

JoSS Article: Volume 6

Structural Knowledge and Success of Anti-Terrorist Activity: The Downside of Structural Equivalence

Maksim Tsvetovat, maksim@cs.cmu.edu

School of Computer Science, Carnegie Mellon University

Kathleen M. Carley, kathleen.carley@cmu.edu

Professor of Computer Science, Technology and Policy

Institute for Software Research International (ISRI), Carnegie Mellon University

This work was supported in part by Department of Defense, the Office of Naval Research(ONR), United States Navy Grant No. 9620.1.1140071, NSF IRI9633 662 and the NSF IGERT 9972762 for research and training in CASOS. Additional support was provided by CASOS - the center for Computational Analysis of Social and Organizational Systems at Carnegie Mellon University. The views and conclusions 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 Department of Defense, the Office of Naval Research, the National Science Foundation, or the U.S. government.

Abstract: Given the increasing threat of terrorism and spread of terrorist organizations, it is of vital importance to understand the properties of such organizations and to devise successful strategies for destabilizing them or decreasing their efficiency. However, intelligence information on these organizations is often incomplete, inaccurate or simply not available. This makes the study of terrorist networks and the evaluation of destabilization strategies difficult. In this paper, we propose a computational methodology for realistically simulating terrorist networks and evaluating alternative destabilization strategies. We proceed to use this methodology to evaluate and conduct a sensitivity analysis of the impact of various destabilization strategies under varying information surveillance regimes. We find that destabilization strategies that focus on the isolation of individuals who are highly central are ineffective in the long run as the network will heal itself as individuals who are nearly structurally equivalent to the isolated individuals "move in" and fill the communication gaps.

Introduction

For reasons of national security it is important to understand the properties of terrorist organizations that make such organizations efficient and flexible, and based on this understanding devise successful strategies to destabilize such organizations or curtail their efficiency, adaptability, and ability to move knowledge and resources. The assessment of destabilization strategies poses a number of key challenges. What does the underlying organization look like? Does it evolve? What strategies inhibit or effect the evolution so that the organization is destabilized? In this paper, we provide an approach to assessing destabilization

strategies that draws on work in organization science, knowledge management and computer science.

Terrorist organizations are often characterized as cellular organizations composed of quasi-independent cells and distributed command. In a sense, this is a non-traditional organizational configuration; hence, much of the knowledge in traditional organizational theory, particularly that focused on hierarchies or markets, does not apply. To be sure, lessons can be learned from the work on distributed and decentralized organizations that provides some guidance. This work demonstrates that such structures are often adaptive, useful in a volatile environment, and capable of rapid response [1] [2]. In other words, we should expect terrorist organization to adapt, and adapt rapidly. This suggests, that in general, they should be difficult to destabilize; however, the traditional organizational literature provides little guidance on how to destabilize the organization.

In general, the organization's form or design profoundly influences its performance, adaptability, and ability to move information [3]. It follows that organizations can be destabilized by altering their design. The one caveat here, is that organizations, particularly more distributed and decentralized ones, are continually evolving [4]. Terrorist organizations are often characterized as dynamic networks in which the connections among personnel define the nature of that evolution. This suggests that social network analysis will be useful in characterizing the underlying structure and in locating vulnerabilities in terms of key actors.

In general, organizations evolve as they face unanticipated changes in their environment, rapidly evolving technologies, and intelligent and adaptive opponents. Over the past decade, progress has been made in understanding the set of factors that enable adaptation and partially validated models of adaptive networks now exist [5]. A key result is that, in the short run, there appears to be a tradeoff between adaptivity and extremely high performance in organizations [6]. This suggests that forcing an organization to adapt should reduce its performance. Thus, even if an actor is no longer key, the mere isolation of that actor may be sufficient to be disruptive. However, to assess this a model of organizational change and network healing is needed.

Since the destabilization of terrorist networks could inhibit their ability to effect harm, there is a profound need for an approach that would allow researchers to reason about dynamic cellular networks and evaluate the potential effect of destabilization strategies. To be useful, such an approach must account for the natural evolution of cellular networks. This situation is further complicated by the fact that the information available on the terrorist network is liable to be incomplete and possibly erroneous. Hence, destabilization strategies need to be compared and contrasted in terms of their robustness under varying levels and types of information error. In other words, it would be misleading to judge destabilization strategies in terms of their impact on a static an unchanging network [7].

These problems suggest the need for a new methodological approach. In this paper, we provide an approach based on the use of a multi-agent network model of the co-evolution of the network of "observers" (the blue network) and the "terrorists" (the red network) in which the observers can capture only partial data on the underlying covert network and the covert network evolves both naturally and in response to attacks by the observers. This approach builds off of

organization theory and social network theory, as well as machine learning and dynamic network analysis. Specifically, we have developed a computational model of dynamic cellular organizations and used it to evaluate a number of alternative strategies for destabilization of cellular networks.

It is important at the outset to note that this examination of destabilization strategies is highly exploratory. We make no claims that the examination of destabilization strategies is comprehensive, nor that the types of "error" in the data that intelligence agencies can collect is completely described. Further, our estimate of the structure of the covert network is based on publicly available data much of which is qualitative and requires interpretation. Thus, this work should be read as a study in the power of an empirically grounded simulation approach and a call for future research. Further, we restrict our analysis to a structural or network analysis and focus on what does the covert network look like, how does its structure influence its performance and ability to pass information, how does it evolve, how can its evolution be altered (its behavior destabilized) through interventions focused on the nodes, and what interventions should be taken given the level of fidelity in the information that we have. Admittedly, in this complex arena there are many other factors that are critical, but they are beyond the scope of this study. Thus, from a straight social network perspective, this study suggests the types of methodological issues that will emerge when working with dynamic large scale networks under uncertainty.

