



ADVANCED DECISION ARCHITECTURES  
FOR THE WARFIGHTER:  
FOUNDATIONS AND TECHNOLOGY



EDITED BY PATRICIA MCDERMOTT AND LAUREL ALLENDER



# ADVANCED DECISION ARCHITECTURES FOR THE WARFIGHTER: FOUNDATIONS AND TECHNOLOGY

---

EDITED BY PATRICIA MCDERMOTT  
AND LAUREL ALLENDER

This book, produced by the Partners of the  
Army Research Laboratory Advanced Decision Architectures  
Collaborative Technology Alliance, also serves as the  
Final Report for DAAD19-01-2-0009.



SECTION I

---

COLLECTING, PROCESSING, AND DISTRIBUTING BATTLEFIELD  
INFORMATION



DYNAMIC NETWORK ANALYSIS APPLIED TO EXPERIMENTS  
FROM THE DECISION ARCHITECTURES RESEARCH ENVIRONMENT

---

Kathleen M. Carley, Ph.D.  
*Carnegie Mellon University, Pittsburgh, PA*

Michael K. Martin, Ph.D.  
*Carnegie Mellon University, Pittsburgh, PA*

John P. Hancock  
*ArtisTech, Inc., Fairfax, VA*

## INTRODUCTION

ADA CTA research is producing experimental results from simulations of intermeshed networks of warfighters and battlefield surveillance assets. These networks form a complex system with behaviors that emerge from patterns of interaction among constituent entities. The simulated interactions are spatially situated, temporally distributed communications among people, robots, and software agents. In general terms, the complex system – parts of which we address in this paper – can be conceptualized as a two-level meta-network that includes interactions among human agents at one level, interactions among artificial agents at another level, and cross-level interactions between human and artificial agents.

The issue addressed in this paper is Dynamic Network Analysis (DNA) of system behavior. For the purposes of this chapter, we are not interested in examining the performance of one system relative to another. Although this type of comparative analysis can be useful for researchers or system designers, it is of questionable use to warfighters. Instead, we are interested in analyses that produce tactically relevant, actionable results that highlight the strengths and weaknesses in the system being

observed. To be useful, the analytical results must foster tactical insight and stimulate battlefield decisions that prudently influence future system behavior.

To this end, we describe two case studies that apply DNA to the simulated battlefield data being generated by experiments in the Decision Architectures Research Environment (DARE). The first case involves intercepts of simulated communications among human agents, which we frame as an exercise in adversarial reasoning. The second case involves simulated communications among surveillance assets (i.e., software agents and robots), which we frame as an exercise in understanding the automated control of a Persistent Coordinated Video Surveillance (PCVS) system. Together, we believe the case studies demonstrate how DNA (Carley, 2002) can foster tactical insight in complex multi-entity scenarios. They also demonstrate that the combination of DNA techniques required for tactical insight may vary according to the type of network being analyzed. Finally, they show how DNA assists in the development, understanding and tuning of software agent systems.

### **ANALYTICAL APPROACH**

One approach to analyzing system behavior employs mainstream statistical techniques on summary measures of performance (e.g., percent of targets tracked, percent of priority targets tracked with acceptable accuracy, track correlation, etc.). This type of analysis, however, does not fully exploit the information generated in DARE experiments – or by battlefield surveillance in general. More to the point, it provides little insight into the identification of levers in the networks underlying system behavior. Once identified, levers can be used to determine actions to take to influence future system behavior.

DNA provides an alternative analytical approach that compliments mainstream statistics. With DNA, the focus shifts from aggregate measures of performance for a collection of battlefield entities to the performance implied by the structure of relations among battlefield entities. This shift is the essence of what it means to view the battlefield from a network science perspective. That is, from a network science perspective we are less concerned with what is normal (e.g., averages, dispersions) about entities in the battlefield (e.g., people, places) and more concerned about detecting substantive patterns in the observed relations among entities. The emphasis on the structure of relations in DNA makes it particularly well-suited to the detection of anomalies and exceptions (e.g., centralities, exclusivities) – the levers with potentially large influ-

ences on system behavior. DNA, therefore, fosters scrutiny of strengths and vulnerabilities in the relations among battlefield entities (i.e., the observed system). With a DNA model, we can identify (among other things) implicit groups of entities, key people and locations, and operationally significant time frames. We can even begin to infer relations among entities where none have been observed.

