable by all users. The solution has many benefits: each client's input is kept private from all the others (this holds assuming all parties are honest-but-curious), and each client's work scales linearly in the length of its own input.

Finally, ITA researchers addressed the problem of broadcasting data over an insecure channel to a dynamically changing population of receivers. Their specific aim was to develop cryptographic techniques that ensured not only the secrecy of the data being transmitted, but also the anonymity of the authorized receivers. They proposed the first broadcast encryption scheme with sublinear ciphertexts to achieve meaningful guarantees of receiver anonymity [36]. In particular, they formalized the notion of outsider-anonymous broadcast encryption (oABE). The degree of anonymity captured in their security model was between the complete lack of protection that characterizes traditional broadcast encryption schemes on one end, and full anonymity on the other end. More specifically, in their setting, recipient identities are hidden from users not authorized to receive the message, but individual recipients might learn who else is getting the same message. They describe a generic oABE construction based on any anonymous identity-based encryption scheme (AIBE). This construction has the advantage of being stateless, and with constant public key size. Additionally, they also obtained an efficient construction with enhanced decryption, where for a given oABE ciphertext, the decryption algorithm executes a single AIBE decryption operation. By relaxing the anonymity guarantees, their constructions could achieve sublinear ciphertexts size and constant public key size.

#### References

- [1] S. Calo, C. Karat, J. Karat, J. Lobo, R. Craven, E. Lupu, K. Ma, A. Russo, M. Sloman and A. Bandara, "Policy Technologies for Security Management in Coalition Networks", in *Network Science for Military Coalition Operations: Information Exchange and Interaction*, Ed., D. Verma, IGI-Global, pp. 146-173, 2010.
- [2] J. Lobo, J. Ma, A. Russo, E. Lupu, S. Calo and M. Sloman, "Refinement of history-based policies", in *Logic programming, knowledge representation, and nonmonotonic reasoning*, Springer Berlin Heidelberg, 2011.
- [3] R. Craven, J. Lobo, E. Lupu, A. Russo and M. Sloman, "Security policy refinement using data integration: a position paper", in Proc. 2nd ACM workshop on *Assurable and usable security configuration* (SafeConfig), 2009.
- [4] A. Bandara, S. Calo, R. Craven, J. Lobo, E. Lupu, J. Ma, A. Russo

and M. Sloman, "Expressive Policy Analysis with Enhanced System Dynamicity", in Proc. *ASIACCS*, 2009.

- [5] A. Bandara, A. Kakas, E. Lupu, and A. Russo, "Using Argumentation Logic for Firewall Configuration Management", in Proc. *IFIP/IEEE International Symposium on Integrated Network Management*, 2009.
- [6] M. Beigi, J. Lobo, K. Grueneberg, J. Karat and S. Calo, "A negotiation framework for negotiation of coalition policies", in *IEEE Policies for Distributed Systems and Networks*, IEEE Policy 2010.
- [7] C. Capar, D. Goeckel, C. Leow, K. Leung and D. Towsley, "Network Coding for Facilitating Secrecy in Large Wireless Networks", in Proc. *IEEE Conference on Information Sciences and Systems* (CISS), 2012.
- [8] B. Bash, D. Goeckel and D. Towsley, "Square root law for communication with low probability of detection on AWGN channels", in Proc. *IEEE International Symposium on Information Theory* (ISIT), 2012.
- [9] W. Zhang, M. Tran, S. Zhu and G. Cao, "A random perturbation-based scheme for pairwise key establishment in sensor networks", in Proc. 8th ACM MobiHoc, September 2007.
- [10] N. Subramanian, C. Yang, and W. Zhang, "Securing distributed data storage and retrieval in sensor networks", in Proc. *IEEE Pervasive and Mobile Computing*, 2007.
- [11] W. Zhang, N. Subramanian and G. Wang, "Lightweight and compromiseresilient message authentication in sensor networks," in Proceedings of IEEE INFOCOM 2008.
- [12] M. Albrecht, C. Gentry, S. Halevi and J. Katz, "Attacking cryptographic schemes based on perturbation polynomials", in Proc. 16th ACM Conference on Computer and Communications Security, 2009.
- [13] S. Reidt, M. Srivatsa and S. Balfe, "The Fable of the Bees: Incentivizing Robust Revocation Decision Making in Ad Hoc Networks", in Proc. ACM Conference on Computer and Communications Security, 2009.
- [14] S. Balfe, K. D. Boklan, Z. Klagsbrun, and K. G. Paterson, "Key refreshing in identity-based cryptography and its applications in MANETs", in Proc. *IEEE MILCOM*, 2007.
- [15] K. Boklan, A. Dent and C. Seaman, "Broadcast encryption in multiple trust domains with dynamic coalitions", in Proc. LATINCRYPT, 2010.

- [16] C. Gentry and S. Halevi, "Hierarchical Identity Based Encryption with Polynomially Many Levels", in Proc. *Theory of Cryptography Conference*, Springer LNCS 5444, pp. 437-456, 2009.
- [17] J. Lobo, J. Ma, A. Russo, Frank Le, "Declarative Distributed Computing", in *Lecture Notes in Computer Science*, Volume 7265, Correct Reasoning, pp 454-470, 2012.
- [18] G. Zhang, Q. Hu Wang, J. Tian, Z. Li, "Design and performance study of a topology-hiding multipath routing protocol for mobile ad hoc networks", in Proc. *IEEE INFOCOM*, 2012.
- [19] J. Ma, F. Le, A. Russo and J. Lobo, "Detecting Distributed Signaturebased Intrusion: The Case of Multi-Path Routing Attacks", in Proc. IEEE INFOCOM, 2015.
- [20] Z. Shafiq, F. Le and M. Srivatsa, "Cross-Path Inference Attacks on Multipath TCP", in Proc HotNets, 2013.
- [21] P. Mardziel, S. Magill, M. Hicks and M. Srivatsa, "Dynamic Enforcement of Knowledge-based Security Policies", in Proc. *Computer Security Foundations* (CSF), June 2011.
- [22] P. Mardziel, S. Magill, M. Hicks and M. Srivatsa, "Dynamic Enforcement of Knowledge-based Security Policies using probabilistic abstract interpretation", in *Journal of Computer Security*, vol. 21, no. 4, 2013.
- [23] M. Srivatsa and M. Hicks, "Deanonymizing Mobility Traces: Using Social Networks as a Side-Channel", in ACM Conference on Computer and Communication Security, 2012.
- [24] P. Mardziel, S. Magill, M. Hicks and M. Srivatsa, "Dynamic Enforcement of Knowledge-based Security Policies", in Proc. Computer Security Foundations (CSF), June 2011.
- [25] P. Mardziel, M. Hicks, J. Katz and M. Srivatsa, "Knowledge-Oriented Secure Multiparty Computation", in Proc. *Programming Languages Analysis and Security* (PLAS), June 2012.
- [26] A. Rastogi, M. Hammer, and M. Hicks, "Wysteria: A Programming Language for Generic, Mixed-Mode Multiparty Computations", in Proc. *IEEE Symposium on Security and Privacy* (Oakland), May 2014.
- [27] C. Liu, Y. Huang, E. Shi, J. Katz, and M. Hicks, "Automating Efficient RAM-Model Secure Computation", in Proc. IEEE Symposium on Security

and Privacy (Oakland), May 2014.

- [28] R. Gennaro, C. Gentry, and B. Parno, "Non-interactive verifiable computing: Outsourcing computation to untrusted workers", in Proc. *CRYPTO*, 2010.
- [29] S. Benabbas, R. Gennaro and Y. Vahlis, "Verifiable Delegation of Computation over Large Datasets", in Proc. CRYPTO, 2011.
- [30] D. Fiore and R. Gennaro, "Publicly verifiable delegation of large polynomials and matrix computations, with applications", in Proc. ACM Conference on Computer and Communications Security, 2012.
- [31] J. Auerbach, D. Bacon, P. Cheng, and R. Rabbah, "Lime: a Javacompatible and synthesizable language for heterogeneous architectures", in ACM Sigplan Notices, vol. 45, no. 10, pp. 89-108, 2010.
- [32] G. Bent, F. Bergamaschi, and H. Hunt, "Computing on encrypted data and its applicability to a coalition operations environment", in Proc. SPIE Defense+ Security, 2015.
- [33] A. Miller, M. Hicks, J. Katz, and E. Shi, "Authenticated Data Structures, Generically", in Proc. ACM POPL, July 2013.
- [34] D. Apon, J. Katz, E. Shi and A. Thiruvengadam, "Verifiable Oblivious Storage," in Proc. 17th IACR International Conference on Practice and Theory of Public Key Cryptography, PKC 2014.
- [35] S. G. Choi, J. Katz, R. Kumaresan, and C. Cid, "Multi-user Noninteractive verifiable computation", in Proc. *Annual Conference of the ITA*, 2012.
- [36] N. Fazio and I. M. Perera, "Outsider-Anonymous Broadcast Encryption with Sublinear Ciphertexts", in Proc. 15th IACR International Conference on Practice and Theory of Public-Key Cryptography (PKC'12), Lecture Notes in Computer Science, Springer 2012.

# 4 Distributed Information Processing in Coalition Networks

In this chapter we describe the work done in the ITA research programme that addresses different aspects of distributed information processing in coalition networks. Information processing in a coalition network needs to be performed on a variety of data sources including, but not limited, to sensor information, structured information from various databases accessible to warfighters, human intelligence as well as information available from sources such as the Internet. A key challenge is how to perform information processing within networks that have unreliable communication, intermittent connectivity and comprise a wide range of information. In addition to the challenges imposed by the physical network there are also the challenges of operating in national coalitions with different types of assets, different policies on how assets can be accessed and used, and different models for information representation and information processing.

The purpose of performing information processing is to provide timely situation understanding by combining data from the appropriate sources in response to mission goals. This not only requires a capability to access and process the data but in many cases to manage the assets that are capable of providing this information. When the ITA research programme began in 2006 much of this analysis was, and to a large extent still is, pre-planned with all the data required to perform the analyses being collected from a variety of sensors and sources, and processed in predefined stovepiped systems with the results being disseminated to the users. The wide range of users, assets and networking technologies found at the edge of the network pose a unique set of challenges in the areas of platform discovery, integration, access and control. Current edge of network systems, including ISR applications, often deliver stovepipe solutions that provide a single point of access to information. This increases cost and operational complexity, limits scalability, prohibits sharing, and reduces the overall value of the assets. As the battlefield becomes more dynamic these types of inflexible systems are not able to keep pace with the changing needs of the commander, and in coalition operations often prevent the sharing of essential information.

It was realized that a radically new approach to information processing was needed, where information could be accessed and processed from anywhere in the network, without the user having to know explicitly where the required data or information processing resources resided in the network. This fully decentralized vision has to take into account how information should be represented, mechanisms for querying for information (pull model) and for exchanging information (push model) and the balance between the two approaches, including hybrid models of push and pull. The information processing has to be performed in an assured way, taking into account different security and policy constraints of coalition operations, and has to be self-adapting to changing mission goals and the added complexities of coalition operations.

In order to achieve this vision it was necessary to invent many different techniques, ranging from approaches that can migrate processing in the network whenever it was needed, understanding how to best access data from many different locations, understand the quality of information that was available, and determine how to find the asset best suited for a specific analysis need, among others. Some of those inventions are enumerated below.

#### **Distributed Dynamic Processing**

In a hybrid coalition network, communication links are unreliable and frequently limited in their capacity. As mentioned above, the general practice in the military community to process information in such networks is to bring all the data back to a single location and to process them there. ITA researchers explored and put forward an alternative view—instead of bringing the data to the processing element, perhaps the processing capability ought to move to the data. In an environment where processing code is smaller than the volume of data that needs to be processed, such a design point could yield much better overall performance.

The main challenge in attaining this vision is the uncertainty where the processing ought to be pushed within a network (modelled as a complex graph) where each of the nodes can be a potential location for information processing. The right node would need to pull the data from the relevant sensors, and obtain the right balance between this push and pull. Further complexities are introduced because the processing function itself is usually a complex graph of many different subfunctions. Further, military networks are highly dynamic, and coalition partners have various restrictions on what can be done where.

An ITA research team from Imperial College London, IBM, ARL and UMass developed new algorithms and architectures that could enable dynamic movement of processing in the network. They defined the information processing application as a logical network of functional nodes, and thus mapped the dynamic distributed problem to that of mapping a logical network on to a physical network [1]. This mapping is applicable in many scenarios, for example, the assignment of network services onto physical devices in wireless or wire-line environments, mapping sensor information fusion elements to nodes in a wireless sensor network, or mapping different components of a cloud based application within an application distributed network. The goal in the mapping is to assign the logical nodes (in the logical network) on to the physical nodes and obtain a physical resource allocation that meets the logical network demands, subject to physical network constraints.

To solve the problem, which is computationally hard (NP-Complete), the researchers initially proposed a two-step approach, the first step providing a set of novel feasibility checks for node assignment (that they proved to be both necessary and sufficient), and the second step a simple and fast algorithm that achieves a feasible logical to physical mapping with high probability of success [1]. Simulation studies demonstrated the efficiency of the proposed approach. In addition to solving the problem for distributed dynamic processing for sensor information fusion, this work also exemplified the case of interoperation between different network layers, which is one of the core aspects of network science.

The migration of processing into the network has one drawback, namely that each processing element can process data sources only from a small region of the network, while a central processing element that processes data from all sources has global visibility. In a tactical environment, where users are mobile, this can cause users to lose connectivity with the processing elements they are associated with. Therefore, the mapping of logical nodes to physical nodes cannot be static, and processing services need to be migrated as the configuration of users changes. The ITA research team explored this challenge in [2], where they modelled the general problem as a *Markov decision process* (MDP). They proved that, in the special case where the mobile user follows a one-dimensional asymmetric random walk mobility model, the optimal policy for service migration is a threshold policy. They obtained an analytical solution for the cost resulting from arbitrary thresholds, and then proposed an algorithm for finding the optimal thresholds. The proposed algorithm was shown to be more efficient than standard approaches for solving MDPs.

The team also addressed other variations of the problems, including approaches that could anticipate the mobility patterns within the network and make predictive assignments of logical nodes to physical nodes [3], as well as addressing services that were composed from other underlying services [4]. Coalition networks have to address different security policies, which affect the right solution for mapping a processing element to the underlying networks, and the researchers proposed an approach to address the constraints imposed by security policies [5].

In further analysis of this problem in the context of hybrid cellular/MANET networks [6], they also demonstrated that under the assumption of stationary and independent requests, it is optimal to adopt static caching (i.e. to keep a cache's content fixed over time) based on content popularity. They also showed that it is optimal to route to in-MANET caches for content cached there, but to route requests for remaining content via the cellular infrastructure for the congestion-insensitive case and to split traffic between the in-MANET caches and cellular infrastructure for the congestion-sensitive case.

# Quality of Information

The concept of QoI, albeit under consideration in enterprise/database systems, was not part of the lexicon in the literature on wireless sensor networks at the start of the ITA programme. Anticipating the central role that data-centric services and open sensor deployment models will play in the evolution of wireless sensor networks (where sensor networks and applications will be deployed independently of each other and bind based on application needs), ITA researchers introduced novel ideas and concepts around QoI which serve as key facilitators for these services. The team working on QoI comprised of scientists from Imperial College London, Pennsylvania State University (Penn State), University of California at Los Angeles (UCLA), IBM and ARL.

Starting from the seminal paper that introduced QoI for sensor networks [7], the ITA research team developed the concepts in further detail to identify principles underlying QoI [8], a metadata model for QoI [9], and provided algorithms and approaches to determine QoI in many sensor information processing contexts. These included incorporating QoI concepts in sensor-based event detection systems [10], sensor network deployment and planning [11], detection of transient signals where only incomplete observations can be made [12], advanced signal

processing techniques [13], target tracking [14], security metadata [15], and privacy of information [16]. The concept of *value of information*, which is related but distinct from the quality of information, was also introduced [17].

Quality of Information had a profound impact on the broader research community. Of specific note is the impact the concept had on the research directions of the Collaborative Technology Alliance in Network Science, where it became a major research theme [18].

#### Mission Optimized Sensor Networks

In a coalition network, dynamic communities of interests are formed to conduct an ad hoc set of operations. Supporting such dynamic communities in an agile manner means the underlying sensing, computing, and communication infrastructure should be optimized to meet the needs of the mission of the community. Since coalition missions tend to be diverse, it is critical to have algorithms and architectures that can quickly adapt the infrastructure to deliver the requirements of the mission in the best possible way.

A team of ITA researchers spanning Penn State, Imperial College London, CUNY, IBM and ARL have explored different facets related to the optimization of sensor networks to best meet the goals of a mission. An understanding into the algorithms that can optimize the infrastructure would allow the right data to get to the right person at the right time.

The underlying framework for optimizing the infrastructure is to assign a utility to a piece of sensor information being received by a person at a given time. The utility would be defined on the specifics of the mission at hand. A utility maximization framework would then enable the optimization of the infrastructure to best serve the needs of the mission. Approaches like distributed *network utility maximization* (NUM) can then be used to determine the appropriate configuration of the network [19].

The ITA research team extended NUM to take into account characteristics such as dynamicity in coalition wireless sensor networks [20]. They extended the framework to cover mission priorities [21], which could be different for different coalition partners. They considered dynamic placement of functions such as data compression and fusion within the network to save power consumption among sensor nodes as they satisfy the mission requirements [22]. They also applied control-theoretic approaches to get insights into the principles that can make a sensor network self-optimizing

#### [23,24].

Another important aspect of optimizing the infrastructure is determining the right placement of sensors to best meet the requirements of a coalition mission. The researchers examined the benefits of dense sensor deployments [25], including naturally providing k-coverage and overcoming inexact sensor placement, the limits of placement variance to obtain target coverage [26], and accommodating variable sensor coverage due to terrain when placing sensors [27].