To ground this paper, a short case description is provided of Al Qaeda with the focus on the network structure. In these two descriptions we draw on both military and organizational theory. This is followed by a discussion of the intelligence agencies engaged in anti-terrorist activity and the possible data and errors in said data. Our intent is to demonstrate, at a fairly high level, the context and the resultant information and modelling problems, not to provide a full analysis for intelligence or military operations. As good science often emerges from attacking hard real world problems, we are trying to provide sufficient detail to understand the basis for the problems that research must address, rather than simply provide a high theoretical description of general data problems. This is followed by a brief discussion of the applicability of traditional social network analysis and the need to take a dynamic network perspective. We then describe a computational model of terrorist organizations as dynamic evolving networks, and anti-terrorist bodies with emphasis on their information collection and destabilization strategies. A virtual experiment is used to examine destabilization strategies and the results are then discussed.

Covert Terrorist Networks - Al Qaeda

Extra-national terrorist groups generally serve to advance the interests of their leaders or direct backers (whether political, religious or commercial) and span multiple nations in their search for operatives and resources. Extra-national terrorist networks may enjoy support of one or several states whose political agendas coincide with the goals of the organization - but ultimately are not dependent on state support due to their ability to find independent financial backing from wealthy sympathizers. Commonly such groups are structured in a way similar to organized crime syndicates and employ networks of quasi-independent cells scattered through the region of operation of the organization as well as other countries that could be used as resource bases, recruiting and training centers.

Al Qaeda, arabic for "The Base," is the largest known extra-national terrorist organization. It is a large dynamic network, estimated to have the support of six to seven million radical Muslims worldwide, of which 120,000 are willing to take up arms [8]. Its reach is global, with outposts reported in Europe, Middle East, East Asia and both Americas. In the Islamic world, its task is to purify societies and governments according to a strict interpretation of the Koran and to use religion as a unification force for creation of an Islamic superpower state.

As Goolsby [9] stated, Al Qaeda extends its reach and recruits new member cells via adoption and of local Islamic insurgency groups. Beginning with provision of operational support and resources to facilitate growth, Al Qaeda representatives work to transform an insurgency group such as Jemaah Islamiyya (Indonesia) from a group seeking political change to a full-fledged terrorist organization executing multi-casualty attacks such as the Bali bombing in 2002 [10].

Al Qaeda's global network, as we know it today, was created while it was based in Khartoum, from December 1991 to May 1996. To coordinate its overt and covert operations as Al Qaeda's ambitions and resources increased, it developed a decentralized, regional structure. Al Qaeda pursues its objectives through a network of cells, associate terrorist and guerilla groups and other affiliated organizations. For instance, the Sudanese, Turkish and Spanish nodes ran clandestine military activities in Europe and North America.

The worldwide nodes appear to have no formal structure and no hierarchy. Assignments are often carried out by individuals and small groups designated for the purpose as "the person responsible." The regional nodes appear not to have a fixed location and move quickly when dictated by the political situation in the region. Al Qaeda shares expertise, transfers resources, discusses strategy and sometimes conducts joint operations with regional terrorist groups.

Although the *modus operandi* of Al Qaeda is cellular, familial relationships play a key role. As an Islamic cultural and social network, Al Qaeda members recruit from among their own nationalities, families and friends. What gives Al Qaeda its global reach is its ability to appeal to Muslims irrespective of their nationality, enabling it to function in eastern Asia, Russia, western Europe, sub-Saharan Africa and North America with equal facility.

Unlike conventional military forces which are often hierarchical and centralized, terrorist militant units are often small, dispersed and seemingly disorganized. Nevertheless, they have been able to effectively counter much larger conventional armies. Large terrorist organizations operate in small, dispersed cells that can deploy anytime and anywhere [11]. Dispersed forms of organization allow these networks to operate elusively and secretly.

The apparent structure of Al Qaeda is not exclusive to such militant or terrorist groups. Indeed, they bear a family resemblance to the structure of other resistance groups. For example, a study published in 1970 by L. Gerlach and V. Hine [12] concluded that U.S. social movements, such as the environmental and anti-war movements in the 1960s, were structured as "segmented, polycentric, and ideologically integrated networks" (SPINs).

"By segmentary I mean that it is cellular, composed of many different groups... . By polycentric I mean that it has many different leaders or centers of direction... . By networked I mean that the

segments and the leaders are integrated into reticulated systems or networks through various structural, personal, and ideological ties. Networks are usually unbounded and expanding... . This acronym [SPIN] helps us picture this organization as a fluid, dynamic, expanding one, spinning out into mainstream society."

The dynamics exhibited by SPINs appear to exist in these social movement groups as well as in various terrorist, criminal and fundamentalist networks around the world [11].