Network science has been hampered historically by a dominantly social perspective focusing on who interacts with whom. However, Carley (2002) argued that these social networks exist within an ecology of networks that can usefully be characterized in terms of the dynamics of the relations among the who, what, where, how, and why. This is known as the meta-network perspective. \*ORA (e.g., Carley, Columbus, DeReno, Reminga & Moon, 2008) is a dynamic network analysis package that can be used to assess multi-mode, multi-plex networks; identify key players, groups and vulnerabilities; enable comparison of two or more networks; and facilitate reasoning about spatio-temporal networks.

\*ORA supports analysis of dynamic networks in many ways: (1) comparisons of temporally ordered snap-shots of static networks, (2) statistical change detection on sequences of networks, (3) trail analysis for trail data and conversion of trail data to networks, (4) simulation of change in networks, and (5) comparative statics for immediate impact assessment. Herein, we make use of statistical change detection and trail analysis, along with \*ORA's visualization capabilities.

### **CASE STUDY1: TERRORISTS IN ADELPHI**

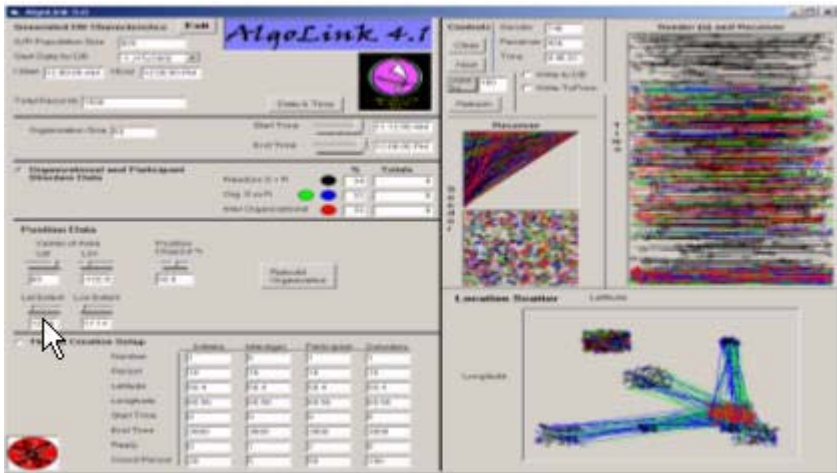
The scenario for case study 1 was framed as intercepts of simulated communications among agents representing terrorists and noncombatants (e.g., pizza delivery guys). The agents communicated via phone and email as they moved about the Adelphi region.

#### **Data Generation**

The data were generated by ArtisTech's AlgoLink simulator. The AlgoLink simulator was originally developed to test the capabilities of message analysis tools to support intelligence analysis requirements in battlefield communications. AlgoLink (see Figure 1.1) facilitates custom construction of entity networks that follow specified communication structures, times, places, durations, and behaviors. It generates both the foreground (network of interest) and the background communication as specified and stitches them together into a single communication record.



Background communication is both structured (e.g. inter-mingled hierarchical and social organizations) and random as specified by the human simulation operator. AlgoLink uses realistic communication and organizational data, timing, and morphologies but contains no information about any real person or organization.



**Figure 1.1**  
The AlgoLink message simulator interface


The data-set created for this experiment was generated by ArtisTech staff looking forward to system-based Intelligent Agent communication behavior analysis. The entities of interest were organized into a small number of “cells” that were uniquely connected and stitched into a larger background “community.” The data-set was geographically centered on the ARL Adelphi campus because the ARTEMIS project is coordinated with the ARL Computation and Information Sciences Directorate research systems. The simulated social structure was morphologically similar to a medium-sized community of intelligent agents acting with a specific purpose in a battlefield Command, Control, and Communications (C3) system. To simulate “an event” that stimulates the communication network, a spike in communication volume was inserted at a selected time. The AlgoLink-generated data-set was a rapid way to assess the feasibility of collaboration between the ARTEMIS-PCVS and CASOS teams.