### Sensor Mission Assignment

A key problem in managing ISR operations in a coalition context is assigning available sensing assets—of which there are increasingly many—to mission tasks. High demands for information and relative scarcity of available assets implies that assignments must be made taking into account all possible ways of achieving an ISR task by different kinds of sensing. The dynamic nature of most ISR situations means that asset assignment must be done in a highly agile manner. The problem is further exacerbated in a coalition context because users do not have an overview of all suitable assets across multiple coalition partners.

A cross-organizational team of researchers from University of Aberdeen, Cardiff University, Penn State and IBM have addressed this challenge by using a knowledge driven approach that used ontologies, allocation algorithms, and a service-oriented architecture. Mission scripts and sensor descriptions were represented as a collection of formal ontologies. The ontologies were based on the *Military Missions and Means Framework* (MMF). The researchers envision these ontologies being fed into a semantic reasoner [28] which would use real-time algorithms, realized as distributed protocols, to select specific sensor instances for the mission tasks.

The research team determined a bound on the utility and devised distributed and centralized algorithms and compared their achieved utility with this bound. The base algorithms considered both static and dynamic settings [29–33]. They also extended the work to incorporate operational costs and energy consumption. The output of the algorithms resulted in sensor bundles that would identify a collection of sensors that maximized the utility of the missions.

The team has implemented a system incorporating these principles [34] that has been used in a number of successful transitions.

### **Dynamic Distributed Federated Databases**

One of the main challenges for supporting networked information processing is that of providing efficient access to distributed sources of data, where the applications requiring the data do not necessarily have knowledge of the location of the data. There is a need for a distributed service registry, where each service holds its own metadata that can be accessed by any other node in the network.

The concept of distributed federated databases began as a new and different way of thinking about the way information is processed in sensor networks. ITA researchers wanted to create an architecture for sensor networks where sensors did not continually push information to an analyst, but instead could be queried on demand. In a dynamically changing bandwidth constrained network environment where coalition operations restricted access to some sensors sometimes, the new model holds many advantages. This new approach required a mechanism to propagate a federated query to each sensor and then to efficiently combine the resulting data. Using a modified *relational database management system* (RDBMS) engine it was possible to take queries in the *Structured Query Language* (SQL) and map them into an optimized set of targeted native queries against a set of heterogeneous data sources. So the initial motivation was to treat sensors as data sources that stored data locally in a file or a small database, and then propagate queries to them from a central location.

It was clear that they needed not just one centralized RDBMS database but some form of self-organizing network of small RDBMS databases, each capable of federating local data source(s). In this case the same SQL query could be propagated to all nodes and locally mapped into targeted native queries of the local data source(s). The researchers brought together concepts from enterprise database federation, self-organization and emergent phenomena, graph algorithms and semantic technologies to create a *dynamic distributed federated database* (DDFD).

The DDFD concept is illustrated in Figure 20 and comprises a set of interconnected vertices (N1, N2,..) each of which is a federated RDBMS engine. Each database engine is able to access internal and external sources as if it was one logical database. External sources may include other RDBMS, including commercial databases, or any other data source such as flat files of data records.

Data can be stored at any vertex in the network of database vertices, with the table in which it is stored being available to other vertices through a logical table mapping. SQL queries can be performed at any vertex requesting information from any other vertex without knowing where the data is located or the route to the data. Since the data location is unknown the query propagates through the network as a controlled flood, such that it propagates to every node exactly once and result sets are returned to the querying vertex. The vertices manage the forward routing of queries so as to minimize the number of copies of the same query. Thus result sets will be received over the shortest paths from vertices that can satisfy the query, with null results (i.e. table found with no data) or zero result sets (no table found along the route) being returned.



Figure 20. The DDFD concept

The fundamental insight that enabled the successful attainment of the DDFD vision was the realization that each piece of sensor information has only one writer. Thus a piece of data will never have to guard against concurrent update by multiple writers, and just a read-only distributed query mechanism is required, which can be built by creating a self-organizing topology. Nevertheless, in order to attain the vision of DDFD, ITA researchers had to resolve several scientific questions first.

One of the scientific questions was the choice of the right approach to selforganization. The team evaluated the properties of the self-organizing graph, such as graph diameter, which led to algorithms that minimized the graph diameter and maximized the resiliency when two separate graphs (e.g. belonging to different coalition partners) merge [35]. Another scientific question dealt with the right measures to handle the complexity of self-organization. The researchers used information entropy as a metric for the effectiveness of self-organization schemes [36,37]. Another issue related to the modelling and evaluation of the performance of such an architecture. They investigated the performance of their self-organization algorithms on a number of topologies, finally focussing on certain topological relationships—including the hypercube topology—to obtain performance bounds on their federation approach [38].

All of these scientific insights resulted in the creation of a prototype DDFD system. That database proved very useful in many different transition contexts. The reader can find more details of the implementation of the DFDD system in "The Gaian Database" on page 103.

### Fault Isolation in Information Networks

Any information source has the potential to be faulty, either due to a malfunction or sometimes due to deliberate tampering. When processing information, faults in the information source or in the underlying network can lead to erroneous effort. Even ignoring malicious subversion, sensor data quality may be compromised by non-malicious causes such as noise, drifts, calibration, and faults. On-line detection and isolation of such misbehaviours is crucial for efficient operation and management by avoiding wasted energy and bandwidth in carrying poor quality data and enabling timely repair of sensors. As a result, principles and approaches that can identify faults within the information network need to be developed so that robust solutions to handle faulty information can be developed.

ITA researchers from UCLA, CUNY, and Imperial College London collaborated to create a two-tiered system [39] for on-line detection of sensor faults. A local tier running at resource-constrained nodes used an embedded model of the physical world, together with a hypothesis-testing detector, to identify potential faults in sensor measurements and notify a global tier. In turn, the global tier used those notifications during fusion for more robust estimation of physical world events of interest to the user, as well as for consistency checking among notifications from various sensors, and generating feedback to update the embedded physical world model at the local nodes. This approach eliminated the undesirable attributes of purely centralized and purely distributed approaches that respectively suffer from high resource consumption from sending all data to a sink, and high false alarms due to lack of global knowledge. A subset of these researchers also addressed the problem of detecting and repairing erroneous (i.e. dirty) data [40] caused by inevitable system problems involving various hardware and software components in a sensor network. As information about a single event of interest in a sensor network is usually reflected in multiple measurement points, the inconsistency among multiple sensor measurements serves as an indicator for a data quality problem. One needs methods that can effectively detect and identify erroneous data among inconsistent observations, based on the inherent structure of various sensor measurement series at a global view.

The researchers developed three models to characterize the inherent data structures among sensor measurement traces, and then applied these models individually to guide the error detection of a sensor network. The first is a Multivariate Gaussian model that explores correlated data changes across sensors in a group. The second is a *principal component analysis* (PCA) model that captures the sparse geometric relationship among sensors in the network. The PCA model is motivated by the fact that not all sensor networks have clustered sensor deployment and clear data correlation structure. Further, if the sensor data shows non-linear characteristics, a traditional PCA model cannot capture the data attributes properly. Therefore, they proposed a third model which utilizes k functions to map the original data into a high dimension feature space and then apply the PCA model to the mapped linearized data. All three models serve the purpose of understanding the underlying phenomena of a sensor network from a global perspective, and then guide the error detection so as to discover any anomalous observations.

Another collaborative effort between Penn State, ARL and IBM [41] examined techniques to address faults that arise within the communication network infrastructure of MANETs. The team characterized the ability to identify faults when large scale clustered failures occur in a network. They developed an algorithm that uses both negative (outage) and positive information (connectivity) symptoms during outages to create a hypothesis list of possible causes. Then, using probabilistic models based on varying amounts of knowledge about the network topology and topography, they rank-ordered the elements on the list. Their results showed that the inclusion of positive information greatly improved the accuracy of the hypothesis list.

# Self-Organization of Services

The natural world is enormous, dynamic, incredibly diverse, and highly complex. Despite the inherent challenges of surviving in such a world, biological organisms evolve, self-organize, self-repair, navigate, and flourish. Generally, they do so using only local knowledge and without any centralized control. Our computer networks are increasingly facing similar challenges as they grow larger in size, but are not yet able to achieve the same level of robustness and adaptability. Many research efforts have recognized these parallels, and wondered if there are some lessons to be learned from biological systems. Within the ITA programme, an effort was undertaken to explore if biologically inspired techniques can be used for self-organization of wireless and sensor networks.

In the collaborative work between University of Cambridge, UCLA, ARL and IBM, the researchers began by exploring why biology and computer network research are such a natural match, and presented a broad overview of biologically inspired research. They concluded that research efforts were most successful when they separate biological design from biological implementation—i.e., when they extract the pertinent principles from the former without imposing the limitations of the latter. This survey work served as a starting point for ITA's work on self-organization, given that it provided insights and lessons learned by other researchers in applying the biological metaphors in computer systems.

Delivering a robust, high-performance application service over static networks, such as the Internet, requires careful configuration and maintenance. Different functional components of the service—such as caching, data processing, or data storage components—must be deployed at appropriate locations based on the resource constraints of the network and the receiving hosts. Although this manual configuration approach works over the relatively static topology of the Internet, the highly dynamic nature of mobile, ad hoc networks would quickly render an initially well-configured service unusable. In a collaborative work, UCLA and IBM researchers confronted this problem by exploring service emergence, a new design paradigm for self-organizing services, where individual nodes automatically adjust the roles they play in the system without any centralized control. From simple, localized interactions between neighbouring nodes, a functional system emerges that meets the desired service goals and requirements.

This work presented a roadmap for designing emergent services. It first described an example surveillance service, which provided both motivation and grounding for service emergence. Based on the study of biological systems, which exhibit intrinsic emergent behaviours, they extracted a few guidelines for use in our engineered network designs. They also discussed one preliminary design attempt, along with evaluation criteria for emergent service designs. In another initiative, ITA researchers designed a self-organizing MANET based on the concept of morphogenesis. Morphogenesis is the process that gives shapes to organisms from an embryonic stage using a process of cell division. Starting from a simple embryonic cell, the controlled division and transformation of the cells into different types leads to the creation of a complex organism. The growth of complex organisms is completely autonomic, and is one of the best examples of self-organizing systems that can be found in nature. In comparison, computer networks of today are much more static and require significant human intervention in order to take on the shape and size that is desired. However, emerging technologies such as virtualization and network computing enables an architecture where computer networks can also develop using morphogenesis. The researchers presented the benefits of morphogenesis in a computer network, and described the architecture for a computer network, which self-organized using these principles.

Another aspect of self-organization of services can be illuminated by exploring approaches for service composition in sensor networks. Service composition is a fundamental method for creating advanced functionality by combining a set of primitive services provided by the system. In sensor networks, resources are constrained and communication among nodes is error-prone and unreliable. Such a dynamic environment requires a continuous adaptation of the composition of services. A team of researchers from IBM and RPI proposed a graph-based model of sensor services that mapped to the operational model of sensor networks and was amenable to analysis [42]. Based on this model, they formulated the process of sensor service composition as a cost-optimization problem that is NP-complete. They proposed two heuristic methods for its solution, a top-down and a bottom-up, and discussed their centralized and distributed implementations.

The team, augmented by further collaboration with ARL also explored other approaches for service composition, including approaches for using policy-based techniques [43] and auction based mechanisms [44].

#### References

- Y. Hou, M. Zafer, K. Lee, D. Verma, and K. Leung, "On the mapping between logical and physical topologies", in Proc. *IEEE Communication Systems and Networks and Workshops*, COMSNETS 2009.
- [2] S. Wang, R. Urgaonkar, T. He, M. Zafer, K. Chan, and K. Leung, "Mobility-induced service migration in mobile micro-clouds", in Proc.

IEEE MILCOM, 2014.

- [3] S. Wang, R. Urgaonkar, T. He, M. Zafer, K. Chan, and K. Leung, "Dynamic Service Migration in Mobile Edge-Clouds", in Proc. *IFIP Networking*, 2015
- [4] P. Novotny, R. Urgaonkar, A. Wolf, and B. Ko, "Dynamic placement of composite software services in hybrid wireless networks", in Proc. *IEEE MILCOM*, 2015.
- [5] E. Ciftcioglu, K. Chan, R. Urgaonkar, S. Wang, and T. He, "Securityaware service migration for tactical mobile micro-clouds", in Proc. *IEEE MILCOM*, 2015.
- [6] M. Dehghan, A. Seetharamz, T. He, T. Salonidis, J. Kurose, and D. Towsley, "Optimal caching and routing in hybrid networks", in Proc. *IEEE MILCOM*, 2014.
- [7] C. Bisdikian, "On Sensor Sampling and Quality of Information: A Starting Point," in Proc. IEEE PerCom Workshop, 2007.
- [8] C. Bisdikian, L. M. Kaplan, M. B. Srivastava, D. J. Thornley, D. Verma, and R. I. Young, "Building Principles for a Quality of Information Specification for Sensor Information," in Proc. *International Conference* on Information Fusion, FUSION 2009.
- [9] C. Bisdikian, L. M. Kaplan, and M. B. Srivastava, "On the quality and value of information in sensor networks", in ACM Transactions on Sensor Networks (TOSN), vol. 9, issue 4, 2013.
- [10] C. Bisdikian, "Quality of information trade-offs in the detection of transient phenomena", in Proc. SPIE Defense and Security Conference, 2007
- [11] S. Zahedi, and C. Bisdikian, "A framework for QoI-inspired analysis for sensor network deployment planning", in Proc. International Conference on Wireless Internet, WICON 2007.
- [12] T. He, M., Zafer, and C. Bisdikian, "Detection of Transient Signals with Incomplete Observations", in Proc. *IEEE MILCOM*, 2008.
- [13] Z. Charbiwala, S. Chakraborty, S. Zahedi, Y. Kim, T. He, C. Bisdikian, and M. B. Srivastava, "Compressive Oversampling for Robust Data Transmission in Sensor Networks", in Proc. *IEEE INFOCOM*, 2010.
- [14] W. Wei, T. He, C. Bisdikian, D. Goeckel, and D. Towsley, "Target

Tracking with Packet Delays and Losses – QoI amid Latencies and Missing Data", in Proc. *IQ2S*, IEEE Percom Workshops, 2010.

- [15] M. Srivatsa, P. Rohatgi, S. Balfe, and S. Reidt, "Securing Information Flows: A Metadata Model", in *IEEE Workshop on Quality of Information* (QoI) for Sensor Networks (QoISN), September 2008.
- [16] S. Chakraborty, K. Raghavan, M. Srivastava, C. Bisdikian, and L. Kaplan, "Balancing value and risk in information sharing through obfuscation", in Proc. *International Conference on Information Fusion* (FUSION), 2012.
- [17] C. Bisdikian, L. M. Kaplan, M. B. Srivastava, D. J. Thornley, D. Verma, and R. I. Young, "Building Principles for a Quality of Information Specification for Sensor Information", in Proc. *International Conference* on Information Fusion (FUSION), 2009. (An extended version of this paper is included in the book "Network Science for Military Coalition Operations: Information Exchange and Interaction" [D. Verma, ed.], IGI 2010.)
- [18] D. Verma, W. Leland, T. Pham, A. Swami and G. Cirincione, "Advances in network sciences via collaborative multi-disciplinary research", in Proc. *IEEE International Conference on Information Fusion*, 2015.
- [19] S. Eswaran, A. Misra, and T. F. La Porta, "Utility-Based Adaptation in Mission-oriented Wireless Sensor Networks", in Proc. *IEEE SECON*, 2008
- [20] S. Eswaran, A. Misra, T. F. La Porta, and K. Leung, "Providing Rapid Adaptation for Dynamic Missions in Mobile Wireless Sensor Environments", in Proc. SPIE Defense and Security Symposium, 2008.
- [21] S. Eswaran, M. P. Johnson, A. Misra and T. F. La Porta, "Distributed Utility-Based Rate Adaptation Protocols for Prioritized, Quasi-elastic Flows", in *ACM SIGMOBILE Mobile Computing and Communications Review* (MC2R), vol. 13, issue 1, January 2009.
- [22] S. Eswaran, M. P. Johnson, A. Misra and T. La Porta, "Adaptive In-Network Processing for Bandwidth and Energy Constrained Mission-Oriented Multi-hop Wireless Networks", in Proc. *IEEE/ACM International Conference on Distributed Computing in Sensor Systems* (DCOSS), 2009.
- [23] S. Eswaran, A. Misra, and T. La Porta, "Control-theoretic Optimization of Utility over Mission Lifetimes in Multihop Wireless Networks", in Proc.

IEEE SECON, 2009.

- [24] J. Tichogiorgos, K. Leung, A. Misra and T. La Porta, "Distributed Network Utility Optimization in Wireless Sensor Networks Using Power Control", in Proc. *IEEE PIMRC*, 2008.
- [25] M. P Johnson, D. Sarioz, A. Bar-Noy, T. Brown, D. Verma, and C. W. Wu, "More is More: the Benefits of Dense Sensor Deployment", in Proc. *INFOCOM*, 2009.
- [26] A. Bar-Noy, T. Brown, M. Johnson, and O. Liu, "Cheap and Flexible Sensor Coverage", in Proc. IEEE/ACM International Conference on Distributed Computing in Sensor Systems (DCOSS), 2009.
- [27] D. Verma, C. Wu, T. Brown, A. Bar-Noy, S. Shamoun, and M. Nixon, "Application of halftoning algorithms to location sensitive sensor placement", in Proc. *IEEE International Symposium on Circuits and Systems* (ISCAS), 2009.
- [28] M. Gomez, A. Preece, M. P. Johnson, G. de Mel, W. Vasconcelos, C. Gibson, A. Bar-Noy, K. Borowiecki, T. La Porta, D. Pizzocaro, H. Rowaihy, G. Pearson, and T. Pham, "An Ontology-Centric Approach to Sensor-Mission Assignment", in Proc. *International Conference on Knowledge Engineering and Knowledge Management* (EKAW), 2008.
- [29] A. Bar-Noy, T. Brown, M. P. Johnson, T. F. La Porta, O. Liu, and H. Rowaihy, "Assigning Sensors to Missions with Demands", in Proc. *ALGOSENSORS*, 2007.
- [30] H. Rowaihy, M. P. Johnson, A. Bar-Noy, T. Brown, and T. F. La Porta, "Assigning Sensors to Competing Missions", in Proc. *IEEE Globecom*, 2008.
- [31] M. P. Johnson, H. Rowaihy, D. Pizzocaro, A. Bar-Noy, S. Chalmers, T. F. La Porta, and A. Preece, "Frugal Sensor Assignment", in Proc. IEEE/ ACM International Conference on Distributed Computing in Sensor Systems (DCOSS), 2008
- [32] M. P. Johnson, H. Rowaihy, D. Pizzocaro, A. Bar-Noy, S. Chalmers, T. La Porta, and A. Preece, "Sensor-Mission Assignment in Constrained Environments", in *IEEE Transactions on Parallel and Distributed Systems*, vol. 21, issue 11, pp. 1692–1705, 2010.
- [33] M. Sensoy, T. Le, W. Vasconcelos, T. J. Norman, and A. Preece, "Resource Determination and Allocation in Sensor Networks: A Hybrid

Approach", in The Computer Journal, 2010.