However, unlike many protest movements, terrorist and criminal networks often wish to remain covert. The need for security dictates that terrorist organizations must be structured in a way that minimizes damage to the organization from arrest or removal of one or more members [13]. This damage may be direct (making key expertise, knowledge or resources inaccessible for the organization) or indirect (exposing other members of organization during interrogations). There are several factors that allow a terrorist organization to remain covert, including:

- Strong religious (in case of Islamic groups) or ideological (in case of Sendero Luminoso and other South American guerilla groups) views that allow members to form extremely strong bonds within a cell.
- Physical proximity among cell members, often to the extent of sharing living quarters, working and training together.
- Lack of rosters on who is in which cell.
- Cell members being given little knowledge of the organizational structure and the size of the organization.
- Little inter-cell message traffic.
- Information about tasks issued on a need-to-know basis, so very few people within the organization know about the operational plans in their entirety.
- Cells are often formed on the basis of familial or tribal ties, or strong interpersonal ties forged in training.

However, a need-to-know information policy can be counterproductive when an organization needs to complete a task that is larger than any one cell. Further, such policies tend to lead to duplication of effort and reduce the ability of one cell to learn from another. To fix these inefficiencies, terrorist organizations have been known to employ "sleeper links" - where a small number of members of each cell have non-operational ties (such as family ties, ties emerging from common training, etc.) to members of other cells [14]. These links are rarely activated, and are used mainly for coordinating actions of multiple cells in preparation for a larger operation.

On the one hand, in part to remain covert, Al Qaeda has structured itself as a leaderless design characterized its organic structure, horizontal coordination, and distributed decision making. However, the need to maintain a strong ideological foundation and resolve coordination issues has led to the need for strong leadership. One apparent solution has been to have multiple leaders diffused throughout the network and engaged in coordinating activities, without central control or a hierarchy among the cells. Whether the leaders are themselves hierarchically organized, even though the cells are not, is less clear.

Under constant pressure from various world governments, terrorist organizations have evolved a structure that appears to be resilient to attacks. However, information on these terrorist organizations, their membership, the connections among the members, and so on is, at best, incomplete. Available information is often obtained during post factum investigations of terrorist acts, and may offer little insight into the "main body" of the organization or the way in which it is evolving.

Cellular Networks and Terrorist Organizations

Substantial intelligence effort is needed to piece together the massive amount of often misleading information - both post factum and "logs" of activity - to generate a picture of the entire organization. Nevertheless, the picture that is emerging suggests that terrorist organizations are organized at the operational level as *cellular networks* rather than as hierarchies [15].

Cellular networks are different from traditional organizational forms in that they replace a hierarchical structure and chain of command with sets of quasi-independent cells, distributed command, and rapid ability to build larger cells from sub-cells as the task or situation demands. In these networks, the cells are often small, only marginally connected to each other, distributed geographically, and may take on entirely different tasks.

Each cell is functionally self-sufficient, and is capable of executing a task independently. Cells are loosely interconnected with each other, for purposes of exchanging information and resources. However, the information is usually distributed on a need-to-know basis and new cell members rarely have the same exact skills as current members. This essentially makes each individual cell expendable. The removal of a cell generally does not inflict permanent damage on the overall organization or convey significant information about other cells. Essentially, the cellular network appears to evolve fluidly in response to anti-terrorist activity.

The fact that covert networks are often built from self-similar and somewhat self-sufficient cells leads to a hypothesis that cells throughout the network contain structurally equivalent [16] and essential roles, such as ideological or charismatic leaders, strategic leaders, resource concentrators, and specialized experts as needed given the modus operandi of the cell or its environment.

Given this hypothesis, we can further reason that operations of a particular cell will be affected in a negative way by removal of an individual filling one of these roles. Using this as a base for further exploration, we venture to show in this paper that cellular networks indeed contain vitally important and structurally equivalent roles, which can be detected through the use of dynamic social network analysis on the organizational MetaMatrix.

Profile of a Covert Network

From a combination of pressures to operate efficiently and pressures to remain covert, an organizational network emerges that combines massive redundancy with secrecy, separating into densely connected cells that are sparsely interconnected with each other through the leaders. No

clear hierarchy emerges from observation of these networks, other than a definite role of a cell leader, who is often the only contact that the cell has with the outside world.

The best profile of the structure of terrorist networks, based on publicly available data, is the following [7]:

- The network consists of cells with very low interconnection between cells.
- Internally, the cells exhibit high degree of connectedness and all-to-all communication patterns.
- There is a very low probability of a tie occurring by chance (0.007).
- The probability of triad closure (link from x to y being more likely if both x and y are linked to third party z) is 0.181.
- Senior members of each of the cells are often also parts of other cells and interact with other senior members on the network.
- Cell leaders are more knowledgeable than other members.
- Cell members have distributed knowledge.
- Cells use information technologies and electronic communication.

Destabilization of Covert Networks

The most common class of attacks against covert networks is comprised of strategies aimed at isolating or incapacitating a particular actor or leader, or an attack on the networks' infrastructure, training or weapons facilities.

For an attack on the infrastructure of a terrorist group to be successful, it has to be carefully targeted to not simply dismantle the command-and-control infrastructure, but to fully disconnect cells from crucial operational information or resource flow. Such precision targeting requires knowledge of the organizational structure of the network (including redundancies and latent links), the task structure (i.e. resource and knowledge requirements) and resource and knowledge distributions within the organization.

Distributed cellular nature of covert networks means that targeted actor attacks require such precision targeting. Targeted actor attacks involve the isolation of actors, where isolation can mean disabling of communications from/to or discrediting actors, incarceration or assassination.

Social Network Analysis and Network Destabilization

Traditional social network analysis (SNA) techniques have focused on analysis of communication networks between individuals. However, most SNA studies have been conducted on single-mode networks (i.e. relationships between people) with binary data (i.e. presence or absence of a connection). Also, most studies have been concerned with analysis of a single network. A further complication is that traditional social network measures are not designed for time-series analysis of dynamic networks. While static analysis may be adequate for slow-changing interpersonal networks, covert networks are characterized by their fluidity and

dynamism. Thus, analysis of covert networks needs to be approached from a dynamic perspective, tracking change inside the network as well as its static parameters.