## Analysis

The AlgoLink output was delivered as an XML file containing a sequence of communications records. Each simulated communication record identified the sender, the receiver, the time the communication occurred, its duration, and whether its content was operationally relevant, irrelevant, or ambiguous. It also contained latitude and longitude for the position of the mobile senders and receivers during each communication.

Our analysis strategy in this case can be generally described as an overview-and-zoom. That is, we first examined the general context of communications activities, and then drilled down to determine important agents, time-frames, and locations. The subset of \*ORA capabilities that proved particularly useful here included geospatial visualization, key player identification, change detection analysis, and the correlation of standard network and geospatial visualizations.


Using \*ORA's geospatial visualization capabilities (e.g., Davis, Olson, & Carley, 2008) we exploited the presence of time-varying, geo-located attributes of the intercepted communications to discover the scenario involved suspicious entities fleeing the Adelphi area (see Figure 1.2).



**Figure 1.2**

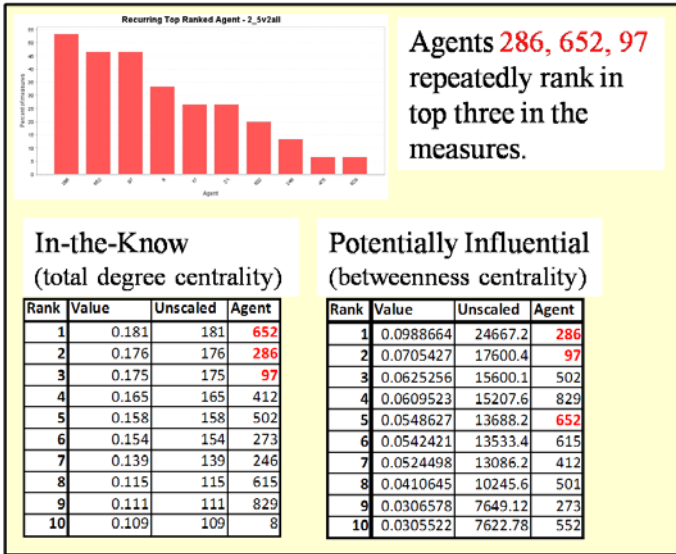
Agents fleeing Adelphi over the time course of the scenario

Using a fuzzy group clustering technique, FOG (e.g., Davis & Carley, 2008), we found that the suspicious entities were organized into five groups with shared members (see Figure 1.3). The interstitial members are likely to contain the coordinators and leaders.



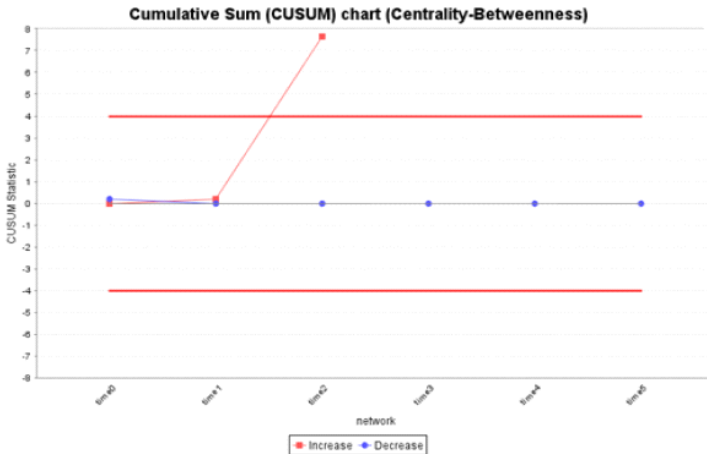
**Figure 1.3**  
Agents organized in 5 groups that share members

To drill down, we first used \*ORA's Key Entity Report to identify the three agents most critical to operations (see Figure 1.4). Because the data were about communications, two different centrality measures were used – degree centrality and betweenness centrality. Degree centrality measures who is connected to most others (i.e., the actor most likely to be “in-the-know”). Betweenness centrality measures who is most likely to be on all the paths by which information flows (i.e., the actor most likely to be influential). This enabled narrowing our focus to a small group of leaders instead of focusing on the set of interstitial members.