- [34] A. Preece, T. Norman, G. de Mel, D. Pizzocaro, M. Sensoy, and T. Pham, "Agilely Assigning Sensing Assets to Mission Tasks in a Coalition Context", in *IEEE Intelligent Systems*, vol. 28, issue 1, pp. 57–63, 2013.
- [35] A. Mowshowitz and G. Bent, "Formal Properties of Distributed Database Networks", in Proc. Annual Conference of the International Technology Alliance, 2007.
- [36] A. Mowshowitz, V. Mitsou, and G. Bent, "Models of Network Growth by Combination", in IEEE/ACM Transactions on Networking, 2010.
- [37] M. Dehmer, and A. Mowshowitz, "On Measuring the Complexity of Sets of Graphs Using Graph Entropy", in Advanced Computational Technologies, Romanian Academy Press.
- [38] G. Bent, P. Dantressangle, P. Stone, D. Vyvyan, and A. Mowshowitz. "Experimental evaluation of the performance and scalability of a dynamic distributed federated database", in Proc. *Annual Conference of the International Technology Alliance*, 2009.
- [39] S. Zahedi, M. Szczodrak, P. Ji, D. Mylaraswamy, M. B Srivastava, and R. Young, "Tiered Architecture for On-Line Detection, Isolation and Repair of Faults in Wireless Sensor Networks", in Proc. *IEEE MILCOM*, 2008.
- [40] R. Zhang, Rui, P. Ji, D. Mylaraswamy, M. Srivastava, and S. Zahedi. "Cooperative sensor anomaly detection using global information", in *Tsinghua Science and Technology*, vol. 18, no. 3, 2013.
- [41] S. Tati, S. Rager, B. Ko, G. Cao, A. Swami and T. La Porta. "netCSI: A generic fault diagnosis algorithm for large-scale failures in computer networks", in Proc. *IEEE Reliable Distributed Systems (SRDS)*, 2011.
- [42] S. Geyik, B. Szymanski, P. Zerfos, and D. Verma, "Dynamic composition of services in sensor networks", in Proc. *IEEE International Conference* on Services Computing (SCC), 2010.
- [43] R. Dilmaghani, S. Geyik, K. Grueneberg, J. Lobo, S. Shah, B. Szymanski, and P. Zerfos, "Policy-aware service composition in sensor networks", in Proc. *IEEE International Conference on Services Computing* (SCC), 2012.
- [44] L. Chen, Z. Wang, B. Szymanski, J. Branch, D. Verma, R. Damarla and J. Ibbotson, "Dynamic service execution in sensor networks", in *The Computer Journal*, 2009.

# 5 Human and Cognitive Issues in Coalitions

Coalition networks are complex collections of human and machine components interacting in many varied contexts, but with the aims of the coalition aligned for particular tasks or activities. Throughout our ITA research the overall challenge has been to develop the fundamental science to underpin this two-way end-to-end socio-technical chain reaching forward from *data-to-decision* and back from *decision-to-data*, resolving complex problems while operating in this coalition environment. Agility has also been a key factor, as well as recognizing that these environments need to operate at the network edge. This is a resource constrained environment with potentially unstable structures, and with human agents who need to have their cognitive load respected, to avoid information overload and mitigate against cognitive biases.

In this chapter we cover the activities of the ITA research programme that have been directly aligned with the complex intersection between human behaviour, human cognition and machine capabilities, both for individuals and groups. Work done by the researchers spans a wide set of disciplines and includes advances in areas such as argumentation theory, collaborative technologies, behaviour simulation, knowledge technologies, shared understanding, information extraction and cognitive modelling.

Some of the major advances made during the course of ITA research are summarized below.

#### Advances in Argumentation Theory

Argumentation theory is the interdisciplinary study of how conclusions can be reached through logical reasoning, i.e., it provides a way to measure whether a claim is based soundly (or otherwise) on the premises. Argumentation theory enables an intelligence analyst to determine whether or not a conclusion is supported by the known facts. In a coalition environment, where different sources have a varying amount of trust that can be ascribed to them, the cognitive load on analysts and decision makers is tremendous. Advances in argumentation theory enable the creation of software agents and other tools that can reduce that cognitive burden.

Researchers from University of Aberdeen, Carnegie Mellon University (CMU), ARL and IBM explored new technologies to incorporate information with uncertainty into argumentation and provenance frameworks. They provided the theoretical glue to accommodate trust, inconsistency and uncertainty in distributed networked information systems. They applied argumentation based reasoning to reorganize facts, answers, and their linkages in a manner that mimics the human mental model. They characterized trust and uncertainty probabilistically using Dempster-Shafer theory, and proposed schemes that can prune the argumentation tree so that only the most relevant arguments and concerns appropriate to the human operator are presented to him or her [1].

Researchers from University of Aberdeen and CMU were the first ones to propose a principled method [2] for linking provenance data with the analysis of competing hypotheses through argumentation schemes. They leveraged existing standards for representing provenance, and enabled analysts to evaluate hypotheses while exploring and assessing the provenance of supporting evidence. While the prior state of the art dealt with approaches to reasoning with uncertain data, there was not sufficient attention paid to reasoning about uncertainty itself. In order to build a complete extensible system for uncertain reasoning, one needs the ability to argue with uncertain data, and the ability to argue about uncertainty. ITA researchers proposed a scheme that addressed this gap [3] and provided the principles that can be used to create argumentation systems that combine logical reasoning and Dempster-Shafer theory.

Further enhancing the ability for argumentation systems to deal with uncertain information, the researchers invented a new method for integrating probabilities and argumentation [4]. Their approach was based on the concept of *Markov Random Fields*. This allowed a seamless integration of symbolic argumentation with probabilistic graphical models that enables forming of hypotheses from inconsistent, uncertain and incomplete facts. It provided a rigorous mathematical foundation towards a computation model that enables effective decision support for linchpin analysis—where the state of the art was limited to using a list of human actionable guidance without a formal model.

# Shared Understanding

Issues of understanding, sensemaking and shared situation awareness have been a common feature of human factors research for the past several decades, as have their collective counterparts: shared understanding, collective sensemaking and shared situation awareness. Within a military coalition context, the notion of shared understanding has emerged as a concept of considerable significance,

Despite the general consensus that shared understanding is important to military coalitions, and that our efforts to configure and engineer the socio-technical environment should be geared to deliver shared understanding, it was clear from the outset that conceptual and scientific progress in this area was fraught with difficulty. Perhaps the biggest problem surrounded the attempt to understand what the notions of understanding and shared understanding actually mean. It was not at all clear whether there is any consensus on what we mean by the term *shared understanding* or how the concept should be differentiated from ostensibly similar concepts in the literature such as collective sensemaking, shared situation awareness, and shared mental models.

One of the key achievements of ITA was in defining the concept of shared understanding precisely, and differentiating it from other existing concepts in the literature.

The attempt of ITA research staff to understand the notions of understanding and shared understanding ultimately resulted in an ability-based conception of understanding [5]. This conceptualization casts understanding as something that is akin to an ability—to understand, on this view, is to be able to do certain things and respond in appropriate ways. The advantage of this ability-based conception of understanding is that it is sufficiently broad to cover a variety of domains in which the notion of understanding is used. Human perception, for example, has been argued to depend on more than just an ability to detect stimuli; it also depends on an ability to make sense of them—to understand them.

The ability-based conception is particularly useful when it comes to understanding the notion of situation awareness. In particular, the ITA researchers [5] suggest that situation awareness should be conceived of as a particular form of ability, namely *dynamic situational understanding*. The object of understanding in this case is, somewhat obviously, a situation, and the performances that manifest dynamic situational understanding are those typically mentioned in relation to situation awareness; i.e., the description of elements of the situation, the provision of explanations as to how the current situation emerged, and predictions as to how the current situation is likely to evolve across time.

ITA research provides this theoretical integration of the notions of understanding, situation awareness and mental models: mental models support the expression of behaviours that warrant the ascription of situation awareness to an agent, and situation awareness is a particular form of understanding, namely dynamic situational understanding.

One advantage of this revised conceptualization of situation awareness is that it helps to resolve a long-standing dispute regarding the state/process duality of situation awareness—the tendency for situation awareness to sometimes be regarded as a state and at other times as a process [6]. By casting situation awareness as a form of understanding, we can see that situation awareness is, in fact, neither a state nor a process. It is neither a state nor a process because understanding is akin to an ability, and abilities are not the sort of things that can properly be conceived of as states or processes [7]. Inasmuch as this thesis is true, it suggests that advocates on both sides of the state/process duality debate for situation awareness are mistaken.

It follows from the ability-based conception of understanding that shared understanding is something akin to a shared ability; one in which multiple individuals are able to exercise the same sort of predictive and explanatory capabilities regarding a common situation. Such abilities are likely to rely on more than just an exposure to common bodies of information; they also depend on the possession of common bodies of background knowledge, experience and training.

The implication for future work is that military coalitions should invest considerable effort in capturing the background knowledge and expertise of military staff. It also focuses attention on the manner in which technologies work to support the explanatory and predictive capabilities of human subjects, the way in which these capabilities are monitored, and the way in which specific explanations and predictions are represented and communicated.

Finally, in terms of the possibility of machine understanding, the ability-based conception of understanding focuses attention on the kind of processes that would be required to yield performances warranting the ascription of understanding to a synthetic agent.

Given the potential role of such algorithms in yielding biologically-based forms of understanding in the perceptual and linguistic domains, it is plausible that similar approaches could prove useful in terms of advancing the state-of-the-art in respect of the machine-based understanding of natural language, sensor information and environmental situations.

# **Controlled English**

A controlled natural language features a limited number of grammatical rules that constrain the interpretation and generation of natural language sentences. While controlled natural languages were known prior to the beginning of the ITA research, the needs for coalition military usage, e.g. a collaborative planning model, require a representation that was customized for its needs. The representation that was required needed to be more accessible than traditional forms such as *Web Ontology Language* (OWL), yet retain sufficient structure and constraints to enable efficient processing in a computer.

The design of the right controlled natural language has a number of concerns for knowledge capture and representation. Proficiency in controlled natural language may require excessive training of the end user. The representation of controlled natural language may be easily parsable by a computer, but if not done appropriately, it may be difficult for a human to comprehend.

A collaborative effort between ARL and IBM resulted in a representation of English [8] that did not suffer from the above limitations. It was designed to easily permit creation of visual tools that reduce the need for training and be easy to comprehend by humans. This representation, *Controlled English* (CE), proved to be a very useful construct, and was successfully used in a variety of application domains by other ITA researchers including sensor mission assignment [9] and information extraction [10].

The current form of the CE language is rich enough to allow complex semantics based on *First Order Predicate Logic* to be expressed within the model both in terms of concepts, properties, relationships and logical inference rules, although there is much opportunity for further enhancements and extensions to the language to improve human readability and also to extend semantic expressivity. There is a compatibility with the *Web Ontology Language Description Logic* (OWL DL) that is achieved through translation (to or from) OWL DL formats as required, and the term *conceptual model* in CE is synonymous with the term *ontology* in the Semantic Web community, although it is interesting to note that CE also contains rules as part of the core language whereas OWL requires extensions (such as SWRL or RIF, the *Rule Interchange Format*) in order to define rules that aren't

part of core OWL.

A major motivation for the CE language was to create a pervasive language representation into which a human user can contribute information about many aspects of their environment. Already mentioned are the ability to create conceptual models and associated information, and the ability to define logical inference rules that utilize and extend the semantic features of these models. In addition the CE language supports a query syntax (similar to the rule syntax, but without conclusion clauses), a reification syntax in which statements about CE statements can be made and the concept of rationale, which is used to explain the premises and rule(s) that were used to infer a specific conclusion and can encompass hypotheses and assumptions. Rationale is not discussed in detail here, but is a very powerful technique that can be used for multiple purposes, not least to help the human users of the system gain an increased trust in the conclusions generated through clear visibility of the rationale graph which allows them to consider all factors involved in the computation of new information. In addition to this the CE language has an in built annotation syntax to allow complex annotations to be linked with CE sentences to aid human understanding and readability as well as a command syntax to allow invocation of specific CE-related commands in a given sequence (for example the loading of sets of sentences or execution of rules or agents).

One minor but extremely important factor with CE is that the human reader and machine agent both consume and process *exactly the same information*. The CE is both human-readable and directly processable (without translation into an intermediate technical format such as XML) by the machine agent. Many other solutions in which a human-friendly representation of information is used will translate this human-friendly format into another technical format before machine processing, introducing a layer for ambiguity and/or misunderstanding to creep in.

With an accompanying implementation of CE Store (page 101), in addition to the fundamental research achievements the CE technology has proved invaluable for several transitions from the ITA programme.

# Advances in Cognitive Modelling

High fidelity cognitive models that are of predictive and explanatory relevance to problem-solving behaviours are applicable to a variety of coalition challenges. The aim of these models is to go beyond mapping the differences in the cognitive properties of groups of agents and to develop models that could be used to predict cognitive performance in context of a specific task. Computer simulation techniques can provide a means to study collective cognitive processes in taskoriented coalition teams, and such cognitive models satisfy the ever-growing need to enhance the cognitive sophistication of the agents and the fidelity of the computer simulations.

The basis for cognitive models for computer simulations was explored in the early parts of the ITA programme with the use of computer simulations featuring relatively simple cognitive agents [11]. Towards the end of the ITA, the focus shifted towards the development of computational cognitive models; i.e. cognitive models that could be used in the context of computer simulation studies. This work initially focused on the computational modelling of specific cognitive processes, such as cognitive dissonance, as a means to understand the interaction between cognitive and social factors in mediating cognitive change [12].

Although the early work [11] served as a useful starting point, it suffered from a number of weaknesses, including the fact that the simulations only targeted a specific kind of cognitive process and they only supported analyses of the interplay between two factors; i.e. the social and the cognitive. These weaknesses prompted a shift towards the use of more sophisticated cognitive agents and an emphasis on additional factors, e.g. technological and informational factors. The result was a series of research and development efforts centred on the use of the ACT-R cognitive architecture [13].

By providing a computational framework for the modelling of human cognitive processes across a range of task contexts, ACT-R helped to address concerns regarding the limited scope of early ITA computer simulations. The use of ACT-R also helped to improve the cognitive sophistication and fidelity of the cognitive agents used in computer simulation experiments. Given the interest in studying cognition in a social context, ITA researchers adopted a computer simulation technique known as cognitive social simulation [14]. This approach is used to study issues of collective performance in tasks that feature the interaction of multiple cognitive agents that are implemented using cognitive architectures.

By relying on the technique of cognitive social simulation, and by developing cognitive models built on top of the ACT-R architecture, ITA researchers were able to develop simulation capabilities in which multiple cognitive agents worked together on a particular collaborative problem solving task [15,16,17,18].

Issues of fidelity and realism in the context of computer simulation studies

continued to be a driving force in terms of the profile of cognitive modelling work throughout the final stages of the ITA programme. The most recent manifestation of this *quest for fidelity* is the use of virtual environments as a means of performing cognitive social simulation experiments. ITA researchers have thus developed a framework for integrating ACT-R with the *Unity* game engine, and tested the framework with virtual robotic platforms that are controlled by ACT-R cognitive models [19]. The use of virtual environments, here, provides a number of relatively new opportunities (and challenges) for the computational study of the sort of extended, embedded and distributed cognitive systems that have been the focus of recent work in empirical and theoretical cognitive science. In particular, the combined use of cognitive architectures and virtual environments may provide a means to study many of the ideas that were first mooted in the context of early work relating to cognitive systems in the ITA. This includes efforts to situate notions of cognitive extension within the socio-technical ecology of the military coalition environment [20,5].

### **Cultural Network Analysis**

A difference in culture is to be expected when coalition forms, since different people come with a widely different background and culture. An understanding of the cultural differences among different groups would be very useful in making different groups collaborate better. ITA researchers from Applied Research Associates and University of Southampton came up with an innovative approach to understand cultural differences between the U.S. and UK armies [21,22].

The researchers proposed a rigorous, end-to-end methodology for modelling culture as networks of ideas that are distributed among members of a population. The method, *Cultural Network Analysis* (CNA), represents an interdisciplinary synthesis of techniques drawn from the fields of cognitive anthropology, cultural and cognitive psychology, naturalistic decision making, and decision analysis.

Cultural network analysis is a technique that stems from work in cultural epidemiological theory. Its aim is to construct mental models based on the data supplied from a sample of individuals. The models themselves are typically represented as a network-based representation of the concepts, causal beliefs and values that are shared by a cultural group. Such models are used as tools to assist with the explanation and prediction of behaviour within cultural groups, and they also provide some indication as to the extent of the cognitive differences that exist between groups. The use of network-based representations is important here because the degree of structural isomorphism across a set of models provides

some indication as to the cognitive similarity of the target cultural groups.

CNA is used to develop cultural models for groups and populations, typically depicted as a network representation of the culturally shared concepts, causal beliefs, and values that influence key decisions. CNA can be usefully employed for a variety of applications, including the design of tools to support multinational collaborative planning and decision making, the development of situated cultural training programmes, and characterizing the cognition of target audiences to support strategic communications campaigns.