Nevertheless, from an organizational perspective, it is important to look beyond social networks. Krackhardt and Carley [17] proposed concentrating knowledge about an organization in a format that could be analyzed using standard network methods, called the MetaMatrix. The MetaMatrix analysis represents organizations as evolving networks in which the nodes in the social network are actively engaged in realistically specified tasks. This conceptualization made it possible to link performance to social networks and ask, at a concrete level, how changes in the social network could effect changes in performance. Carley [18] [19] generalized this approach and extended the perspective into the realm of knowledge networks, enabling the researcher to ask how changes in the social network could effect changes in the distribution of information and the resultant impact of knowledge disruption strategies on organizational performance. By taking an information processing perspective, we are explicitly linking knowledge management and social networks [1] and enabling network evolution through learning mechanisms. From a conceptual and data perspective, this means that we examine the co-evolution of all networks in the MetaMatrix as described in table 1. Moreover, we explicitly focus on the fact that the organization, and so these underlying networks, evolve.

Table 1: Metamatrix of Organizational Knowledge

	People	Knowledge	Tasks
People	Structural knowledge: Command and control structures, information pathways and relationships between organization members.	Knowledge Distribution: Who has access to what knowledge within the organization.	Task Assignment: Who does which tasks within the target organization.
Knowledge		Knowledge Precedence: Which types of skills go together.	Skill Requirements: Which skills are needed to accomplish a particular task.
Task			Task Precedence:

			On a tactical level, the sequencing and precedence of tasks that the target organization can accomplish.
--	--	--	--

A number of social networks metrics have been proposed for identifying the key actors who should be targeted in order to destabilize covert networks. Such metrics include, but are not limited to, those focused on centrality, random attacks, and from a more dynamic network perspective, cognitive demand [7].

Identifying an actor as key, using one of these metrics and then isolating that actor is a destabilization strategy. We now consider several such strategies.

The *centrality approach*, consisting of measuring the centrality [20] of each node in the network, then selecting a small number of most central nodes as targets for further action, is an intuitive approach to finding a core group of leaders within a terrorist network.

However, it is known from available intelligence that terrorist networks function in tightly connected cells and maintain only loose connections with the rest of the organization. Therefore, a search for highly central individuals is more likely to turn up a large number of agents that do not constitute the leadership circle, but are members of a densely connected cell. Moreover, as Borgatti [21] stated, none of the centrality metrics is guaranteed to disconnect the network into discreet components.

Bienenstock and Bonacich [22] have conducted a simulation study on vulnerability of networks to random and strategic attacks. The study suggests that as average connectedness of each individual node rises and high betweenness nodes are methodically attacked, the impact on overall performance of the network is minimal. However, if neighborhoods (nodes connected to a high-centrality node) are attacked along with the node, the opposite is true.

The implication of that result is that the cells of covert networks that are connected by a few individuals with high betweenness are very vulnerable to discovery of these individuals.

Johnson et al. [23] show in their study of Antarctic winter crews that in order for an organization to exhibit high morale and operate efficiently, the positions of formal and informal (charismatic) leader of the network have to be occupied by the same person, and this person needs to be highly embedded within the network. The results of this study suggest that perhaps the structural position of a gatekeeper is not important to the functioning of an isolated cell. However if two

cells of the organization are to function in concert, the best position for the charismatic leader is in a gatekeeper role.

The *cognitive load* approach described by Carley [6] combines static measures of centrality with dynamic measures of information flow, task performance and resource distribution. These measures are based on the meta-matrix knowledge about the organization and have been shown to accurately detect emergent leaders. Consequently, cognitive load metrics can potentially be useful for detecting key members of terrorist networks.

Based on the foregoing review of the literature we have identified a suite of destabilization strategies. Each strategy identifies actor criticality in a different way. All strategies rely on data in one or more cells in the meta-matrix. The identified strategies are:

- **Highest degree centrality:** Isolate one agent from the covert network that has the highest degree centrality [20].
- **Highest betweenness centrality:** Isolate one agent from the covert network that has the highest betweenness centrality [20].
- **Highest cognitive load:** Isolate one agent from the covert network that has the highest cognitive load [6], where cognitive load is computed as a linear combination of:
 1. Number of people person i interacts with / total number of people in the group;
 2. Number of subtasks person i is assigned to / total number of subtasks;
 3. Sum of number of people who do the same tasks person i does / (total number of tasks * total number of people);
 4. Sum of negotiation needs person i needs to do for each task / total possible negotiations, where a negotiation corresponds to the amount of information or resources that an agent needs to complete a subtask that it is assigned to, but doesn't have (and thus has to obtain from another agent by negotiation).
- **Highest task accuracy:** Isolate the best performing agent in the organization. This corresponds to standard police practice of arresting agents implicated in commission of a terrorist act.
- **Amount of unique knowledge:** Isolate the agent that has the highest expertise.

When a destabilization strategy is applied, an actor is identified and isolated. This results in one or more changes in the underlying networks in the meta-matrix and possibly a cascade of future changes [24]. Since the overall network is a complex adaptive system there is no guarantee that such cascades will destabilize the overall network, particularly in the long run. Thus, an examination of these destabilization strategies needs to be done in a dynamic context.