**Figure 1.4**  
Key Entity Report identifies 3 important agents


We next asked, is something happening? One way of answering this is to see whether there is a change in standard behavior. Using the Change Detection Report (see Figure 1.5), we identified period 3 as the time-frame in which operations most likely occurred.



**Figure 1.5**  
Change Detection Report signals time period 3 is different


The network change detection analysis (McCulloh & Carley, 2008; 2009; McCulloh, Webb, Graham, Carley & Horn, 2008) extends change detection from operations research, where it has been used on variable level data, to relational data. This is a statistical approach for detecting small persistent changes in organizational behavior over time using statistical process control techniques applied to network summary statistics. Period 2 was the point at which organizational behavior changed, leading to radical difference by period 3. This appears to have been a planning-execution phase shift.

Examination of individual-level metrics for the three key players and network-level metrics (e.g., centrality, betweenness, efficiency, connectedness) corroborated our interpretation that period 3 was operationally significant. Individual-level metrics indicated that actor 286 engaged in extensive coordination at period 2, passing the reigns of control to actor 652 at period 4 (see Figure 1.6).



**Figure 1.6**  
Individual-level metrics converge on period 3

Examination of network-level metrics (see Figure 1.7) showed that the group was generally a very distributed structure that coordinated into a centrally controlled, more efficient, unit at period 3. Then, it went back to its

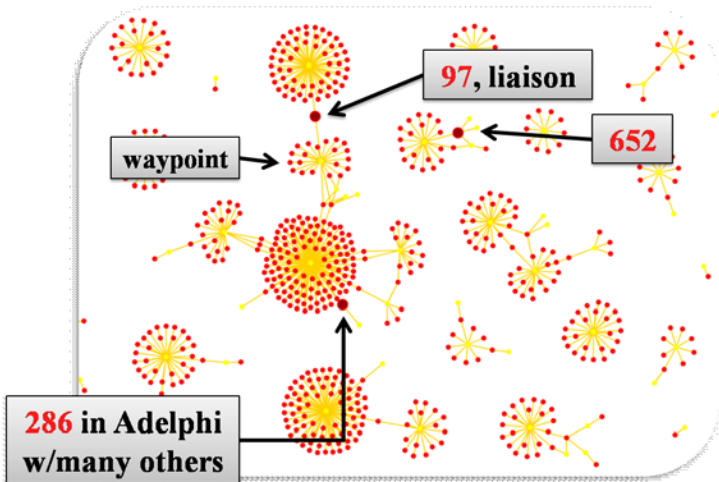


**Figure 1.7**

Network-level metrics converge on period 3

standard form and became more “hidden” in the socio-communication environment.

Having identified the key players and time period, we focused our analyses on discovering what may have happened. Examination of the agent x location network for period 3 (see Figure 1.8) indicated a large cluster of suspicious entities in the Adelphi area (including key player, Agent 286), a fairly large cluster of suspicious entities (including another




**Figure 1.8**

Agent x Location Network for period 3

key player, Agent 97) in what appears to be a staging area, an apparent waypoint between the staging area and the cluster of suspicious entities

in Adelphi (also visited by Agent 97, who seems to be a liaison), and a runner (key player, Agent 652) who visits many locations with few suspicious entities present.


Another key advance, developed as part of the CTA, was the capability to move between trail data – who was where when – and networks. When we examined the trails visualization for period 4 (i.e., the period immediately following the apparent operation, perhaps a period of initial surveillance), we saw that the three key players were never in the same place at the time; Agent 652 was again running, whereas the activities of Agents 97 and 286 were restricted to one or two areas. The apparent coordination handoff from Agent 286 to Agent 652 is related to 652's increase in spatial movement and coordination needed due to increased movement. In the trails visualization (see Figure 1.9), time progresses down the y-axis. Geographic regions form vertical bins along the x-axis. Arrows are plotted as agents move from region to region (or within regions), and in this case are color coded to the three key players.