One example of the application of cultural network analysis to US/UK military coalitions within the context of the ITA was done by a team from Applied Research Associates and University of Southampton [23]. They attempted to understand the cognitive differences between US and UK military planners when it came to evaluations as to what makes a good plan. The results of this work revealed a number of important differences between US and UK planners; in particular, the ideal military plan for a British planner was one which served an *enabling function*, whereas the ideal military plan for an American planner was one which served a *controlling function*. The five areas of differences in the mental models of planning between the two countries are shown in Figure 21. These differences could result in misunderstanding and miscommunication within coalition planning teams.



Figure 21. Cultural difference in planning between US and UK

When developed, CNA was a very different method than the existing methods for understanding cultural differences. Those methods included ethnographic methods focusing on qualitative analysis of single cultural group, or psychological methods that attempt to capture cultural differences in a few general dimensions such as individualism or collectivism. Those methods were limited in their ability to capture and represent cultural commonalities and lacked the precision needed to enable the design of complex cognitive systems. CNA addressed these gaps.

CNA was an instance of cognitive modelling done by ITA in the earlier years, where the focus was on analysing differences in coalition practices. Similar cultural modelling was done to analyse other differences in coalition practices described below.

#### Analysis of Coalition Differences

Coalitions bring together humans and organizations that have evolved independently. As a result, there is a pronounced difference in the way members of a coalition think, act and behave. These differences manifest themselves in many different ways, and could result in operational inefficiencies. Therefore, one of the key focus areas of ITA research has been to understand differences between different groups that make up a coalition.

ITA researchers have analysed a variety of differences in the different teams that make up a coalition, including differences in processes, difference in linguistics, and differences in the semantic information used by different teams.

The research [24,25] on conversations between US and UK coalition partners conducted by ITA researchers from The Boeing Company, Dstl, ARL, and Systems Engineering & Assessment Ltd showed that there are five main types of linguistic sources of miscommunication among coalition partners. These are the use of acronyms and jargon, use of slang and colloquialisms, miscommunication associated with the medium of communications, misinterpreted speech, and denotation versus connotation in languages. Many of the relevant issues were largely pragmatic in nature. This was the first team to use computational pragmatics to analyse miscommunication in coalition contexts.

ITA researchers from Applied Research Associates and The Boeing Company developed a top-down approach for modelling and analysing the similarities and differences between related complex processes [26,27]. The method was derived from coordination theory, which is the general body of theory regarding how people or software agents coordinate their activities. Coordination theory states that collaboration occurs in order to manage the dependencies between tasks. Different groups of people performing the same tasks may choose different
ways for coordinating them. These choices can yield work processes that appear widely divergent, even though their purposes are essentially identical. The use of coordination theory makes both the similarities and differences readily visible across processes.

The method developed included use of trade-off matrices to compare alternative coordination mechanisms and a derivation tree to reveal the overall similarities and differences between different complex processes. This method would be very useful to coalition members who must work together despite differences in language, culture, policies, and organizational processes. The researchers validated the method by applying it in the military decision making domains as well as in engineering change management.

Another type of difference that can frequently arise in military coalitions is a difference in ontologies. Ontologies can be used for a variety of human computer interactions, and in several applications of semantic web technologies. It is highly likely that different coalitions would develop their own ontologies for various functions. Even when the ontologies target the same domain of discourse and are semantically similar, there would be a difference between them. This difference can hamper interoperability.

ITA researchers at IBM and University of Southampton investigated techniques to address the difference in ontologies, and developed a portable ontology alignment solution [28]. The solution isolates a fragment of source ontology that is relevant to an ontology mapping, uses information available to identify and capture the fragments that are relevant for use in a specific context. Thus, it minimizes the amount of ontology information that needs to be used, and also avoids spurious differences that are not relevant to the specific use-case.

# **Collaborative Planning Model**

Planning is a core military activity, which is carried out by a large, culturally diverse, hierarchical and geographically distributed teams each with different specializations. At the beginning of the ITA research, planning tools had found only limited traction. Among the different reasons that were identified for the lack of adoption, the most prominent was the unwillingness of humans to give up the planning process, preferring tools that only automated repetitive tasks.

ITA researchers from IBM, Dstl, ARL, and The Boeing Company tried to develop technologies that could help alleviate these challenges, which led to the

development of a collaborative planning tool [29] that did not suffer from these issues. Recognizing the fact that a single planning tool is not appropriate for all members of the distributed team, they opted for creating a representation of the plan [30] together with its associated artefacts, which allows the communication of a shared representation and understanding of the plan.

In addition to supporting the plan and the artefacts, the collaborative planning model also supported the concept of rationale [31]. Incorporating the rationale information improves the usability, evaluation and communication of plan-related decisions to different echelons, branches or coalition partners. It also emphasized the importance of a variety of other types of information, e.g. assumptions, provenance, certainty, trust etc.

Thus, the research into collaborative planning model led to insights into the principles that can lead to the development of effective collaboration planning tools in the future.

### References

- [1] M. Sensoy, G. de Mel, A. Fokoue, T. Norman, J. Pan, Y. Tang, N. Oren, K. Sycara, L. Kaplan and T. Pham, "Reasoning with uncertain information and trust", in Proc. SPIE Defense, Security, and Sensing, 2013.
- [2] A. Toniolo, F. Cerutti, N. Oren, T. J. Norman and K. Sycara, "Making informed decisions with provenance and argumentation schemes", in Proc. *International Workshop on Argumentation in Multi-Agent Systems*, May 2014.
- [3] Y. Tang, N. Oren, S. Parsons and K. Sycara, "Dempster-Shafer Argument Schemes", in Proc. International Workshop on Argumentation in Multi-Agent Systems, May 2013.
- [4] Y. Tang, A. Toniolo, K. Sycara, and N. Oren, "Argumentation Random Field", *International Workshop on Argumentation in Multi-Agent Systems*, May 2014.
- [5] R. Smart, T. Huynh, D. Mott, K. Sycara, D. Braines, M. Strub, W. Sieck, and N. Shadbolt, "Towards an Understanding of Shared Understanding in Military Coalition Contexts", in Proc. *Annual Conference of the International Technology Alliance*, 2009.
- [6] R. Rousseau, S. Tremblay, and R. Breton, "Defining and modeling

situation awareness: A critical review", in *A Cognitive Approach to Situation Awareness: Theory and Application*, S. Banbury and S. Tremblay (Eds.), Ashgate Publishing Company, Aldershot, England, UK, 2004

- [7] G. Baker, and P. Hacker, "Understanding and Meaning: Essays Pt. 1", Analytical Commentary on the Philosophical Investigations, Blackwell Publishers Ltd, Oxford, UK, 1980.
- [8] D. Braines, D. Mott, S. Laws, G. de Mel, T. Pham, and C. Giammanco, "Controlled English to facilitate human/machine analytical processing", in Proc. SPIE Defense, Security and Sensing, 2013.
- [9] A. Preece, D. Pizzocaro, D. Braines and D. Mott, "Tasking and sharing sensing assets using controlled natural language", in SPIE Defense. Security, and Sensing, 2012.
- [10] D. Mott, D. Braines, S. Poteet, A. Kao, and P. Xue, "Controlled Natural Language to facilitate information extraction", in Proc. Annual Conference of the International Technology Alliance, 2012.
- [11] P. R. Smart, T. D. Huynh, D. Braines, and N. R. Shadbolt, "Dynamic Networks and Distributed Problem-Solving", in Proc. *Knowledge Systems* for Coalition Operations (KSCO), 2010.
- [12] P. R. Smart, and K. Sycara, "Collective Sensemaking and Military Coalitions", in *Intelligent Systems*, vol. 28, issue 1, pp. 50-56, 2013.
- [13] J. R. Anderson, D. Bothell, M. D. Byrne, S. Douglass, C. Lebiere, and Y. Qin, "An integrated theory of the mind", in *Psychological Review*, vol. 111, issue 4, pp. 1036–1060, 2004.
- [14] R. Sun, "Cognitive social simulation incorporating cognitive architectures", in *Intelligent Systems*, vol. 22, issue 5, pp. 33-39, 2007.
- [15] P. R. Smart, and K. Sycara, "Cognitive Social Simulation and Collective Sensemaking: An Approach Using the ACT-R Cognitive Architecture", in Proc. International Conference on Advanced Cognitive Technologies and Applications (COGNITIVE), 2014.
- [16] P. R. Smart, D. P. Richardson, K. Sycara, and Y. Tang, "Towards a Cognitively Realistic Computational Model of Team Problem Solving Using ACT-R Agents and the ELICIT Experimentation Framework", in Proc. International Command and Control Research Technology Symposium (ICCRTS), 2014.

- [17] P. R. Smart, Y. Tang, P. Stone, K. Sycara, S. Bennati, C. Lebiere, D. Mott, D. Braines, and G. Powell, "Socially-Distributed Cognition and Cognitive Architectures: Towards an ACT-R-Based Cognitive Social Simulation Capability", in Proc. Annual Conference of the International Technology Alliance, 2014.
- [18] Y. Tang, C. Lebiere, K. Sycara, D. Morrison, P. R. Smart, "Cognitive and Probabilistic Models of Group Decision Making" in Proc. *Conference on Behavior Representation in Modeling and Simulation*, 2015.
- [19] P. R. Smart, and K. Sycara, "Situating Cognition in the Virtual World" in Proc. International Conference on Applied Digital Human Modeling, 2015.
- [20] P. R. Smart, and K. Sycara, "Place Recognition and Topological Map Learning in a Virtual Cognitive Robot", in Proc. *International Conference* on Artificial Intelligence, 2015.
- [21] W. Sieck, L. Rasmussen, and P. Smart, "Cultural network analysis: A cognitive approach to cultural modeling", in *Network Science for Military Coalition Operations: Information Extraction and Interaction*, pp. 237–255, 2010.
- [22] W. Sieck, "Cultural network analysis: Method and application", Advances in cross-cultural decision making, pp. 260–269, 2010.
- [23] W. Sieck and J. Patel, "Cultural Issues in coalition planning", in Proc. International Conference on Knowledge Systems in Coalition Operations, IEEE Press, 2007.
- [24] S. Poteet, P. Xue, J. Patel, A. Kao, C. Giammanco, and I. Whiteley, "Linguistic sources of coalition miscommunication", in NATO RTO HFM-142 Symposium on Adaptability in Coalition Teamwork, 2008.
- [25] P. Xue, S. Poteet, and A. Kao, "Conversation Analysis of Coalition Communication in Network Centric Operations", in *Network Science for Military Coalition Operations: Information Exchange and Interaction*, 2010.
- [26] S. E. Poltrock, M. Klein and M. Handel, "Understanding process differences: Agreeing upon a single way to skin a cat", in Proc. *IEEE Conference on Knowledge Systems for Coalition Operations* (KSCO), May, 2007.
- [27] S. Poltrock and M. Klein, A coordination-theoretic model of the military decision-making process. In 1st Annual Conference of the International

Technology Alliance, 2007.

- [28] Y. Kalfoglou, P. R. Smart, D. Braines, and N. R. Shadbolt, "POAF: Portable Ontology Aligned Fragments", Proc. International Workshop on Ontologies: Reasoning and Modularity (WORM) hosted by the European Semantic Web Conference (ESWC), 2008.
- [29] A. Bahrani, J. Yuan, D. Chukwuemeka, D. Masato, T. Norman and D. Mott, "Collaborative and context-aware planning", in Proc. *IEEE Military Communications Conference* (MILCOM), 2008.
- [30] T. Klapiscak, J. Ibbotson, D. Mott, D. Braines, and J. Patel. "An Interoperable Framework for Distributed Coalition Planning: The Collaborative Planning Model", in Proc. *Knowledge Systems for Coalition Operations*, 2012.
- [31] D. Mott, and C. Giammanco, "The use of rationale in collaborative planning", in Proc. Annual Conference of the International Technology Alliance, 2008.

# 6 Exploiting the Science

One of the strategic aims of the NIS ITA was to achieve rapid and broad exploitation of the science in both defence and civil domains (Figure 22). The vision for the ITA sought to leverage the observation that "experience has shown that the development of new technologies is best achieved when governments work in close co-operation with technology providers in industry and academia from the very earliest stages of research." This was to be achieved by "a new way of doing business: forming an international alliance of government, industry and academia" [1].

This chapter addresses how the science conducted within the NIS ITA fundamental research component has been exploited. The discussion is naturally limited as some of the exploitations are of a sensitive nature. The chapter has three major sections: the first describes the types, speed and breadth of exploitation; the second describes some of the key publicly available (primarily open source) components; and the third presents a narrative description of a sub-set of the exploitations.



Figure 22. Collaborative science to research exploitation

# Enablers

The need to conduct both fundamental research and rapid exploitation, whilst complying with UK and US regulations, meant the ITA programme consisted of two parts:

- 1. A fundamental research component which undertook deeply collaborative research, the results of which would be published in the public domain
  - A technology transition component, consisting of separate UK (MOD/ Dstl) and US (ARL) agreements that would provide for the application of the fundamental research results to defence and security applications (and included the ability to address any security or export issues)

The existence of the transition component, and its associated contracts, provided an efficient route to enable the exploitation of the knowledge developed within the ITA research team. Its existence also acted as a positive incentive to the organizations within the ITA. This was reinforced due to the inclusion of transition opportunity assessment criteria within the fundamental research programme formulation and peer-review assessment. Moreover, the challenge-led nature of the research meant that, due to its relevance, it would be potentially exploitable.

# **Broad Achievements**

This section addresses the types, speed and breath of exploitation. It does not address the impact of the ITA science within the scientific domain, which is addressed elsewhere in the book.

# Enhanced S&T Capability

The relationships and ways of working developed within the ITA have enhanced both UK and US scientific and technical capability in government, industry and academia. As noted in the Foreword "[p]ossibly the most significant outcome of the NIS ITA has been the strong bonds that have formed between the US and UK leadership and researchers that will endure long after the programme ends." This is clearly an impact which cannot easily be measured, but it is supported by the peer-review assessment that the ITA is "an outstanding example of true, deep and enduring International Research Collaboration".

One part of this soft (as in hard to measure, but no less important) impact has been a thorough understanding of the unique challenges associated with defence coalition

operations, which has been built within the academic and industrial researchers through their direct interaction with government researchers and military subjectmatter experts (SMEs). For example, all-source analysts from ARL and Dstl played an integral part in ITA research and transitions. This increase in the researcher's knowledge and understanding means that the future research they undertake is likely to be more applicable to the military coalition domain; particularly as the military coalition domain often represents a more challenging case that, if solved, means the research can be applied to less challenging domains too.

Another aspect of the soft impact, as is shown in some of the detailed examples discussed below, is the way in which UK and US transition activities have been able to build upon each other due to these relationships and ways of working.

A more measurable impact is that the NIS ITA programme has supported 64 Ph.Ds. and 14 Master's. However, even in this case the story can be more complex and the impact higher, as over the 10 years of the programme researchers have moved through different roles and moved between organisations.

#### 10-years within the NIS ITA: a researcher's journey...

2006 to 2012: Ph.D. student at UK academic institution
2010: ITA summer intern at ARL and IBM Research (US)
2012 to mid-2013: Postdoctoral Research Fellow concurrently at a UK academic institution, ARL and IBM Research (US)
mid-2013 to 2015: Research Staff Member at IBM Research (US)
2016+: Research Staff Member at IBM Research (UK)

#### Research Exploitation: Broad Types and Measurement

In addition to the enhancement of S&T capability (including trained staff with enhanced awareness, established relationships and ways of working) there were more direct exploitations of the science created. These exploitations again contain softer (as in very hard to measure) and harder exploitations of the science.

The hard exploitations are instantiated as technologies, usually in software form, which are tested and developed within applied research and lead to products and services within the military and civil domains. While we can measure the number of transitions placed by Dstl and ARL using the technology transition component (or their monetary cost or level of effort in person-days) this is a crude metric that

does not capture their impact. Further, any such measurement would be a significant underestimate as it would ignore any exploitations that utilised a different route (including industry intra-mural research), or exploited ITA technology that has been placed into the public domain.

#### Public domain exploitation: an example

Dstl researchers attended a research output briefing by a UK defence prime contractor that was not part of the NIS ITA. At this briefing it became clear that the UK organization was utilizing an ITA technology placed by IBM in the public domain on IBM's developerWorks® site, and had derived significant value from it.

An element that can be measured is the number of patents submitted based on ITA research. At the time of writing, at least 53 patents based on ITA research have been filed.

The ultimate expression of ITA technologies that are within the public domain are the set of open-source software components and libraries that have been matured, tested and validated by the transition activities undertaken by different parties within the Alliance.

The softer exploitations are those where the *knowledge* generated by ITA research is exploited when used to inform the development of new concepts and systems.

#### Exploiting the knowledge: a Dstl example

Following the publication of Gen. Barrons paper on "Warfare in the Information Age" (WitIA), Dstl was tasked with supporting Joint Forces Command in developing a set of concepts that embraced its vision.

Dstl researchers working on the NIS ITA supported this task and drew on the knowledge and understanding they had of future socio-technical information systems to assure the scientific feasibility of the concepts developed, and that feasible disruptive concepts were considered.

*Knowledge* exploitations are those where it would be impossible to label the system as having "ITA technology inside". Instead the research may well have set performance bounds on a potential system—for example allowing effective decision making on the likely cost-effectiveness of pursuing improvements in

its performance—or identified that a highly disruptive change in the status quo was feasible. In these cases, it is possible to lay out the narrative of a particular exploitation but it is frequently hard to fully substantiate and impossible to capture them all.

### **Rapid Exploitation**

The traditional view of research and development timelines is that fundamental research would not actually appear within defence and civil sector products and services within a period of 10 years, and could easily take 30 years to mature. However, the fast moving nature of *Information Communications Technology* (ICT) research and development means that this is not true within a significant fraction of industry civil sector programmes. Component technology refresh is fast, with new generations coming along every 2–3 years (cf. computer processor technology), but major generational change is slower (cf. 4G networks), and significant paradigm changes are slower still (cf. mainframes to personal computers to thin client devices linked to the Cloud).