NetWatch: A Multi-Agent Network Model of Covert Network Surveillance and Destabilization

NetWatch is a multi-agent network model for examining the destabilization of covert networks under varying levels and types of surveillance. Computational models, particularly, multi-agent network models, are a valuable tool for studying complex adaptive systems like organizations in general [14] [25] and covert networks in particular [7].

In multi-agent models, social behavior grows out of the ongoing interactions among, and activities of, the intelligent adaptive agents within the system. From the meta-matrix perspective, actions of each agent or actor are constrained and enabled not just by the activities of other agents but by what resources or knowledge they have, what tasks they are doing, the order in which tasks need to be done, the structure of communication and authority, and so on. Further, the agents are intelligent, adaptive and computational information processing systems.

The goals of NetWatch are to:

- Simulate the communication patterns, information and resource flows in a dynamic covert cellular network;
- Model the process of gathering signal intelligence on a cellular network and evaluate a variety of heuristics for intelligence gathering;
- Model and evaluate strategies for destabilizing a covert network based on intelligence obtained;
- Model reactions of a covert network to these destabilization strategies.

Agents in NetWatch

The Multi-Agent Network paradigm is based upon the following postulates:

- The simulation consists of agents.
- Agents are independent, autonomous entities endowed with some intelligence.
- Agents are cognitively limited.
- Agents can learn knowledge about the world and referential knowledge about other agents, with a limited learning capacity.
- Agents can forget.
- Agents communicate asynchronously and deal with asynchronicity (i.e. deadlocks, delays, etc) in an autonomous manner.
- Agents do not have accurate information about the world.
- Agents do not have accurate information about other agents.
- Unless required by the simulation domain, there is no central mediating entity to resolve the conflicts.
- Unless required by the simulation domain, the agents do not use predefined geometrical locations or neighborhoods.
- Agent communications are governed by a number of common protocols, including these for knowledge and resource exchange, task execution and reporting of status or results.

The social and cognitive underpinnings of the actors and the network in which they operate are based upon the CONSTRUCT model of the co-evolution of social and knowledge networks [18] [24]. The agents in the model perform a classification task that is information-

intensive (i.e. requires a large amount of knowledge to complete without guessing). In the beginning of the simulation, agents are endowed with relatively little knowledge and must engage in learning behaviors in order to increase their task performance. Agents learn by interaction: trading facts with other agents or asking direct questions in hope of getting an accurate answer. Agents also forget little-used facts.

In keeping with the research in cognitive science, the agents representing humans are both cognitively and socially constrained [26] [27] [28] [29]. Thus, their decision-making ability, actions, and performance depend on their knowledge, structural position, procedures and abilities to manage and traverse these networks.

Unlike Construct agents, the NetWatch agents are implemented as non-deterministic finite automata, with states of the automaton representing low-level behaviors and transitions governing the way the agent switches between them. Some transitions are deterministic, others rely on probabilistic equations.

Low-level behaviors include chatter, knowledge seeking, resource seeking, task execution and information reporting.

Chatter is the simplest of the low-level behaviors. It can be thought of as non-goal-directed socializing, where some information is exchanged but it may or may not be relevant to the task the agent is engaged in. Partners for chatter interaction are randomly picked from the agent's ego network (peer group). Chatter uses the *Knowledge Exchange Protocol* (see section 4.3) but its messages have the lowest priority.

Knowledge and Resource seeking behaviors use the same protocol as chatter, however assign a higher priority to the messages. Communication partners are determined by estimating the probability of a successful interaction, informed by the MetaMatrix representation of the agent's ego network. Processes that govern selection of partners are described in section 4.2.

Task execution is described in detail in section 4.4. It is governed by a simple challenge-response protocol that is executed over one time period. Task messages have the highest priority in the system and will preempt both knowledge exchange and chatter messages.

It is important to note that due to asynchronous execution of agents and multi-tiered message priorities, it is possible that some interactions will never complete or will complete after a significant delay. Each agent stores incoming unprocessed messages in a queue sorted by message priority. Thus, if an agent is overwhelmed with tasks or goal-oriented information exchange, most chatter requests will never be processed.

To prevent deadlocks, each of the messages is time-stamped at the time of sending, and interactions are set to time out after a fixed number of time periods. Also, agents are capable of handling multiple interactions at the same time, with task preempting based on priority of incoming messages. For example, if an agent was in the middle of a chatter interaction when a resource request or a task request arrived, the chatter will not be resumed until higher-priority interactions have been finished.

Formal and Informal Networks in NetWatch

In NetWatch, the formal structure of the organization is specified as a directed weighted graph that specifies the communication channels that are open as well as their throughput or cost of communication. The directed nature of the graph allows one to specify one-way relationships and chain-of-command relationships.

The beliefs about the informal structure are individual to every agent, and also consist of a weighted directed graph. However, when an agent joins a network, its informal relationship graph is empty, and it must learn about the informal network before it can be used for communication.

In NetWatch, the agents' interactions are governed by the formal structure of the organization, and agents' beliefs about the informal structure.

The agents communicate solely on the basis of networks that they belong to. Each of the networks is represented as a directed graph structure representing probability of communication or social proximity:

$$Net_i = A, P$$

$$A = a_i : \text{Set of agents}$$

$$P = p_{ij} : \forall a_i, a_j \in A p_{i,j} = \text{probability of communication}$$

The agents do not have access to full information about the network, but rather every agent

$a_k \in A$ can only access a probability vector $P_i = p_{ki}$ where p_{ki} is a probability of agent a_k communicating with all agents $a_i \in A$. This means that each agent may only know who it may interact with or is close to - but does not have access to interaction patterns of any other agents.