**Figure 1.9**

\*ORA Loom visualization of trails for 3 key players during period 4

Finally, correlating a standard agent x location network visualization with a geospatial visualization for the end of the scenario we found that Agent 286 was alone with a single movement between two locations in Adelphi, Agent 97 was holed up with a sizable group of suspicious entities north of Adelphi, and Agent 652 was alone but on the run (see Figure 1.10).



**Figure 1.10**

Correlation of agent x location network and geographic location at scenario end

Given the pattern of communications and movement during the scenario, two courses of actions appear reasonable: (1) scour Adelphi for a bomb, IED, etc. planted during operations in period 3, or (2) go after dispersed suspicious entities. With respect to action 2, Agent 286 may be an easy target with direct knowledge of the operations that occurred during period 3. Targeting the location where Agent 97 is hiding, however, will yield more suspicious entities. Note that the simulated data does not include communications content so specification of the event is not possible, but the applied DNA analysis accurately identified the time, place, and lead entities in the simulation.

## **CASE STUDY 2: MOVEMENT IN THE PERIMETER**

The scenario for this case centered on an automated surveillance system (i.e., the ARTEMIS-PCVS system) that is responsible for identifying moving entities within the perimeter of a Blue Force research compound. The system divides the perimeter into four areas of responsibility, where each area is assigned to a “Tasking Agent” that is responsible for surveillance (see Figure 1.11). Tasking Agents get simulated movement reports from simple video analysis algorithms. Tasking Agents then prioritize “targets” and assign mobile robotic assets to pursue and identify the targets.

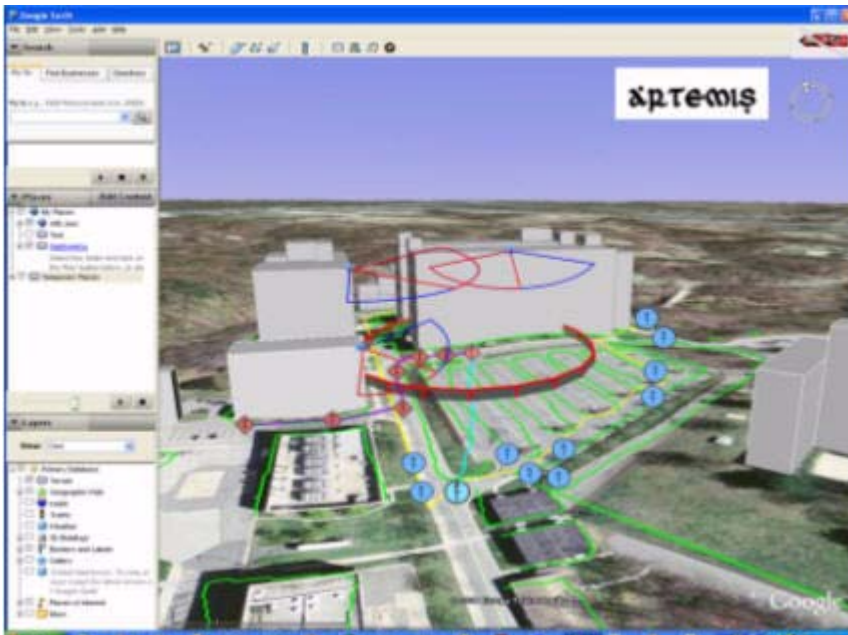




**Figure 1.11**  
Tasking Agent Areas

### Data Generation

The data were generated by ArtisTech's ARTEMIS-PCVS system prototype (see Figure 1.12). This prototype uses hundreds of small reasoning



**Figure 1.12**  
3D View of ARTEMIS-PCVS scenario on the GIS ARL Adelphi testbed model

algorithms encapsulated in software agents. The agents communicate with humans, agents, and other system elements and even create and delete other agents with frequencies that are dependent on their reactions to the sensed environment (simulation).