In keeping with this, ITA science has been exploited within a number of civil sector products and services. Where these products and services are used by defence then the natural update cycle of *commercial off-the-shelf* (COTS) technology means that US and UK defence are already using ITA technology within fielded military systems. An example is the utilisation of ITA policy management technologies within some *U.S. Department of Defense* (DoD), MOD and NATO network management systems.

Indeed ITA technology is inside a number of industry offerings to both the civil and defence sectors including several IBM products and technologies including Indeed ITA technology is inside a number of industry offerings to both the civil and defence sectors including *IBM*® *SmartCloud*® *Application Performance Management* (APM) and *IBM*® *PureApplication*® *Software*.

Similarly, when MOD procurement was responding to Urgent Operational Requirements during operations in Afghanistan, then ITA technology was exploited through military off-the-shelf (MOTS) routes. An example is the visual query builder that was developed during an applied research task assessing the utility and feasibility of exploiting ITA research into DDFDs as part of the response to the IED threat in theatre. The visual query builder became part of a MOTS solution which was procured and deployed to Afghanistan.

### **Broad Exploitation**

MOD and DoD have exploited research from across the full breadth of the NIS ITA research programme. At one extreme, ITA research that developed a methodology to undertake cultural network analysis (page 86) resulted in a 2-year applied research project (March 2010–March 2012), funded by the *Office of Naval Research*, into *Extremist Ideological Influences on Islamic Terrorist Decision Frameworks*. This project was undertaken by Applied Research Associates in the US and University of Southampton in the UK. At the other extreme, ITA research into network tomography (page 23) resulted in an applied research project for the *Defense Threat Reduction Agency* (DTRA) undertaken by ARL, IBM US and Penn State University.

# Public Domain and Open-Source Technologies

This section discusses the ITA technologies available for wide exploitation due to their availability within the public domain, in some cases as open-source software. It ignores the scientific advances which are similarly available.<sup>1</sup>

Making ITA technology available as open-source software has the great advantages of: making it widely available to defence and civil sector supply chain ecosystems, avoiding procurement programmes becoming locked into a single supplier, and supporting innovation.

ITA open-source software comprises:

- · Controlled English (CE), consisting of the CE Store and the CE Node
- Edgware Fabric
- Gaian Database

Before they became fully open-source, earlier versions were frequently placed in the public domain.

A number of other ITA technologies which are available within the public domain, but not available as open source, are available on the IBM developerWorks® web

<sup>1</sup> There are too many of these to list here. An example is the ITA research enhancements to the ACT-R model (page 84) that have been made available to the wider scientific community.

site, including the Watson Policy Management Language (WPML).2

The majority of the technology maturation, testing and validation required to enable ITA technologies to be made available within the public domain and/or as open-source software has occurred through transition tasks funded by the MOD and DoD: primarily by Dstl and *Defence Equipment and Support* (DE&S) in the UK, and by ARL and the CWP in the US.<sup>3</sup> However, not all of the development has been funded in this manner: a significant fraction has been achieved using other funding sources available to industry and academia.

#### **CE** Store

CE is an unambiguous subset of English that can be both directly processed by a machine and understood by a human. Thus both humans and machines can work with the same symbolic representation of the world; it is not necessary to transform the human representation into a computer language that can only be used by a very small number of technology experts.

The user describes a domain model in terms of concepts, properties and relationships, and then populates this model by stating facts and rules: all in the form of CE sentences. The CE Store stores the domain model and facts, and can reason using the rules to create new facts and insights.

The CE Store is an open-source technology—available on GitHub—with which CE can be created and experimented upon for the representation of knowledge and the application of reasoning. Fact extraction from natural language can also be performed using the features of CE and the CE Store. An *application programming interface* (API) is provided for programmatic agents that convert incoming text into sentences, parse those sentences into raw parse trees, and turn the parse trees into phrases, all expressed using CE sentences. CE rules are then applied to extract facts, their relationships, and their properties. Facts are expressed as CE sentences in the context of a domain model described by the user.

<sup>2</sup> WPML is based on ITA fundamental research and was matured by an ARL funded transition; it has been further exploited into a COTS product which is in use within DoD. MOD and NATO network management systems.

<sup>3</sup> The Coalition Warfare Program is an Office of the Secretary of Defense (OSD) initiative to support international cooperative technology development that enables coalition forces to operate more effectively across the full spectrum of multinational operations. The goals are to accelerate delivery of high-quality solutions for the warfighter, to improve US interoperability with coalition partners, and to strengthen global partnerships.

#### **CE** Node

CE Node is a lightweight CE processing environment implemented in JavaScript that can be easily deployed in a variety of contexts, including Web browsers, mobile apps, and servers. CE Node is lightweight in the sense that it does not aim to be a fully-fledged CE engine—for example offering only limited inference and natural language processing—and requires relatively little network bandwidth to download and operate. Once loaded, a CE Node instance can function independently without any network connection. This makes it well-suited to deployments at the network edge.

**UK Command** Users and Assets in the Area of US Command and Control Users and Control Users Operations \* \* \* \* \* \* \* \* . . 0 .... 1 Fabric Bus **UK Forward US** Forward Operating Base Users and Assets Operating Base Users and Assets

#### **Edgware Fabric**

Figure 23. Edgware Fabric logical bus

Edgware Fabric (the open source name of the ITA Information Fabric) is a lightweight agile service bus that provides many of the features found in an enterprise service bus (such as discovery, routing, a registry, and message transformation) but which is built for resource constrained, dynamic and/or unreliable environments. Thus it integrates systems at the very edge of the network into a service-orientated architecture running on (or alongside) the devices that it connects. It is designed to be self-managing; it tracks which systems are connected, what services they offer and when they are being used. The discovery protocol enables neighbouring nodes to quickly form into a bus that spans an ad hoc network of communicating nodes, and can be viewed logically as a bus, as illustrated in Figure 23. Actors simply request data from the bus for one or more

data feed services (hardware or software components), requiring no knowledge of the structure of the network itself.

New software services can be deployed on the bus, enabling local processing of information at or near its point of origin, saving valuable network bandwidth and helping to manage the large volumes of data that can be generated by edgedevices such as sensors.

Edgware Fabric enables information from the edge to be easily integrated into existing applications and used in new and innovative ways. It can be used with existing hardware (e.g. physical assets such as sensors), software and networking technologies.

#### The Gaian Database

Federating and aggregating information distributed across a coalition, or indeed within any organisation, is a major operational challenge. Doing so efficiently, transparently and with minimal management overhead has been an unachieved goal, particularly in resource constrained and dynamic environments. The Gaian Database addresses this challenge.



Figure 24. Expanded Gaian Database node

Gaian embodies the concept of a *dynamic distributed federated database* (DDFD). This is a self-organizing network of federated nodes that combines ideas from data federation, distributed databases, network topology and the semantics of data. It is an information virtualization middleware component that is ideally suited to the ad hoc queries and processing operations that are needed to maximize business intelligence.

Gaian uses a *store locally query anywhere* (SLQA) paradigm giving global access to data from any participating node. Moreover, Gaian makes it possible for a set of heterogeneous data sources to be accessed as a single federated database, including sources as diverse as SQL and non-SQL databases, document repositories, spreadsheets and text files (Figure 24).<sup>4</sup> Applications can transparently perform database queries across a multiplicity of data sources in a single operation.

Access to data and the flow of data can both be controlled using formal policy based mechanisms that provide fine-grained management of security constraints. This is achieved using distributed *policy enforcement point* (PEP) and *policy decision point* (PDP) components at all database nodes where policy is to be enforced. The concept is shown in Figure 25.



# Figure 25. Distributed policy management using WPML and the Gaian Database

<sup>4</sup> The federation of non-SQL databases was funded by a Dstl transition contract.

The Gaian Database uses an extension to the standard Kerberos protocol to maintain security and access control within its distributed environment.<sup>5</sup>

The key to achieving efficient query performance is the way in which the nodes logically connect themselves together in order to minimise the cost of performing the distributed database operations. The mechanism used is based on the ITA fundamental research on network growth and emergent graph properties. The attachment mechanism results in a connected graph structure that has predictable properties that directly impact on database query performance, and which scale efficiently with network size. Overall this ensures optimal performance with minimal overhead.

Gaian does not replace existing systems; instead it federates them in a transparent, scalable and secure manner. It introduces a new agile model of information integration that revolutionizes the way that coalitions and organizations can access and exploit the information held within their IT systems. Its small footprint and efficiency make it ideal for use everywhere from the enterprise to mobile and other constrained environments.

#### Other ITA Technologies-Some Examples

In April 2013 IBM US released, as open source, a *fully homomorphic encryption* (FHE) library with an open challenge to the academic community to improve its performance. ITA research into FHE has resulted in the placing of additional components into this open-source technology.

Several algorithms for tracking information flow release developed within the ITA have been open-sourced by University of Maryland as extensions to probabilistic programming languages.<sup>6</sup> These algorithms allow information release tracking over dynamic secrets, using methods of mixed-mode SMC computations, while guaranteeing aggregate information release and game-theoretic models of information release (between an information provider and information consumer).

In other instances where the ITA research has involved developments of existing open-source software, the ITA research has, naturally, been fed back into the open-source software concerned. An example of this is the ITA research that

6 https://github.com/plum-umd/qif

<sup>5</sup> The extension to Kerberos was funded by a joint UK and US transition project funded by ARL, CWP, and Dstl. This enabled demonstration of fine grained policy controlled access to distributed and federated data at the *NATO Intelligence Fusion Centre* (NIFC).

has extended the capabilities of ACT-R, which has been placed back within the ACT-R open-source community (page 84).

### Watson Policy Management Library

WPML is an implementation of the algorithms for policy analysis developed as part of the basic research within the ITA programme; the transition was funded by ARL. The library implements the basic framework for policy management that underlies the prevailing policy management paradigm within the community. This policy architecture consists of four components: a policy management tool, a policy repository, a PDP, and a PEP. The policy management tool provides a user interface for the creation and definition of policies. The policy repository provides mechanisms for storing policies and retrieving them as needed by the decision points. The PDPs are modules in the system responsible for selecting and evaluating policies stored in the repository. The PEPs are elements that are responsible for enforcing the outcome of those policy evaluations.

WPML provides a set of tools and APIs that implement the policy architecture. A generalized policy model able to support arbitrary policy languages sits at its core. Policies may be one of two types: authorization policies that provide true/false results, and obligation policies that enable conditional execution of some action. Finally, a language-specific policy string is used to determine the set of required instance data for the policy to be evaluated successfully.

WPML has been integrated with a number of other ITA software assets, including the Gaian Database and Edgware Fabric.

# **Exploitation Narratives**

This section contains a narrative description of a sub-set of the transitions, focussing on those funded by the DoD or MOD.

# The Pathfinder Series of Applied Research Projects

There were three *Pathfinder* applied research projects—funded by Dstl between January 2008 and July 2011—which tested the potential value of semantic technologies to intelligence analysts. The projects were a collaborative effort between IBM UK, University of Southampton, and LogicaCMG UK Ltd.

The first Pathfinder project, a proof-of-concept demonstration, took data from raw intelligence reports through to the representation of processed information in a collaborative environment. The project utilised semantic web languages for representation formats, reasoning, and rule engines and developed a rich set of complex semantic representations for the core domain models (ontologies), set of rules, truth-values, assumptions and rationale to support situation awareness and decision-making amongst intelligence analysts. (See: Smart, P. R., "Controlled natural languages and the semantic web" 2008.)

The second Pathfinder project focused on the way in which human agents could interact with machine agents, and the underlying representations in such an environment. This was a pinch-point in such systems where the environment and underlying models are potentially changing rapidly, and in which the analyst must be able to deeply configure and direct the system in order to maximize the benefit for themselves and their collaborators in their generic operational context (Figure 26).



Figure 26. Operational context for the Pathfinder transition projects

This represented the first applied research using CE and it was able to significantly simplify the architecture by replacing the semantic web languages previously used—the *Web Ontology Language* (OWL) and the *Semantic Web Rules* 

*Language* (SWRL). The project was thus able to utilise CE as the sole language for representation and reasoning, as it was both directly machine-processable and human-readable without transformation to alternative forms.

In addition to the focus on the CE language the project also integrated a number of emerging *Web 2.0* practices, particularly exploiting the *actions-of-others*. This was used to demonstrate collective intelligence in the environment, specifically to enable multiple users within the environment to be aware of the relevance of their own actions in the context of everyone else's activities.

The final Pathfinder project focused on the feasibility of achieving information extraction from unstructured textual reports using *natural language processing* (NLP). The novelty in the approach taken was to try to use the CE representations developed in the second Pathfinder project to drive the NLP pipeline. This involved defining models of language, documents, reports and other contextual factors in addition to the traditional domain models for the area of operation, such as IEDs, narcotics production, or influence within tribal societies. The project identified a basic pattern to achieve this, and this was an extension to the field of *ontology-based information extraction* (OBIE). Additionally the purity of the CE language and the approach taken meant that the system had the potential to represent and direct far more of the extraction pipeline than is the case in most OBIE systems.

### Information Fabric and Related Technology Transitions

There have been a number of transition activities over the course of the NIS ITA programme that have had at their core the ITA Information Fabric. These have matured the Information Fabric and added to its capabilities (particularly by incorporating ITA policy research), all brought together into the current open-source Edgware Fabric. These transition activities began in 2007 and the latest was in 2015. They have been funded by a mixture of ARL and CWP in the US, and Dstl and DE&S in the UK.

One particularly important transition in maturing the Fabric and enhancing its capabilities was the *ITA Sensor and Policy Software Tools and Protocols for Networking of Disparate ISR Assets* project jointly funded by ARL, CWP and Dstl (between 2009 and 2013). The project included team members from ARL, Dstl, IBM US and IBM UK; it was supported by the *U.S. Central Command* (CENTCOM).

The goal of the project was to develop a set of sensor and policy software tools and protocols that would enable rapid assembly/dynamic control of ISR assets, and policies that govern the sharing of those assets and the dissemination of data and information to support multiple and dynamically changing coalition missions. The project exploited two key technology components researched and developed within the ITA, namely the Sensor Fabric (the original name for the Information Fabric) and the Policy Management Toolkit [2].

The team participated in Empire Challenge 2010 at Ft. Huachuca (July 2010), an ARL field experiment at Camp Roberts in California (June 2011) and a Dstl field experiment at Porton Down in the UK (February 2012). The Porton Down experiment utilised the Information Fabric to link together a set of heterogeneous sensor assets—provided by a range of different companies, including some small and medium enterprises, as well as more established defence suppliers—so that tips and cues could flow between sensors and feed data to a common operational picture in the experimental Ops Room. The final demonstration at Pershore in the UK (March 2012) linked together the UK *Persistent Wide Area Surveillance* (PWAS) *Capability Concept Demonstration* (CCD) 2—which included a variety of networked sensors—with US sensors utilising a different sensing modality, and policy-controlled information dissemination [3].<sup>7</sup>

A subsequent set of UK transitions were funded by DE&S as part of the UK Land Open Systems Architecture (LOSA) Research, Experimentation and Development (RED) programme. These utilised the Information Fabric to enable data and service sharing between three platform architectures (vehicle, solider and base) within LOSA. These transitions took place between 2012 and 2015, and included several major UK defence companies who were not part of the ITA (as well as IBM UK and Dstl). The transition included the creation of an initial draft Defence Standard (Def Stan 23-013) for the Common Open Interface Land (COIL) in 2014, which was based upon the Information Fabric, and subsequent testing that a number of suppliers could produce interoperable software solutions to the standard. These LOSA transitions were instrumental in leading IBM, with support from Dstl and ARL, to make the Information Fabric available as the open-source software Edgware Fabric.

In parallel with LOSA RED, another Dstl transition utilised the Information Fabric within a proof of concept test of the Air domain *JBRIDGES* concept. Subsequent to this, a further Dstl transition in 2015 developed the concept, benefits and roadmap for realisation of a *Joint Tactical Information Service* that would deliver a tactical service bus providing information as service across the tactical Air, Land

<sup>7</sup> The PWAS CCD was deployed overseas to an operational theatre as an Operational Concept Demonstrator (OCD).

and Sea domains.

### Gaian Database and Related Transitions

Similarly to the Information Fabric, there have been a number of transitions that have had at their core the ITA Gaian Database. These have matured Gaian and added to its capabilities (particularly by incorporating ITA policy research) into the current open-source Gaian Database. These transition activities began in 2009 and the latest was in 2014. They have been funded by a mixture of ARL, CWP, Dstl and DE&S.

The initial set of transition tasks were funded by Dstl and included: (i) evaluation of the feasibility of utilising Gaian on current legacy architecture, with current bandwidth constraints; and (ii) evaluation of the utility of Gaian to federate information (both structured and unstructured).

These transition tasks included the development of the *Visual Query Builder* (VQB). VQB allows analysts to assemble a query by dragging and dropping from a customisable palette of entities (e.g. people, places, dates, products) extracted from structured or unstructured data sources. Once entities have been added they can be linked and have search terms specified to create a query. For example connecting a person called "James" and a place called "London" together and running the query will result in a list of all people called *James* who have been mentioned in the context of a place called *London*. Multiple entities can be joined together to create powerful queries that if created in a language such as SQL would require many complex statements to be written.

The *ITA Policy Controlled Dissemination* was the next major transition project (2011-2013); it was another joint transition funded by Dstl, ARL and CWP, with support from TRANSCOM and NORTHCOM. The resulting technology transitioned to the *NATO Intelligence Fusion Centre* (NIFC). The project comprised team members from ARL, Dstl, IBM US and IBM UK.

The project exploited the Gaian Database, the Visual Query Builder and the Policy Management Toolkit from the joint Information Fabric *ITA Sensor and Policy Software Tools and Protocols for Networking of Disparate ISR Assets* project.