Each agent also possesses a belief matrix that it uses to store any information it learns about interrelationships of other agents within the network. However, this information is far from complete and is often inaccurate.

The directionality of the network also means that the communication may be asymmetric - thus allowing full representation of command networks as well as (more symmetric) friendship networks.

For example, in NetWatch, a cellular organization like Al Qaeda can be represented as a cellular network structure. The formal network consists of small densely connected cells that maintain a small number of connections to other cells. The ties in experimental networks are generated from a profile of a cellular organization, such as one described in section [2.2](#). The profile contains the following information:

- **For communication networks:** Mean and standard deviation of size of cells, connection density inside cells and outside cells, density of one-way links and probability of triad closure;
- **For knowledge networks:** Amount of common knowledge (doctrine), distributed knowledge (group member specialties) and specialized (expert) knowledge;
- **For resource networks:** Amount of resources needed to accomplish tasks, amount of common and distributed resources;
- **For task networks:** Branching factor and depth of the task precedence network.

The profile is used to create a probability distribution for each edge within a network, thus generating a space of random networks that all conform to the original profile. A number of sample organizations is then drawn from that space and run through the simulation, and mean and standard deviation of each of the resulting variables are taken.

Processes Governing Communication

Each of the agents in NetWatch maintains a perception of its surroundings, via the notion of MetaMatrix [1](#). The perceptive MetaMatrix consists of the agent's ego network (agents that it is directly connected with), agent's own knowledge, resources and task assignments, and is augmented by the agent's perception of other agents' ego networks, knowledge, resources and task assignments.

However, an agent may only learn of other agents outside its ego network via interaction with agents that are in its ego network - and therefore any agent's perception of other agents' networks or knowledge is generally inaccurate. Moreover, it has been shown [\[30\]](#) that knowledge of people outside a person's ego network decreases exponentially as graph distance between the actors increases.

In the context of a cellular organization, this translates to agent's initial knowledge of its network including its cell (because of dense communication patterns inside cells) and a small number of agents outside the cell with whom cell members regularly communicate. Agents may later acquire further knowledge of the organization through interactions.

The choice of communication partner at every time period is based on two factors: *Social proximity* of the agents and their *motivation to communicate*. Social proximity is defined as closeness of a relationship between two agents, scaled between 0 and 1 where 0 means "no relationship" and 1 is "very close relationship."

Motivation to communicate is computed on the basis of *homophily* (relative similarity) and relative expertise.

We define homophily to be based on a measure of relative similarity between agent i and agent j : the amount of knowledge that i and j have in common divided by the amount i shares with all other agents, or

$$RS_{i,j} = \frac{\sum_{k=0}^K (S_{ik} * S_{jk})}{\sum_{j=0}^J \sum_{k=0}^K (S_{ik} * S_{jk}) : wq}$$

where S_{ik} is 1 if agent i knows fact k and 0 otherwise.

Relative expertise is defined as RE_{ij} = how much agent i thinks j knows that i does not know divided by how much i thinks all others know that i does not know, or

$$RE_{ij} = \frac{\sum_{k=0}^K ((1 - S_{ik}) * S_{jk})}{\sum_{j=0}^J \sum_{k=0}^K ((1 - S_{ik}) * S_{jk})}$$

In both cases, agents operate on their beliefs about what the other agents know. Thus, their predictions of relative expertise or similarity can be inaccurate. However, as interaction progresses and agents learn more and more about each other, they learn an increasingly complete picture of their world.

Processes Governing Knowledge Exchange

In a multi-agent network, the agents do not have perfect knowledge about the world. The only way to obtain information about the world is via interaction with other agents - either through direct query or through information exchange.

Tracing back to its roots with Construct [18] model, the NetWatch model is based upon the concept of knowledge, knowledge manipulation and learning. In NetWatch, each agent's knowledge is represented by a bit string. A value of 1 in the position n means that the agent knows fact n and the value of 0 means that it does not.

At the start of the simulation, the agents are endowed with some initial knowledge (typically within 2%-10% range). This allows for only a minimally acceptable performance (and thus a very low utility), giving agents an incentive to communicate with other agents and attempt to gain more information.

To learn new facts, the agents execute the **Construct Knowledge Exchange Protocol**. For ease of description, we shall refer to the parties in knowledge exchange as Alice (agent A) and Bob (agent B). Note that Alice and Bob can be any two agents $a_i, a_j \in A$.

1. **Determine who to communicate with:** Alice does this by evaluating Relative Similarity (Eqn. 4.2) or Relative Expertise (Eqn. 4.2) of every agent accessible through the Alice's social network (i.e. $p_{A_i} > 0 \quad \forall i, a_i \in A$ for $a_i \in A$). After the probability of communication for each of the agents is computed, Alice throws a dice that reflects the computed probability vector and determines an agent to communicate with, or Bob.
2. **Determine what to communicate:** This is done by weighing information seeking vs. similarity-driven communication. If Alice is in information seeking mode, it chooses at random a part of the knowledge string that is not known (i.e. bits are set to "false") and queries the agent chosen in step 1. In similarity-based communication, Alice chooses a part of the known knowledge string and sends it to its counterpart.
3. **Determine proper response:** On receipt of a query, Bob determines if it should answer it by checking whether the sender of the query is a part of its network and whether it has the knowledge in question - and, if all is good, sends a reply. If Bob does not know the facts requested, it checks its internal belief matrix and may respond to Alice with a name of another agent (Clare) that may be better suited to answer Alice's question. In this case, the agents exchange referential data.