The particular experimental run analyzed here was conducted to determine whether Tasking Agent reasoning and communication about sharing mobile robotic assets was functioning as expected. The scenario included a short time of quiescence in the compound, followed by the injection of a relatively large number of moving targets that moved about the compound using reasonable paths but freely crossing areas of surveillance responsibility. When targets cross areas of responsibility the complexity of reasoning increases and requires that Tasking Agents transfer (handoff) tracking and even possibly “lend” robotic assets to other Tasking Agents. The act of lending the asset involves communication to notify an adjacent Tasking Agent of an “incoming” unidentified entity, and a multi-message Tasker-Global Tasker “handshake” to transfer control. There were two of these handoff events in the scenario. This scenario was a simple one to facilitate early collaboration between ARTEMIS and CASOS staff.

## **Analysis**

As with case study 1, the data were delivered as an XML file comprised of a sequence of communications. The communications were between software agents or software agents and robots.

Although the data sets were superficially similar (i.e., logs of communications records), case study 2 presented several challenges not present in case study 1. These challenges arose primarily because the data for this case were relatively impoverished. DNA requires large data sets and we did receive more data for case study 2 than for case study 1. However, the extra data were of little benefit because they provided little information regarding the structure of the system we were analyzing. In terms of structural information, the extra data were mostly redundant.


The analysis was complicated by the fact that the scenario included only two instances of the target handoff event –the signal we were to detect. In structural terms, this means that we were looking for a change in network structure that involved (per ArtisTech’s description of the handshake) three links at most. Thus, the statistical change detection analysis used in case study 1 was of no use. Such a small change in structure would not be detected as being statistically significant.

The analysis was further complicated by the absence of geo-coordinates for the mobile robots. Therefore, \*ORA's geospatial visualization capabilities could not be employed. The insight that can be gleaned from visualizations of agents x locations, as used in case study 1, was also lost. Without location, we also lacked any means for constructing trails data to examine who was where when.

Finally, the goal of analysis in case study 2 differed from case study 1. In case study 1, we employed DNA techniques designed to identify important entities, locations, and times. The purpose of the analysis, therefore, was to identify centralities. In contrast, the purpose of the analysis for case study 2 was to identify exclusivities (i.e., the two Tasker-Global Tasker handshakes).

Given the impoverished data set and our goal of detecting only two instances of the handoff among a small set of agents and robots, the analysis strategy we adopted for this case was one of converging operations (to "kick-start" a DNA of richer data in the future). ArtisTech personnel (with their knowledge of the system) manually analyzed the data. CASOS personnel applied DNA techniques to the data. The following discusses only two of the issues we addressed.

To find the handoff where one Tasking Agent loaned a robot to another Tasking Agent, we relied on ArtisTech's identification of message-types that indicate such a handoff occurred. Three message-types were related to the handoff: RemoveIdentity, ReportPositionToSelectedTasker, and ReportPositionToSelectedTaskerReturn. Finding the handoff was then simply a matter of using \*ORA's Sphere of Influence capabilities to visualize which agents were involved. Figure 1.13 shows the three important messages, along with the agents that sent and received them. It can be seen that Tasking Agent 3 is positioned in the network differently than Tasking Agents 1, 2, and 4. Furthermore, we see that Tasking Agent 3 was the only agent to send the ReportPositionToSelectedTaskerReturn message (as indicated by the arrow on the link between the agent node and the message node). Global Tasking Agent 7 received this message and sent a RemoveIdentity message to Tasking Agent 3. The sphere of influence also indicates that the ReportPositionToSelectedTasker message was probably not uniquely related to the handshake.