The main purpose of the project was to demonstrate how the policy based access control mechanism within the Gaian Database could be used to provide finegrained access control to information held within the different data repositories. It used the extended Kerberos mechanism to authenticate the users, and to provide single-sign-on to the different data sources. The activity also successfully demonstrated how policies could be changed dynamically, restricting or relaxing policies in response to operational requirements. Achieving this in the accredited single security domain of the NIFC *Battlefield Information Collection and Exploitation System* (BICES) network was a major achievement (Figure 27). However, an obvious extension of this capability was to demonstrate how the model could be extended to data sources that exist in different security domains and indeed to multi-level security domains. In a subsequent transition project, it was demonstrated how the Gaian Database can be used to perform queries across security domains in an accreditable manner.



Figure 27. NIFC deployment of the Gaian Database

The project demonstrated the federation of the different data sources on the NIFC BICES network, without the need for replication and with no central repository of information, acting as a *virtual knowledge base* (VKB). The sources federated across this set of activities included:

- Combined Information Data Network Exchange (CIDNE)
- Coalition Shared Data Server (CSD)
- Tactical Ground Reporting System (TIGR)
- NATO Intelligence Toolbox (NITB)
- Microsoft SharePoint®

In parallel with the NIFC transition, the Gaian Database and the Visual Query Builder became part of a MOTS solution which was procured and deployed to Afghanistan.

### Collaborative Planning Model (CPM)

This transition project (November 2011 to September 2012) was an evaluation of the *Collaborative Planning Model* (CPM), a rich ontology developed in the CE language to support multi-level coalition planning for military operations. The transition project involved Dstl, IBM UK (funded by Dstl) and the *NATO Consultation, Command and Control Agency* (NC3A) organisation in The Hague.

NC3A had designed the *Tool for Operational Planning, Force Activation* and Simulation (TOPFAS) strategic planning tool that supports the NATO *Comprehensive Operations Planning Directive* (COPD). The purpose of the evaluation was to demonstrate how CPM can act as an interoperable plan representation format for plans created in TOPFAS and distributed to other members of the NATO coalition to be imported into their country-specific planning tools and processes.

CPM was developed within the ITA research programme to support developing a shared understanding of a plan, including how the plan details can be communicated and understood across different planning cells and different members of a potentially multinational coalition (Figure 28). In addition, such a planning representation must also be capable of communicating a Commander's intent and its rationale to military planners and operations staff.

The TOPFAS planning tool provides documentation outputs for planners to brief their Commander but there is no capability for the plan to be exported in an electronic form that can be imported into other software tools. IBM worked with NC3A to define an XML Schema that represents the planning objects in TOPFAS. The XML is transformed into CE and through the application of rules developed as part of the project, a representation of the TOPFAS plan expressed as CPM is created. This is then available as a rich representation of the plan that can be transformed into alternative representations.

The project successfully demonstrated the export of plans from TOPFAS and the import into other tools with the exemplar tool being Fujitsu's *openJOP* which was being developed for the MOD.

basic logic and rationale	Agent, Assumption, ConceptualSpace, Container, Entailment, Inconsistency, PossibleWorld, Proposition. PropositionIndex Quantity, ReasoningStep, Set, Triple Verbinding, WorldState
general	ConceptualThing, Constraint, Synchronisation, Context
temporal	Precede, TemporalConstraint, TemporalEnity, TimeInterval, TimeLine, TimePoint
space	Area, Elevation, Line, Point, SpatialConstraint, SpatialCoordinateSystem, SpatialEntity, SpatialIntersection, SpatialLocation, SpatialUnion
resources	Resource, ResourceAllocated, ResourceCapaility, ResourceConstraint, ResourceQuantity, ResourceSet
actions	Activity, Effect, Precondition
collaborative problem solving	ChoicePoint, Collaboration, Commitment, Communication, ConstrainViolated, Decision, GoalSpeification, Influence, Issue, JointPersistantGoal, MutualGoal, Problem, Solution, Trust
planning	Allocation, Evaluation, EvaluationCriterion, InitalState, Plan, PlanTask, PlanTaskDescription, PlanTaskTemplate, PlanningProblem, PlanningProblemContext, ResourceCommitment, ResourceReq, TaskCommitment
military planning	Terrain, Brigade, Division, Field Artillery, RotaryWing, Mission, Intent Area, DecisionPoint, ResourcePool

Collaboration required a common, shared, model of concepts as well as the state of the problem solving process (including rationale)

Figure 28. Multi-layered conceptual models to support collaborative planning in the CPM

#### Management of Information Processing Services (MIPS)

In this Dstl funded transition project (two phases: November 2011 to December 2012, and January 2014 to July 2014) IBM UK applied a number of emerging technologies from the ongoing ITA research programme to provide a service-based infrastructure for information analysts to support them in maximising their success rate in achieving analytical goals. The infrastructure enabled fact extraction and analysis from disparate coalition information sources via rapidly composed information processing services, deployed and integrated using the Information Fabric (later open-sourced as Edgware Fabric). The facts obtained from this process were persisted in a CE Store against which long running CE queries, specified by the analysts, were issued. Analysts were automatically informed when information was available that satisfied their queries, and were supported in identifying relevant source documents/material (when necessary). A key aim was to ensure that analysts were not required to read and assess every

possible source of information, and could instead focus on key material identified by the infrastructure from a large source corpus.

This approach enabled rapid, agile composition of processing services and information sources in response to both long- and short-term (sometimes fleeting) analytical requirements. A visual user interface enabled specification of the information processing functionality required (information extraction) and the expression of analytical queries. The project was also able to bring in a number of capabilities from the earlier Pathfinder series (most notably some of the information extraction services from the third Pathfinder transition, and the use of the CE language as a representation format for many elements of the MIPS environment). The project was successfully developed and transitioned to Dstl; Dstl then used the environment for various ongoing experimental activities with international partners [4].

### Actionable Intelligence Technology Enabled Capability Demonstration (AI-TECD)

ARL and the U.S. Communication-Electronics Research, Development and Engineering Center (CERDEC) Intelligence and Information Warfare Directorate (I2WD) began a joint effort in 2014 to develop and transition ITA technologies for the Actionable Intelligence Technology Enabled Capability Demonstration (AI-TECD). The AI-TECD is focused on technologies for small units to improve capabilities for sending and receiving critical tactical intelligence and prevent surprise encounters without increased physical or cognitive burden on the soldier.

The project exploited ITA research into *Sensor Assignment for Missions* (SAM) including the OWL ontology and mixed-reasoning matching and bundling algorithms. This is done to match mission requirements to sensor capabilities from a collection of disparate *Open Standards for Unattended Sensors* (OSUS) enabled ISR assets and provide users with sensing options at the tactical edge [5, 6].

The technology was used to assist edge users perform mission planning and matching with OSUS sensor assets at the final AI-TECD field trial at Fort Dix, New Jersey in July 2015.

In the second part of the AI-TECD transition effort, the CE Store were used to develop *ontology on the fly* capability to augment the sensor ontology (in OWL). This has the potential to provide a much-needed agile capability for the edge user to customize the domain ontology for his local missions without being an

ontology expert.

#### Coalition ISR Assets Interoperability (CIAI)

Technologies developed within the ITA program were exploited by ARL for the CWP project on *Coalition ISR Assets Interoperability* (CIAI) between ARL and *Defence Research Development Canada* (DRDC) in Valcartier, Quebec that was initiated in October 2015. The objectives of the CIAI project were to: (i) develop open agile architecture for *plug-n-play* interoperability and ISR interoperability standard(s), (ii) develop mission-driven resource management and mission programming for command and control of ISR assets, (iii) develop policy-controlled C2 and information dissemination, and (iv) demonstrate, in an operational environment, the viability of the ISR interoperability operational concept, validity interoperability architecture, and enhanced situational understanding.





For the second objective, CIAI built on the SAM algorithms and tools developed by the ITA team from University of Aberdeen, Cardiff University, ARL and IBM US. It also be exploited SAM's *Military Mission and Means Framework* (MMF) originally developed by ARL to enable the optimal matching from available ISR information resources (means) to the mission-relevant information needed (mission) for enhanced ISR applications (Figure 29).

For the third objective, CIAI will leveraged WPML tools developed by IBM US that were exploited previously by the two CWP projects with the UK.

### NATO Summit 2014: Understanding the Local Public Reaction

The following is extracted from a press release regarding this particular transition project which was self-funded by Cardiff University and IBM UK between July and September 2014.<sup>8</sup> It was undertaken to show the potential value of using a rich human/machine collaborative environment with conversational capabilities applied to social media data for large events or protests in order to understand local public reaction to the event and the policing of the event:

One week on from the NATO Summit, researchers at Cardiff University studying community reaction to the NATO Summit have found that despite general negativity on social media towards the event, there were some significantly positive local reactions.

The team of Social and Computer Scientists conducted an experiment using a combination of computing and social science methods to gauge local community reaction to the NATO Summit.

This experiment – part of on-going work in collaboration between Cardiff University School of Computer Science and Informatics and the Universities' Police Science Institute (UPSI) – examined "mass" reactions to events on social media whilst also homing in on specific incidents as they were unfolding on the ground.

The software and methods used for the study were able to detect specific incidents and potentially disruptive events - including spontaneous unscheduled protests - as they were happening.

"As these potential incidents were being identified, we were able to direct our teams on the ground to obtain accurate information,

8 For the full press release see http://www.cs.cf.ac.uk/newsandevents/natostudy.html

such as the size of crowds, their mood, where they were going, and so on," explained Professor Alun Preece from the Cardiff School of Computer Science & Informatics.

"Working with our research colleagues at IBM Emerging Technologies, we were able to capture information in a knowledge base that can be employed to answer questions, to help make decisions, and to manage an ongoing situation like the Summit in real-time."

The success of the demonstration led to the founding of the *Open Source Communications Analysis Research Centre* (OSCAR) at Cardiff University which has attracted funding from UK Police organisations [7].

#### Hudson and Watson at Wimbledon 2015

*Hudson* is a self-funded (by IBM) transition of ITA CE Store and CE Node. It was used at Wimbledon in 2015 by IBM together with the *IBM Watson Engagement Advisor* (WEA) to answer questions about Wimbledon posed using a natural language interface.

The questions posed can generally be split in to two categories: (i) language comprehension, such as "why do the players wear white?"; and (ii) statistical questions, such as "what is Andy Murray's average serve speed in semi-finals on grass?". Comprehension questions would be answered from information in unstructured written documents, whereas the statistical ones would best be answered through Wimbledon's relational databases.

WEA, with sufficient documents and training, answered the comprehension questions. It did not, however, answer structured statistical queries. Instead Hudson used the CE store to enable the exploitation of Wimbledon's relational databases to answer the statistical queries. The CE Store was also used to provide a chat like interface for people to interact with WEA and Hudson.

# **Final Note**

There are many other transitions whose stories could have been told, but space and sensitivity prevent the listing and explanation of all the transitions. Instead this chapter has focused on telling the stories of several of the most important (from a defence point of view) transitions, highlighting the open-source software available for wider exploitation, the speed of exploitation (into products and services used in the real world) and noting the breadth of ITA science that has been exploited so far.

It is clear how the deep collaboration between the members of the Alliance, and particularly between the UK and US, has supported effective interaction between transitions which have been able to build on prior transitions and thus achieve results that neither nation could have achieved alone.

# References

- T. Killion, P. Sutton, M. Frame and P. Gendason, "A New Paradigm in International Collaboration: The US-UK International Technology Alliance in Network and Information Sciences", in *RUSI Defence Systems*, pp. 46-49, June 2007,
- [2] T. Pham and G. Cirincione, "Sensor, Data and Information Sharing for Coalition Operations", in *Knowledge Systems for Coalition Operation* (KSCO), January 2012.
- [3] T. Pham, R. Young, G. Pearson, F. Bergamaschi and D. Conway-Jones, "Networking & Interoperability of Disparate Coalition Assets – A Coalition Warfare Program (CWP) Project", *NATO SET-186/IST-112 Symposium*, Quebec, Canada, May 2012.
- [4] D. Braines, J. B. Ibbotson, and G. White, "MIPS: A service-based aid for intelligence analysis", in Proc. SPIE, vol. 8758 875809-13.
- [5] T. Pham, G. de Mel, J. Schoening, and R. Ganger, "Sensor and Information Fusion for Actionable Intelligence at the Tactical Edge", *IST-*126/SET-189 Symposium, Norfolk, VA, May 2015.
- [6] J. Schoening, T. Pham, et al, "PED fusion via enterprise ontology", in Proc. SPIE 9464, Ground/Air Multisensor Interoperability, Integration, and Networking for Persistent ISR VI, 94640D (20 May 2015)
- [7] A. Preece, W. Webberley, D. Braines, "Tasking the Tweeters Obtaining Actionable Information from Human Sensors", in *SPIE Defense, Security,* and Sensing, International Society for Optics and Photonics, May 2015.

# 7 External Perspective

*R. Srikant (US lead peer reviewer), Mike Underhill (UK lead peer reviewer)* 

This chapter has been written by the peer review team as part of their final peer review of the NIS ITA programme at the Annual Fall Meeting of the ITA (AFMITA) at the University of Maryland in September 2015. The remaining content of this chapter is their material and expresses in their own words a number of perspectives on the ITA programme throughout the ten-year period.

# The NIS ITA

The Network and Information Science International Technology Alliance (ITA) is a jointly established research alliance formed from US and UK government, industrial and academic members for the purpose of conducting research to develop underpinning technology applicable to network-centric warfare and to enhance US and UK capability to conduct coalition warfare. It is a bilateral UK Ministry of Defence (MOD)–U.S. Army collaboration and an integrated US/UK industrial academic consortium. It was originally a 5-year programme started in May 2006 and in 2011 was extended to 2016. It builds on the success of UK Defence Technology Centres and US Collaborative Technology Alliances.

The Programme Objectives and Assessment Criteria are as follows:

- Science and Technology Achievement. Shown by the technical quality and significance of the advances made under the programme.
- Military Value to Stakeholders. Shown by the value to the US and UK stakeholders of the advances made under the programme both now and in

the future.

 US/UK Collaboration. Shown by the extent to which the programme has delivered greater value to both governments through effective and true collaboration across international boundaries, organizational boundaries, and technical area boundaries.

# Summary of the 2015 Annual Fall Meeting

The peer reviewers were invited to the 2015 Annual Fall Meeting (AFM), which took place on 15–17 September 2015 in the College Park Marriott at the University of Maryland. As this was the last AFM of the ITA Network Information Science 10-year programme, the reviewers' remit was to focus on eliciting highlights from the entire 10-year ITA programme, interact with the researchers to encourage successful completion of the programme in 2016, and identify further the Critical Success Factors (CSFs) of this exemplar and unique US/UK research collaboration.

The AFM schedule covered 2 days and was followed by a Capstone planning meeting, to which the peer reviewers were invited to participate. On Tuesday, 15 September 2015, after introductions from Seraphin Calo, George Vongas, and Greg Cirincione, and invited talks from Jeff Singleton and Sue Toth, there were presentations of the Technical Area 5 (TA5) programme by Don Towsley and Technical Area 6 (TA6) by Alun Preece. In the afternoon, there were 12 demonstrations and 35 short papers presented as posters. The Capstone Planning Meeting followed on Thursday, 17 September.

Because this was an interim review prior to the termination of the ITA programme in 2016, the peer review panel was not provided with review documentation prior to this AFM.

The AFM itself was found to be well planned with a single stream of lectures and demonstrations with short paper posters. This was ideal for the peer reviewers and seemed to be very good for the researchers and other participants. A single stream encourages cross-fertilization between the Technical Areas. Tutorials, Focus Groups, and Technical Area Planning meetings should perhaps always be kept to a separate extra day. In the long paper session, the clear designation of the task number and technical area for each long paper was welcomed. It would have been helpful if the short paper posters had been annotated in the same way. Most of the long paper presentations included very useful introductory slide(s) placing

the presented work in its historical context in and relevance to the programme. Particularly for this reason, it would have been useful for the reviewers to have had copies made available to them before or immediately after the sessions.

The peer review panel comprised three members from the US and four from the UK. All three US panellists were from academia. Two of the UK panellists were from academia and two were independent with strong historical connections to academia or the MOD. The skills of the panellists were wide ranging, covering all the areas being reviewed.

# Some Reflections on Peer Review 2008-2016

The first peer review took place in 2008, which was 2 years after the start of the ITA programme in 2006, with further full reviews in 2010, 2012, and 2014. There were also informal reviews in 2009, 2011, and 2013. In February 2011, the reviewers assisted with the shaping of the second 5-year follow-on phase of the ITA.

Importantly, there was a rolling 2-year plan throughout the programme, the Biennial Planning Process (BPP) run by the ITA consortium, at which the Peer Review Reports, both formal and informal, provided a significant part of the input. The peer reviewers have been very pleased by the way their evaluations and suggestions were taken into account in the BPPs.

The peer review process was established at the first formal peer review in 2008 and in subsequent years, small but useful refinements were made, keeping the basic process essentially unchanged.

For the first phase of the ITA programme 2006–2011, the work was divided into four Technical Areas, TA1 to TA4. For this reason, a total of eight independent peer reviewers were chosen, four from the US and four from the UK. Particularly for the formal reviews, there were also independent government peer reviewers, one from the US and one from the UK.

However in the 2008 review, it became apparent that the four TAs were operating essentially as silos/stovepipes with little, if any, interdisciplinary work between them. Throughout the rest of Phase 1 of the ITA, until 2011, the peer reviewers urged better integration of and collaboration between the technical areas, and were very pleased that several of their suggestions on how this might be improved were taken up. Phase 2 of the ITA programme (2011 to 2016) was reshaped by combining TA1 and TA2 to become TA5, and TA3 and TA4 to become TA6.

Taking them out of their silos/stovepipes enabled the level of interdisciplinary work even including between these new combined technical areas (TA5 and TA6) to achieve an appropriate and commendable level. As a result, toward the end of the programme, although dealing with separate, well-defined domains, the reviewers felt unable to compare the performance of TA5 and TA6 independently, as they had become so well integrated.