On receipt of knowledge, Bob determines if the knowledge is useful (i.e. whether it is already known) and whether it came from one of the agents in its network (and thus can be trusted). If all is good, the agent will choose some knowledge from its knowledge base and send it back.

4. **Update internal knowledge base:** On receipt of the reply, Alice determines the usefulness of the reply and uses that to update its internal knowledge of Bob ("Bob knows fact π ") as well as its knowledge base ("I now also know fact π ").

If Alice receives referential data, it uses that to update both its knowledge of Bob ("Bob does not know fact π " and "Bob knows Clare") and its knowledge of Clare ("Clare may know fact π ").

This may be followed by a query to Clare - which may or may not be honored.

Note that Clare may not have been originally a part of Alice's network - but now, through Bob, Alice has learned about her existence. Thus, agents within the organization use referential data about each other to form an informal network.

Due to asynchronicity of communication, the agents may not be able to conduct a knowledge exchange transaction in one time period. It is also possible that agents may be too busy to be able to respond to a query, and may either delay or terminate the transaction. The knowledge exchange protocol, however, provides for a robust deadlock resolution, allowing agents to detect a transaction that is deadlocked, terminate it and start anew, finding a different party to communicate with.

Tasks and Organizational Performance in NetWatch

The simulation paradigm is task-independent. The task is merely defined as a function that maps a problem vector and agent's knowledge and resource vectors onto a result vector.

In NetWatch, we measure agent performance as accuracy in performance of a ternary classification task. The classification task is represented by a vector of binary values. An agent can only access bits in the task vector that correspond to non-zero values in agent's knowledge vector. The task is then decided by a "majority rule."

An agent's decision accuracy is computed by taking a series of classification tasks and comparing agent's decisions to "true decisions" - computed by applying a majority rule to the tasks assuming "perfect knowledge" or access to all bits of the task string. Task performance is measured as a percentage of correctly decided tasks.

The agents in NetWatch are engaged in a knowledge-intensive ternary classification task, identifying targets (represented as bit strings) and determining whether they should be attacked, treated as friendly or ignored. Each individual agent's performance on each particular task depends on its level of knowledge: an agent can see specific parts of the target bit string only if it has appropriate facts in its knowledge base.

While appearing simplistic, performance in classification tasks have been shown [31] to correspond, in aggregate, to organizational performance in real-world cases, thus making classification tasks a suitable substitute for more complex tasks for purposes of simulation modelling.

Simulation Design

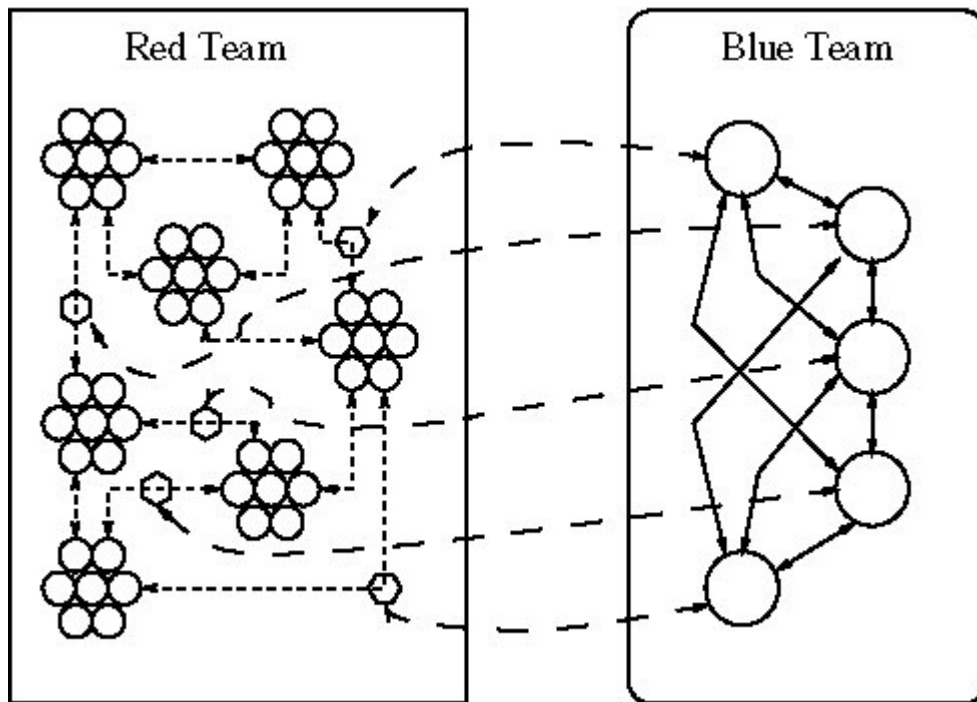


Figure 1: NetWatch Simulation Design

The simulation consists of several networks of agents (see figure 1): the **Red Team**, representing the covert network of a terrorist organization, the **Blue Team** representing the anti-terrorist or law enforcement forces, and a set of instrumentation agents that observe and document the behavior of other agents for later retrieval and processing.

Red Team




Figure 2: Red Team: A Cellular Covert Network (click the figure to zoom)

The Red Team, or the Covert Network consists of a set of small fully-interconnected cells of agents with little interconnect between them, mimicking the organizational structure of a terrorist organization described in [section 2](#).