**Figure 1.13**

Sphere of Influence visualization showing robot handoff

In anticipation of richer data sets, the second issue we examined was whether we could use DNA techniques to partition agents into foreground and background agents. To this end, the ArtisTech team produced a meticulous message trace analysis and event identification as ground truth for the CASOS team. Even as fully knowledgeable designers of the communication logic this analysis and documentation took more than 4 hours using simple text search and a numeric message frequency analysis provided by CASOS. The complexity of this analysis underscores the need for network analysis techniques for a more complete system experiment analysis.

Per ArtisTech’s analysis, agents could be divided into foreground agents that were substantively related to the scenario and background agents that existed simply to make the simulation run. Our task at CASOS, therefore, was to employ one or another DNA technique to separate foreground from background agents. We found that the Newman grouping algorithm worked well, separating foreground and background agents into groups that closely approximated the manual analysis. Seven of ten background agents were correctly classified. But there was disagreement between ArtisTech judges regarding whether one of the three misclassified agents was a background or a foreground agent. One foreground agent out of 29 undisputed foreground agents was misclassified. That the Newman grouping algorithm corresponds fairly well with the judgments of domain experts is promising, and suggests a line of research concerning the “psychological validity” of Newman grouping.

## SYSTEM UNDERSTANDING AND TUNING ACTIVITIES

ArtisTech postulated the additional benefit of performing DNA on the output logs from the internal system communications; the analysis supported understanding of *how* the team and system achieved the measured performance. The ARTEMIS research team was in the process of setting up system model experiments to study the advantages that automated reasoning could add to widely used numeric image and video processing for the purpose of PCVS. As we modeled the reasoning and set it into the simulation we have conducted many test runs. We can easily see when the macro system behavior is as designed and when it deviates from intentions. However, to determine why the deviations occur and how the reactive agent networks achieve intended system performance requires message level analysis. ArtisTech shared mixed-result test logs with CASOS specifically to facilitate the identification of expected and unexpected results, and how they arise.

Results from initial analyses in both case studies diverged from expected results. In hindsight this is hardly surprising. Design of a complex, multi-agent system with emergent reactive system behavior may not be possible without the support of network analytics.

In both cases we found evidence of pragmatic, technically correct, programming practices that interfered with substantive DNA of the communications logs. ArtisTech confirmed that the experiment runs that were analyzed were not considered final or expected to be entirely correct, merely typical in content and form. Whether we view this as a verification issue (i.e., building the thing right) or a validation issue (i.e., building the right thing) depends on perspective. Given that verification is primarily an internal activity; the degree to which programmers achieved the intent behind specifications is a matter of interpretation. From the perspective of network analysis (i.e., users of the output logs), however, we can note a minor short-fall in validation. Specifically, communication records that are not relevant to the scenario being analyzed should be filtered from the log to increase usability.

In the initial analysis, we experimentally illustrated that DNA approaches provide a capability to analyze complex message sets to find particular behavior patterns of interest. In the second analysis we began forging a collaborative research approach designed to unearth the analytic steps – the combinations of DNA techniques – required to analyze different types of messaging behavior.

## LESSONS LEARNED AND FUTURE DIRECTIONS

Open collaboration between data providers and network analysts created a beneficial gap between expected and observed system behavior. As data provider, ArtisTech developed the multi-agent system and environment simulations, designed the scenarios, and conducted the simulation-based experiments – with expected system behaviors in mind. The nonlinearities inherent to complex systems comprised of interacting agents, however, make it notoriously difficult to predict emergent and reactive behavior, and are indeed the reason computer simulation is necessary. As network analyst, CASOS received militarily relevant battlefield simulation data without prior knowledge of expected system behavior. The challenge to CASOS, therefore, was to use DNA to characterize what happened in the mysterious scenarios received – and CASOS observations initially diverged from ArtisTech default expectations. To resolve the discrepancy, we used DNA to examine why the simulations did not behave as expected. Thus, the postulated benefit of the ArtisTech-CASOS collaboration was the use of network analytics to gain insight into the performance of multi-agent systems. ArtisTech, with CASOS, now plans to use this approach and DNA tools to build a set of reusable system experiment analysis methods that will be applied to understand how variant Human/Automation PCVS experiments achieve observed results using networks of communication and behaviors.