The formal reviews included detailed "traffic light" assessments by the peer reviewers based on comprehensive Peer Review Reports and presentations made as part of the AFMs. The usual format was presentations from each of the Technical Area leaders and full attendance of the paper and poster sessions delivered by the researchers. If there were multiple lecture or tutorial streams, the reviewers split their attendance according to expertise to ensure full coverage. In the latter stages of the programme, "demonstrations" featured increasingly heavily, at the original request of the reviewers. The demonstrations were seen to be very important for showing the utility of the research work, collaboration between tasks and Technical Areas, and any "potential or actual transition opportunities."

The traffic light assessment methodology adopted from the start was applied to five (four for the US) peer review assessment criteria: scientific quality, relevance, collaboration, potential for technology transition, and technical risk/challenge. These criteria were assessed for every project in all the Technical Areas against a four-level traffic-light assessment scale for each of these criteria with the ratings being (red = stop, amber = redirect the work, green = good, and blue = outstanding).

Peer reviewer consensus on these was always achieved, sometimes after considerable debate, sometimes by looking at the projects down to separate "task" level, sometimes by giving the relevant domain expert in the peer review panel the "casting vote," and sometimes allowing "abstention" when a reviewer did not feel sufficiently expert in the area to be able to venture an informed opinion for consensus.

# **Critical Success Factors**

The overall success of the NIS ITA programme and the listing of benefits, expected and unexpected, has raised the question of what are the CSFs, seen and unforeseen, that have made this programme successful. The following list of CSFs is a reconsideration and revision of a list in the 2014 Peer Review Report. That report also contained an initial list of "lessons learned" as seen by the peer

reviewers. What is notable is the continuous improvement throughout the ITA programme due to these CSFs and the way in which they addressed lessons learned.

This list of CSFs is supporting evidence for the thesis that this ITA programme is "an exemplar for any future collaborative research programme". Arguably, they should be incorporated in any future such programme.

- Any Industry Prime contractor for a bilateral multi-party research programme should have well-linked research facilities both in the US and UK.
- Any future programme should be flexible enough to allow continuous improvement, termination of blind alleys with redirection to promising areas, and the pursuit and exploitation of new discoveries. It is very important to recognize that good research is unpredictable at the outset. (This has been the key to the success of the NIS ITA programme.) The 2-year rolling refresh of plan, the BPP, has proved to be a very good method of achieving this flexibility, maximizing progress but at the same time retaining the necessary balance of stability required for truly "blueskies" future-looking research.
- An open-minded research culture directed toward addressing the generic problems of the overall programme is essential. Identifying and defining "hard problems" to be addressed can provide the necessary challenges to the researchers.
- As far as possible, there should be multi-way collaboration among all participants in every project. There were six categories that had to be well linked: DoD/MOD laboratories, US/UK academe, and US/UK industry. (This makes 15 possible two-way collaborative linkages!)
- There should also be multidisciplinary cross-sector teams and collaboration between any scientific, engineering, or operational specialties within the entire programme being undertaken.
- An important CSF is "critical mass" for all teams down to the task level. The ITA achieved this very successfully.
- The peer reviewers believe that they have been able to contribute significant value to the ITA programme and therefore their existence was an important CSF. Thus, they strongly recommend that there should be independent peer

reviewers at the outset and throughout any such programme.

- The five (four for the US) peer review assessment criteria of scientific quality, relevance, collaboration, potential technology transition, and technical risk/challenge have proved to be an inspired choice and are highly recommended for any collaborative research programmes. Also recommended is four-level "traffic-light" assessment scale for each of these criteria: red = stop, amber = redirect the work, green = good, and blue = outstanding. Peer reviewer consensus on these is important, but sometimes this may mean giving the relevant domain expert the "casting vote."
- It is important that there should be a peer review culture of constructive criticism and suggestions, wherever possible including face-to-face discussion with researchers. (Peer reviewers are by definition chosen for expertise in relevant areas.) Peer review should never be just a detached one-way arms-length assessment. Constructive feedback to the researchers is essential and integral to the programme.
- The availability of common scenarios for the testing, integration, and assessment of the research is highly recommended. It automatically informs the researchers of the potential value of their research so that it can be better directed.
- The availability of a common Experimentation Facility is also significant asset. This also acts to integrate the programme and promote dialogue and collaboration between researchers in the tasks and technical areas.
- An annual forum/meeting for informal or formal peer review, for cross-project interaction of researchers, the initiation of new technical collaborations, and as a spur to the achievement and reassessment of research milestones, has proved to be essential for the ITA programme and is highly recommended for any bilateral US/UK research programme.
- For a collaborative, multi-party, international, geographically dispersed research programme, such as the ITA, regular face-to-face meetings have proved to be essential. The annual "boot camps" have been a most productive facilitation of this.
- The reviewers are very pleased to hear of the plans for creating an accessible archive of all the technical papers and relevant papers from the entire ITA programme from 2006 to 2016. Such a "legacy" archive system

is important and highly recommended for future research programmes of this kind. The reviewers are also interested that a book suggesting lessons learned in the running of such an international partnership programme is being developed; this would be a valuable product. It is also a part of putting on permanent accessible record the technical and scientific "lessons learned." Both of these should prove to be additional CSFs in the future.

# Concluding Remarks in September 2015

There is no doubt that the NIS ITA programme has been a great success. This is what the reviewers have consistently identified, as can be consistently seen from reviewers' comments. We have given substantial opinion and evidence in our successive reports and these are on open record. The final review this year confirms the upward trend.

Without doubt, the programme can now be regarded as an exemplar for undertaking any future collaborative research programmes. It is important that the "lessons learned" and the CSFs should be captured and recorded before the programme terminates.

Once again, the peer reviewers found the experience enjoyable and instructive. They continue to be gratified by the way that many of their recommendations and constructive suggestions have been taken into the programme over the years.
## 8 Achieving an Outstanding International Collaborative Research Alliance

An important objective, and in some sense the most important objective of the U.S. and UK Government vision for the ITA was the need for a true, deep, and persistent enduring collaboration between all of the Alliance partners from the different sectors (government, industry and academia). A multi-disciplinary as well as cross-sector approach was foreseen as necessary to generate world leading innovative and disruptive scientific advances in Network and Information science [1]. The ITA independent peer review panel identified that the ITA is "an outstanding example of true, deep and enduring International Research Collaboration" that has "significantly advanced the state-of-the-art in network and information science through multi-disciplinary research".

Indeed the ITA programme was a unique international endeavour that brought together a large number of US and UK research institutions from academia, industry and government to work together on a collaborative research agenda for a decade. Given the size, complexity and diversity in technical focus and cultural differences between different organizations, there were significant challenges in realizing the aspirations for the programme.

There are several important lessons related to execution of cross-sector multidisciplinary international collaborative research programmes that we learned during execution of the ITA. In this chapter, we summarize these lessons for the benefit of others who may need to lead similar programmes in the future.

These lessons exploit and combine the insights from peer reviewers evaluating the programme ("External Perspective" on page 119), the personal experiences of the programme leadership, ideas resulting from interactions with other peer collaborative programmes in the US and UK, and the experiences of government researchers who have engaged in other multi-organization collaborative single-country programmes.

There are many challenges associated with the execution of a complex programme such as the ITA that seeks to achieve a *"true, deep and enduring collaboration"* and to build a collaborative culture where the collaborative behaviours are incentivized. This collaborative environment required energy, leadership, and persistence to achieve. We begin this chapter with a discussion of some of those challenges, and then enumerate some of the different initiatives that proved to be critical factors for the success of the programme.

## Challenges of International Collaborative Alliances

The main challenge in the ITA was to find approaches to achieving the common goal of advancing the science of coalitions through deep and persistent multidisciplinary research collaboration, while understanding the different constituents in the Alliance and their organizational, cultural and, technical differences. These differences caused a number of tension points, each with their own trade-offs, which needed to be harmonized.

In this section, we outline the main salient tensions that arose in the ITA and describe some of the factors and approaches that we have found to successfully address them.

#### Tension between Blue Skies and Targeted Research.

The overall rationale and premise for the ITA created an inherent tension between blue-skies and targeted research since it had dual mandates to both provide the fundamental underpinnings for network and information science, and to exploit those advances in the state-of-the-art to enhance coalition operations. Advancing fundamental science can be blue-skies research which is driven purely by a desire to increase human knowledge in a specific domain, and may or may not have relevance or bearing with any practical applications of the knowledge that is discovered. On the other hand, targeted research is the type of fundamental science that is done with a specific goal and knowledge discovered in that exploration can have an immediate application, e.g. it can be useful to the needs of a military coalition operation currently in the field.

Both blue-skies research and targeted research have their advantages. A blueskies exploration is high risk, but can have a potentially high pay-off in the future, although it may be hard to articulate its impact in a precise manner. Targeted research has the advantage that its benefits and impact on the objectives are much more direct and clear. However, targeted research is more incremental and longer-term pay-offs are likely to be much more limited as evidenced in studies of different approaches to medical research funding.

Stakeholders, researchers, and managers in the programme had varying preferences, interests, and perspectives that needed to be balanced to meet the overall goals of the programme.

#### Tension between Scientific Freedom and Accountability

Researchers and scientists in the ITA would love to have complete freedom to perform the research that they believe is most appropriate to solve the challenges that they identify, and that broadly fits within the scope of the ITA research programme. The nature of fundamental research is that its outcomes are unpredictable, insights often unexpected, and innovation paths are diverse. On the one hand a laissez-faire and unconstrained management approach would seem to enhance scientific freedom and give scientists the flexibility and freedom to innovate, as well as sufficient laxity in the time period, so that they have a good probability of making progress on the scientific challenges. On the other hand research progress must be assessed, and focus on the agreed upon scientific challenges maintained, to achieve the goals of the programme.

Technical review was critical to the success of the ITA, but must be done without stifling the freedom of the scientists, and without becoming too onerous. Within the ITA programme, we had to experiment with several models for accessing and reporting scientific progress before we ended up with an approach that had the appropriate balance between accountability and scientific freedom. Another important issue was ensuring that the research was being done on the right scientific challenges. As scientific advances are made within the global research community, some of the problems that appear challenging at first may no longer remain relevant. Progress in related areas made within or without the research programme may identify new scientific challenges worthy of investigation and, an associated re-prioritization of the current research agenda. In this context, it is worth reflecting on how much the wider world has changed over the ten years of the programme (e.g. rise of smartphones and on-line shopping). Lastly a programme of this size cannot possibly address all of the technology needs of future coalition operations, so choosing to address the most critical scientific challenges and gaps is critical.

The right balance between providing researchers sufficient freedom to pursue their

<sup>8</sup> Achieving an Outstanding International Collaborative Research Alliance 129

exploration without overburdening them with excess oversight, and appropriate revision and focus of the research agenda were key to the ITA success.

#### Tension between Group Needs and Individuality

Each organization is formed for its own individual needs. Industrial research labs are generally created by commercial entities to create technologies that will lead eventually to a profitable outcome. Academic research labs are generally created by universities with the goal of attracting the best student talent, as part of their mission to improve the training and the knowledge created by them. Government research labs are created so that they can increase the technical capabilities and effectiveness of the branches of government that they support. These organizations can achieve their goals individually, but can achieve significantly more if they harmonize their approaches toward a common goal and if their diverse strengths are exploited.

Individual researchers have a further sharper focus on their individual goals and interests. Each researcher has their own personal interests and incentives. These interests are driven by a complex set of factors such as the past experiences of the researcher, the culture prevailing in the employing organization of the individual, the technical perspective of the researcher, and their previous history in collaborating or investigating research problems. It is to be expected that the values and motivations of each organization and individuals will be different.

The Alliance itself was formed with a specific set of goals and objectives that do not exactly align with the needs of any individual Alliance member. An example of such a tension comes from the needs of individual researchers, many of whom are well-known as leading experts in their field. While such researchers may be willing to collaborate with other members of the Alliance, they may not have enough time available to do so. It is also quite possible that a strong researcher may make more progress if they were working alone, whereas collaboration will slow them down since they would need to mentor members lacking sufficient expertise in the area. In this case, the needs of the individual researcher is partially overlapping with the needs of the Alliance (to do great science) but also diverges partially (surge ahead alone versus collaboration). This tension is a natural outcome.

Managing the tension between individual needs and aligning them with the overall Alliance needs was an important factor in the success of the programme.

#### Tension between Depth and Breadth of Research

An international collaborative research programme, especially one whose scope covers multiple technical areas and scientific disciplines, faces a further challenge in managing the depth of research work undertaken in each technical area and discipline versus exploiting the synergies across different technical areas and disciplines to come up with new insights. Both of these approaches have their own merits, one results in deeper exploration of a subject, while the other can lead to new insights and results that transcend any individual area or discipline.

A similar tension exists in the size of the Alliance, measured in the number of participating principal investigators. In general, a larger alliance is more likely to provide a larger pool of experts and more opportunities for synergistic collaborative research to emerge. On the other hand, there is a danger of being spread too thin, not having sufficient resources to collaborate, and losing focus if the alliance becomes very large.

A similar tension exists in determining the set of experts that are chosen in each technical area. With the explosion of scientific knowledge in the modern age, every technical area is fairly broad. One can pick experts to form the Alliance that are all working on closely related topics. That makes the task of collaboration fairly easy. However, this depth comes at the cost of the breadth of new ideas and new collaborative activities that could be launched if the experts came from different areas.

Important success factors for ITA were to (i) motivate researchers to collaborate across disciplines and technical areas and to avoid their natural tendency to gravitate towards doing deep research in a specific technical area; and (ii) ensure critical mass of individual researchers, disciplines, and organizations.

## Success Factors

In this section, we set out some of the ways that we addressed these challenges based on the lessons learned from the NIS ITA—both those that worked, and those that failed to yield the desired results.

# Addressing the Tension between Blue-Skies and Targeted Research

Within the ITA programme, there was a constant tension between the desire to permit blue-skies research in different areas without too many restrictions, and the opposite desire to have research work that would address the challenges of coalition operations in a more direct manner. In order to resolve these two opposing requirements, the following approaches were used, and each of these approaches helped in making the programme successful.

#### Stakeholder and User Engagement

An important aspect of the ITA programme was continuous engagement with the stakeholders of the programme-including potential end users of the technologies. enabled particularly by the government partners-to ensure that the programme was aligned to their needs and addressing challenges that they considered vital for their operations. This process included raising awareness, consulting with and getting buy-in from end users and others on the technical results that were achieved from the programme. Stakeholders and specialists were invited to meetings of the Alliance to provide insights into their operational needs and experiences. The sharing of this experience helped the researchers to understand the context and needs of coalition operations in personal ways that impacted how researchers viewed their research challenges. Another way to engage with ITA stakeholders was through special interest group days in the UK where transition issues related to the programme were discussed (primarily with government and industry stakeholders). This stakeholder engagement created an environment in which researchers were thinking about stakeholder context, thereby enabling blue-skies research problems to be formulated within that context, thus increasing the probability of impact.

#### Hard Problem Definitions and Scenario Development

In order to obtain the right balance between blue-skies research and targeted research, and to focus the ITA resources on the most critical problems, an analysis of the key technical challenges and scientific gaps for coalition operations was conducted. This produced a set of high-level hard problems that could be addressed by the ITA, and that were used when formulating the research programme. These hard problems were collaboratively developed and adapted by the Alliance leadership and researchers.

To place these hard technical challenges in context and to understand the assumptions and limitations of the coalition environment, the ITA undertook an effort to define relevant military coalition scenarios. By appropriate definition of the scenarios, and defining vignettes within the scenarios that required different types of scientific understanding and technologies, the problems that were required to be addressed were identified. The availability of common scenarios

for the logical experimentation, integration, and assessment of the research helped inform researchers of the potential value of their research so that they can selfsynchronise and place their work in context.

The definition of scenarios was an important step in identifying potential relevance of blue-skies research to the goals of the programme. As an example, one of the blue-skies research activities undertaken in the programme was the development of the theory of network tomography—approaches that could identify internal attributes of a network by only examining end-to-end measurements. At first glance, this may not appear relevant to the context of a coalition. However, as the science was developing, the scenarios made it clear that tomography would be extremely useful in monitoring coalition networks were visibility into partner networks is limited and later to be extended to the placement of distributed processing elements. The identification of hard problems and scenarios helped highlight the value of a fundamental break-through idea.

This created an open-minded research culture, which was directed toward addressing the generic problems of the overall programme. This culture was an essential ingredient contributing towards the success of the programme.

#### **Common Experimentation**

While basic science is the main focus of most research programmes, and validation of research results is always a key part of the scientific process, experimentation across multiple disciplines is difficult. It is also relatively hard for stakeholders and users to understand how that science may transform to have a meaningful impact in an applied context. The output of research programmes includes items such as new mathematical models and new algorithms that are hard to relate to practical systems, and it can be hard to understand how they interact (including compatibility and synergy). However, identifying a set of experiments that position the work in basic science and its practical applications can often be useful. One can run experiments that show which parameters or adjustments to a scientific model yield an improved performance in the context related to real world applications, or one can conduct experiments to understand how a newly invented algorithm improves the performance over other algorithms in a specific applied context.

Experiments are particularly important to verify the science. They can also be used to discover interaction effects between different approaches, and as a bridge between basic science and transition. Within the NIS ITA, we developed an experimentation framework and facility that allowed scientists to run experiments,

<sup>8</sup> Achieving an Outstanding International Collaborative Research Alliance 133

as well as store current experiments and rerun them with new parameters or models in the future. This facility was also useful in creating demonstrations of new technology (a demonstration just being an experiment in which all variables are pre-set).

The availability of a common experimentation facility was a useful utility that helped us improve our basic science explorations, particularly by enabling experimentation on complex interaction effects (which the U.S. National Research Council Committee on Network Science for Future Army Applications identified as a key research need [2]) and also in improving the probability of transition.

# Managing the Tension between Scientific Freedom and Accountability

Finding the right balance between scientific freedom and sufficient accountability in the research programme was one of the tough problems, and the balance between them was learned over the course of the ITA programme. The following measures were the factors that led to what we believe is the right balance between scientific freedom and oversight.