With regard to demonstrating the tactical relevance of DNA, the availability of military scenario data is invaluable. It provides opportunities to combine extant DNA techniques into analytic strategies that produce results warfighters can use. Generally, we found that the family of DNA metrics and visualizations that have been implemented in \*ORA over the past 10 or so years provides an ample basis for conducting tactically relevant, actionable analyses on battlefield data.

Two \*ORA capabilities were particularly helpful during this effort: geo-visualization and network change detection. The (not-so-simple) act of placing networks on a map puts abstract social networks into a concrete, spatial context. It provides explicit information about where the action is taking place. The capacity to detect structural changes among temporally ordered networks provides explicit information about when the action is taking place.

The case studies also helped to identify several areas where future development could improve \*ORA's tactical relevance. The improvements would generally support delivery of one or more Tactical Insight Reports. As envisioned, these reports would contain the outputs of all

DNA techniques that contribute to a particular analytic strategy. From the two case studies described above, it appears that correspondences among networks of the relations among entities, networks of geospatially anchored entities, and networks distributed over time will play a central role in such reports.

For ArtisTech the goal of DNA applications to system analysis and monitoring is to investigate the creation of a generalized data structure and DNA method combination approach that will allow the encoding of expected or detected message and behavior patterns. The initial application of this concerns post experiment review. The more far-reaching implications of such DNA mechanisms extend to real-time embedded system monitoring to increase distributed system security and user trust.

### **ACKNOWLEDGEMENTS**

This work is part of the Dynamics Networks project at the center for Computational Analysis of Social and Organizational Systems (CASOS) of the School of Computer Science (SCS) at Carnegie Mellon University (CMU). For \*ORA, for the parts used in this analysis, the geo-spatial components were supported by ARO and ERDC-TEC W911NF0710317, the \*ORA-Loom by ONR N00014-06-1-0104, the over time change detection basic analysis by ARI W91WAW07C0063 and the ORA implementation of change detection by ARL DAAD19-01-2-0009, the fourier analysis by ONR N00014-06-1-0104, the key entity analysis was supported by ONR N00014-06-1-0104, extensions to handle communications networks by ARL through the Communications and Networks (CN) Collaborative Technology Alliance (CTA) 20002504 and the immediate impact assessment by ONR N00014-06-1-0104. The basic research conducted here to explore how to use these technologies for adversarial behavior and to link to the ArtisTech ARTEMIS-PCVS system, AlgoLink simulator, DARE was supported by the Army Research Lab through the Advanced Decision Architecture (ADA) Collaborative Technology Alliance (CTA) - DAAD19-01-2-0009. The views and proposals contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the Office of Naval Research, the Army Research Office, the Army Research Institute, the Army Research Lab or the U.S. government.

## REFERENCES

Carley, K.M. (2002). Smart agents and organizations of the future. In L. Lievrouw & S. Livingstone (Eds.), *The handbook of new media* (pp. 206-220). Thousand Oaks, CA: Sage.

Carley, K.M., Columbus, D., DeReno, M., Reminga, J., & Moon, I.C. (2008). *ORA User's Guide 2008* (Tech. Rep. No. CMU-ISR-08-125). Pittsburgh, PA: Institute for Software Research, School of Computer Science, Carnegie Mellon University.

Davis, G.B., & Carley, K.M. (2008). Clearing the FOG: Fuzzy, overlapping groups for social networks. *Social Networks*, 30 (3), 201- 212.

Davis, G.B., Olson, J., & Carley, K.M. (2008). *OraGIS and loom: Spatial and temporal extensions to the ORA analysis platform* (Tech. Rep. No. CMU-ISR-08-121). Pittsburgh, PA: Institute for Software Research, School of Computer Science, Carnegie Mellon University.

McCulloh, I., & Carley, K.M. (2008). *Detecting change in human social behavior simulation* (Tech. Rep. No. CMU-ISR-08-135). Pittsburgh, PA: Institute for Software Research, School of Computer Science, Carnegie Mellon University.