#### Competitive Research Plan Formulation to Adapt to Scientific Advances and Changing Coalition Needs

Any research programme should be flexible enough to allow continuous improvement, termination of blind alleys with redirection to promising areas, and the pursuit and exploitation of new discoveries. It is important to recognize that good research is unpredictable at the outset and that failure often advances the state of knowledge. The ITA utilized a 2-year rolling research formulation process to develop a Biennial Program Plan (BPP). It has proved to be a very good method of achieving this flexibility, maximizing progress but at the same time retaining the necessary balance of stability required for truly blue-skies future-looking research and for Ph.D. students.

The programme plan definition was a lightweight process to minimize the overhead and effort in project proposal and formulation, and also an open process to reinforce the collaborative culture. Each programme plan definition was preceded by a short document collaboratively developed by the Alliance technical leadership team that identified the technical challenges and high-level guidance for scoping the collaborative project proposals. This open and competitive process focused on encouraging innovative research ideas, collaboration among researchers, and multidisciplinary approaches. It proved to be a good way to build

trust and collaboration between researchers and to adapt the research programme based on the state of the science, user feedback, and peer-review feedback.

The peer review was also synchronized with the 2-year BPP, so that a full peer review was conducted every two-years with a lighter weight review in the other years. A full peer review involved a detailed assessment of the research aligned to the peer-review criterion. Presentations, posters, and individual interactions between the peer reviewers and the researchers occurred at the annual fall meeting.

#### Independent Peer Review

One of the important components of the ITA programme was the inclusion of independent external peer review that provided advice to the US and UK Collaborative Alliance Managers who were responsible for the overall Alliance programme.

The programme benefited greatly from rigorous peer-review by a senior team of external academic, industry and government experts from both countries, with deep experience in relevant fields. The peer reviewers frequently gave constructive advice to the Alliance leadership for changes and refocusing of effort that enhanced the overall programme.

The peer-review panel provided independent assessment and evaluation of each project with respect to: (i) the technical merit of the approaches adopted; (ii) the operational relevance of the problems addressed; (iii) the synergistic value of collaboration; (iv) the likelihood of exploitation (technology transition); and (v) whether the science is innovative and advances the state of the art.

#### Lightweight Reporting Process

As the ITA was a novel international collaboration, it took a while to develop an appropriate reporting process that met US and UK governance requirements. The process was initially too burdensome, but by addressing this as an Alliance we were able to develop an effective yet lightweight replacement. This collected key metrics on a quarterly basis with a simple roll-up process from individual projects to a programme-wide scale. We also simplified the technical reports by making it a collection of the outputs produced by the scientific research.

#### Managing the Tension between Individuality and Group Needs

The balancing of the tension between the needs of the individual organization and that of the research Alliance was critical to meeting the ITA goals of deep collaboration to advance the state-of-the-art in network and information science.

#### Consortium Formation: the Articles of Collaboration

The ITA was a novel international research collaboration. As such it was important to define and adopt an effective, yet efficient process for managing the collaboration among consortium members. This included agreement on the administration of the consortium, e.g. determining how membership could be changed, ownership of the intellectual property created by the consortium, and procedures to be followed for fiscal administration etc. This was attained through the creation and agreement of a set of *articles of collaboration*, which were agreed during the consortium's formation.

#### Role of the Consortium Managers

The success of a programme like the ITA depends on the innovation and collaboration of individual researchers but also on the persistence of leadership. The Consortium Managers are critical to this success. Their role is not just to manage the day-to-day operations and mechanisms of the complex international endeavour, but also to provide scientific leadership. These leaders needed to have the scientific stature and leadership skills to: (i) attract, retain, and bring the best out of top researchers in both countries, (ii) encourage researchers to stretch into new areas of investigation, (iii) encourage and enhance collaboration especially with disciplines that do not typically work together, and (iv) create an environment where healthy competition breeds innovation.

A critical issue is finding ways to keep the top researchers, those with plenty of funding, students, and research problems, engaged. The Consortium Managers were critical in providing leadership to make the case to these researchers that the downsides of the processes, collaboration requirements, and competition were far outweighed by the collaborative benefits, the unique scientific challenges, and the impact that their research could have on future coalition operations.

#### Staff Rotation and Exchange

A key element of developing deep, persistent collaboration is the face-to-face interactions that occur during staff rotations and exchange. Collaborative benefits were realized from a variety of exchanges to ARL and among consortium organizations. Another useful exchange involved cadets from the U.S. Military Academy at West Point collaborating with researchers at IBM UK's Hursley Laboratory.

Early in the programme IBM looked for ways to mitigate the time demands on some of the ITA's most prominent researchers. Despite good intentions, it was difficult for them to spend enough time on active collaboration. In order to create a more collaborative environment, the summer internship programme was instituted. Through summer internships at industrial research locations, the graduate students working with various busy professors were brought together into a single location. This enabled creation of long lasting friendships and research collaboration that lasted beyond the graduation of the students and their leaving the Alliance formally.

A wide-ranging array of new collaboration resulted from face-to-face interaction, including peer-to-peer senior researcher collaboration, mentorship of students or of junior government staff, and partnerships of many kinds that may not have been possible without this unique venture. Staff rotation, exchange, and the intern programme helped the ITA to enhance collaboration, especially those involving multinational teams.

#### A Boot Camp for Collaboration

One of the most successful elements of the ITA strategy to produce fundamental advances in network and information sciences to enhance decision making for coalition operations, was the annual Boot Camp. This face-to-face meeting rotated annually between the US and the UK and has significantly enhanced collaboration. The goals of this meeting were to provide an unstructured environment where collaboration, innovation, and collegial interaction flourished. The meetings deepened existing collaboration, started new collaboration especially among researchers with disparate backgrounds, broke down barriers between scientific disciplines, allowed deep and unconstrained interactions with military users, identified new research approaches, and energized the research team.

The meeting was structured in a way that encouraged the challenging of ideas but was not a technical review of the programme. It encouraged the raising of new ideas even if the problems were not fully formed, so as to provide a collegial sounding board. Sessions encouraged dialogue, not just presentation of ideas. Topics were proposed from across the Alliance, and the topics were selected based on the interest shown by other researchers. Agenda slots were held open so that topics that arose during the Boot Camp would have space and time for discussion of the newly identified issues. Cross-research task sessions were held to identify linkages and to investigate if the joint study of aspects of their research could advance the science. At the end of the Boot Camp a short informal session was held to report new ideas and planned follow-on activities in plenary.

The annual Boot Camp has been described by many participants as the most rewarding activity in the ITA. It has led to much collaboration, research projects, and personal friendships that would not have been possible with more structured meetings.

#### Managing the Tension Between Depth and Breadth of Research

Managing the tension between research that goes deep into one specific scientific discipline, an activity that is more likely to produce impactful results in a specific area, and across multi-disciplinary research, which can provide new fundamental insights, is an important challenge in any multi-disciplinary programme. In any collaborative research programme, there should be multidisciplinary cross-sector teams and collaboration between the scientific, engineering, or operational specialties within the programme should be pursued.

#### Ensuring Critical Mass

In any task that is undertaken, a critical mass of research funding and resources needs to be provided. This applies for multi-disciplinary research as well as for the research going deep into any one specific technical area. The availability of critical mass ensures sufficient opportunity for making significant progress. Coupled with peer-review based assessment, this allows one to determine whether a specific investigation is making progress. In combination with critical resource allocation, the periodic adjustment of the programme ensures that the right balance between deep technical work in one area and cross-disciplinary work is attained.

#### Cross Cutting Research Themes

For the ITA to be successful it needed to break down the barriers to collaboration between the disparate scientific disciplines. One way to break down barriers, build relationships, and develop mutual understanding was to identify cross cutting research themes that encouraged researchers across the ITA technical areas to work together on difficult scientific challenges requiring a multidisciplinary approach to make fundamental progress.

## Summary

The ITA programme has proven itself to be highly successful, with a significant impact in advances in fundamental science, successful transitions of its technology,

trained scientists, and in international multi-disciplinary collaboration. This success was in part a result of the lessons that we learned and applied throughout the programme.

### References

- T. Killion, P. Sutton, M. Frame and P. Gendason, "A New Paradigm in International Collaboration: The US-UK International Technology Alliance in Network and Information Sciences", *RUSI Defence Systems*, pp. 46-49, June 2007.
- [2] Committee on Network Science for Future Army Applications; Board on Army Science and Technology; Division on Engineering and Physical Sciences; National Research Council, "Network Science", National Academies Press, ISBN 978-0-309-10026-7, 2005.

## A NIS ITA Leadership

Collaborative Alliance Managers	
John (Jay) Gowens (ARL) John Pellegrino (ARL) Greg Cirincione [Acting] (ARL)	Jack Lemon (MOD) Fiona Cotter (Dstl) Graham George (Dstl) George Vongas (Dstl)
Programme Managers	
Dinesh Verma (IBM US)	David Watson (IBM UK)
Technical Area 1: Network Theory	
Ananthram Swami (ARL) Don Towsley (UMass) Kang-Won Lee (IBM US)	Tom McCutcheon (Dstl) Stuart Farquhar (Dstl)
Technical Area 2: Security across a	System of Systems
Greg Cirincione (ARL) Dakshi Agrawal (IBM US)	Trevor Benjamin (Dstl) Robert Bourne (Dstl) John McDermid (Univ. of York)
Technical Area 3: Sensory Informat	ion Processing and Delivery
Tien Pham (ARL) Thomas La Porta (Penn State)	Gavin Pearson (Dstl) Flavio Bergamaschi (IBM UK)
Technical Area 4: Distributed Coali	tion Planning and Decision making
Mike Strub (ARL) Cheryl Giammanco (ARL)	Jitu Patel (Dstl) Nigel Shadbolt (Univ. of Southampton) Graham Bent (IBM UK)
Technical Area 5: Coalition Interope	erable Secure & Hybrid Networks
Greg Cirincione (ARL) Don Towsley (UMass) Mudhakar Srivatsa (IBM US)	Trevor Benjamin (Dstl)
Technical Area 6: Distributed Coali Making	tion Information Processing for Decision
Tien Pham (ARL)	Gavin Pearson (Dstl) Alun Preece (Cardiff Univ.) Dave Braines (IBM UK)

### Metrics

*Intellectual property:* at the time of printing 53 patents have been filed over the course of the NIS ITA programme. Full details and the latest information is available on the NIS ITA web site at *http://nis-ita.org*.

*Degrees awarded:* a total of 64 Doctor of Philosophy degrees have been awarded to students working on the NIS ITA programme, along with a further 14 Master's degrees:

Degree Awarded	2007	2008	2009	2010	2011	2012	2013	2014	2015
Ph.D.	7	5	6	13	3	11	5	5	9
Master's	2	3	1	3	1	0	4	0	0

## Bibliography of Scholarly Publications

A full and definitive list of the scholarly publications generated throughout the NIS ITA program can be seen at http://nis-ita.org, including details of the authors and organizations involved in each publication, and where they were presented or published. The purpose of this web site is to provide a full science library of publications that will persist beyond the completion of the NIS ITA programme, leaving the rich legacy available for further reference and exploitation by the wider scientific and commercial community.

At the time of publication, the NIS ITA had produced well over 100 peer-reviewed journal papers and more than 600 peer-reviewed external conference papers.

A key feature of the programme was the annual ITA conference that served as an excellent opportunity to publicize ongoing research progress, and also socialize some of the more disruptive or unusual ideas through provocative short papers. Throughout the 10 years of the NIS ITA more than 500 long and short papers were produced and presented at these regular internal conferences, that were accompanied by numerous workshops and demonstrations.

The headline publication numbers are, however, only the beginning of the story of the NIS ITA's legacy. The data from this record shows a rich set of collaborations

between authors from academia, industry and government in both the US and UK, evidencing the highly collaborative nature of the programme.



Figure 30. Visualizing NIS ITA paper co-authors

For example Figure 30 shows the NIS ITA co-author relationships between a government author (shown in green at the centre of the graph) with their academic (blue) and industrial (orange) co-authors, the thickness of the lines indicating the strength of the collaboration as derived from the number of papers co-authored. The data also shows the degree of collaboration and influence that the programme had beyond the organizations and individuals that made up the core Alliance. At the time of publication over 140 organisations from 24 countries had contributed to ITA publications, with over 700 individual co-authors.

These details and many more can be found in the online science library for the NIS ITA at http://nis-ita.org.

## C Glossary

ACT-R	Adaptive Control of Thought-Rational	
ADS	Authenticated data structure	
AIBE	Anonymous identity-based encryption	
AI-TECD	Actionable Intelligence Technology Enabled Capability Demonstration	
API	Application programming interface	
ARL	U.S. Army Research Laboratory	
ASP	Answer Set Programming	
AWGN	Additive white Gaussian noise	
BICES	Battlefield Information Collection and Exploitation System	
BPP	Biennial Program Plan	
CCD	Capability Concept Demonstration	
CE	Controlled English	
CENTCOM	U.S. Central Command	
CERDEC	U.S. Communication-Electronics Research, Development and Engineering Center	
CIAI	Coalition ISR Assets Interoperability	
CIDNE	Combined Information Data Network Exchange	
CIM-SPL	Common Information Model Simplified Policy Language	
СМИ	Carnegie Mellon University	
CNA	Cultural network analysis	
Col	Community of Interest	
COIL	Common Open Interface Land	
COIN	Counterinsurgency	
COPD	Comprehensive Operations Planning Directive	
COTS	Commercial off-the-shelf	
СРМ	Collaborative Planning Model	
CSD	Coalition Shared Data Server	
СТА	Collaborative Technology Alliance	
CUNY	The City University of New York	
CWP	Coalition Warfare Program	
DDFD	Dynamic Distributed Federated Database	

DRDC	Defence Research Development Canada
DE&S	Defence Equipment and Support
DoD	U.S. Department of Defense
Dstl	UK Defence Science and Technology Laboratory
DTC	Defence Technology Centre
DTN	Disruption tolerant network
DTRA	Defense Threat Reduction Agency
FHE	Fully homomorphic encryption
FPGA	Field-programmable gate array
GWSS	Green wave sleep scheduling
HIBE	Hierarchical identity-based encryption
I2QS	International Workshop on Information Quality and Quality of Service for Pervasive Computing
I2WD	Intelligence and Information Warfare Directorate
ICT	Information Communications Technology
IBE	Identity-based encryption
IDRM	Inter-Domain Routing Protocol for MANETs
IED	Improvised explosive device
ISP	Internet Service Provider
ISR	Intelligence, surveillance and reconnaissance
ITA	International Technology Alliance
KBSP	Knowledge-based security policy
LOSA RED	Land Open Systems Architecture Research, Experimentation and Development
LPD	Low probability of detection
MPD	Markov decision process
MAC	Media access control
MANET	Mobile ad hoc network
MIPS	Management of Information Processing Services
MMF	Military Missions and Means Framework
MOD	UK Ministry of Defence
MOTS	Military off-the-shelf
МРТСР	Multipath TCP
МТО	Memory-trace oblivious
NATO	North Atlantic Treaty Organization

NC3A	NATO Consultation, Command and Control Agency			
NCO	Network-centric operations			
NEC	Network-enabled capability			
NGO	Non-governmental organization			
NIFC	NATO Intelligence Fusion Centre			
NIS ITA	Network and Information Sciences International Technology Alliance			
NITB	NATO Intelligence Toolbox			
NLP	Natural language processing.			
NS CTA	Network Science Collaborative Technology Alliance			
NUM	Network utility maximization			
oABE	Outsider-anonymous broadcast encryption			
OBIE	Ontology-based information extraction			
OCD	Operational Concept Demonstrator			
ORAM	Oblivious RAM			
OSCAR	Open Source Communications Analysis Research			
osus	Open Standards for Unattended Sensors			
OWL	Web Ontology Language			
OWL DL	Web Ontology Language Description Logic			
PCA	Principal component analysis			
PDP	Policy decision point			
Penn State	Pennsylvania State University			
PEP	Policy enforcement point			
РНҮ	Physical layer			
PVO	Private voluntary organization			
PWAS	Persistent Wide Area Surveillance			
Qol	Quality of Information			
RDBMS	Relational database management system			
S&T	Science and Technology			
SAM	Sensor Assignment for Missions			
SINR	Signal-to-interference-plus-noise ratio			
SLQA	Store locally query anywhere			
SMC	Secure multi-party computation			
SME	Subject-matter expert			
SWRL	Semantic Web Rules Language			

SQL	Structured Query Language	
ТА	Trusted authority	
	Technical Area	
TAL	Technical Area Leader	
ТСР	Transmission Control Protocol	
TTL	Time to live	
TIGR	Tactical Ground Reporting System	
TOPFAS	Tool for Operational Planning, Force Activation and Simulation	
UCLA	University of California, Los Angeles	
UMass	University of Massachusetts, Amherst	
VKB	Virtual knowledge base	
VOS	Verifiable Oblivious Storage	
VQB	Visual Query Builder	
WEA	IBM Watson Engagement Advisor	
WIBE	Identity-based encryption with wildcards	
WPML	Watson Policy Management Language	
XACML	Extensible Access Control Markup Language	
XML	Extensible Markup Language (XML)	



The Network and Information Sciences International Technology Alliance, initiated in 2006, marked a new paradigm in international collaborative research. It brought together a diverse set of researchers from academia, industry and government in the US and the UK to collaborate in making advances in the field of network and information sciences, geared towards coalition operations. It broke new ground in several areas, established a new way of working, and focused on high-risk basic research addressing hard problems. In reviews conducted by peers from the scientific, academic and governmental establishment, the Alliance was widely praised for its high level of collaboration and the quality of scientific achievements.

As this research Alliance draws to a conclusion in 2016, after ten years of successful collaborative exploration, it is time to reflect on the achievements made and lessons learnt. This book provides an overview of those achievements, enumerating the scientific advances, how some of those advances were applied to create new systems, and most importantly the lessons learnt from running a successful collaborative international research program.