The plot on [figure 2](#) shows a covert network generated using parameters in [section 2](#) and bears a striking resemblance to the structure of terrorist networks as shown by Valdis Krebs [[14](#)].

In the same time, driven by the incentive to increase task performance accuracy, the agents are engaged in learning behavior using the protocol described in [section 4.3](#).

If left uncontested, the members of the Red Team proceed to learn all accessible facts, thus increasing the performance of the organization as a whole (see [figure 3](#)).




Figure 3: Knowledge Diffusion Through a Cellular Network

Blue Team

The Blue Team is an Anti-Terrorism organization consisting of a small number of fully interconnected law enforcement agents.

The goals of the Blue Team are:

- Learn as much as possible about the structure, task assignments and knowledge distribution of the Red Team.
- Use the knowledge obtained to remove or isolate Red Team members, aiming to impair Red Team's performance as much as possible.

The Blue Team has very little access to the actual information about the Red Team. Its only source of information is a set of wiretaps placed on the communication network of the Red Team.

Wiretaps

A wiretap in NetWatch is an agent that selectively intercepts messages from the message stream and routes them to one or more Blue Team members, functioning on a subscription model. When a Blue Team agent receives a message through a wiretap, it records the origin and destination of the message. If a wiretap were capable of capturing all communication, this would allow the Blue Team agent to create a full and accurate picture of the Red Team network.

However, the wiretap agent (as well as real wiretaps) is not capable of capturing all relevant messages. Based on empirical information about the use of wiretaps, the wiretap efficiency is set between 5% and 25%.

Given the expense and political difficulty of wide use of wiretaps, law enforcement agencies utilize a number of strategies to maximize the yield of useful information from the minimum number of wiretaps.

In this paper, we study three different wiretap strategies:

- **Random:** Any message in the message stream has an equal probability of being intercepted. This is similar to the Echelon signal intelligence gathering.
- **Targeted:** Only messages addressed to or from a specific person can be intercepted, similar to the way standard police wiretaps are done.
- **Centrality-Based Roving:** Every 50 time periods, the most central person is picked for surveillance.

As the Blue Team agents receive messages from the wiretap agents, they use the message *To:* and *From:* fields to build a representation of the organizational network of the Red Team, or the *Learned Network*. Other information collected includes data on task performance and knowledge as communicated by the agents on the Red Team.

Network Destabilization Tactics

In a subset of the experiments, the Blue Team agents not only collect information about the Red Team but also attempt to influence the performance of the Red Team by finding vulnerabilities in the covert network and attacking them by isolating or terminating agents within the covert network.

In this paper, we test a number of strategies for finding key individuals within the covert network, including

- **Random:** A base-line strategy; isolate one random individual from the network.
- **Highest degree centrality:** Isolate one agent from the covert network that, by the data of the Blue Team, has the highest degree centrality.
- **Highest betweenness centrality:** Isolate one agent from the covert network that, by the data of the Blue Team, has the highest betweenness centrality.
- **Highest cognitive load:** Isolate one agent from the covert network that, by the data of the Blue Team, has the highest cognitive load, where cognitive load is computed as a linear combination of:
 1. Number of people person i interacts with / total number of people in the group;
 2. Number of tasks person i is assigned to / total number of tasks;
 3. Sum of number of people who do the same tasks person i does / (total number of tasks * total number of people);
 4. Sum of negotiation needs person i needs to do for each task / total possible negotiations.

- **Highest task accuracy:** Isolate the best performing agent in the organization.
- **Amount of knowledge:** Isolate the agent that has the highest expertise.

Performance Metrics

The Red Team performance metric is based on the average accuracy of the classification task (as described in section 4.6, or, for any given task t ,

$$Accuracy_t = \frac{\sum_{a_i \in A} |Result_{t,a_i} - Result_t|}{\bar{A}}$$

where \bar{A} is the number of agents in the system, $Result_{t,a_i}$ is the decision made by agent a_i in task t and takes values of 2 for "Hostile Target", 1 for "Neutral Entity" and 0 for "Friendly Entity" and $Result_t$ is the correct decision for task t , taking on the same values.

The performance of the Blue Team is subject to two measures: a measure of intelligence gathering accuracy and a measure of isolation strategy effectiveness.

The effectiveness of intelligence gathering is measured as a *Hamming Distance* between the knowledge of the Red Team as collected by the Blue Team (a.k.a. the *Learned Network LN*) and the actual network of the Red Team AN , computed as

$$HD = \sum_{i=0}^{\bar{A}} \sum_{j=0}^{\bar{A}} |LN_{i,j} - AN_{i,j}|$$

or sum of absolute values of differences between the two networks represented as binary adjacency matrices (where 1 signifies presence of a tie between two agents and 0 signifies absence thereof).

Effectiveness of wiretapping strategy can thus be measured as the *first derivative* of the Hamming Distance, signifying the speed of learning.

The effectiveness of isolation strategy is measured as the difference between baseline performance of the Red Team (i.e. without any action by the Blue Team) and performance of the Red Team in presence of an anti-terrorist task force of the Blue Team.

One must note, however, that the networks in question are dynamic and thus there can not be an absolute and static performance metric for any of the teams outside the time series data.

Virtual Experiments

We have conducted two virtual experiments with NetWatch simulation. For both experiments, the simulation was configured as follows:

Red Team:

- 50 agents organized into cells
- Mean cell size is 6
- Connection probability within cell: 95%
- Connection probability outside cell: 5%
- Initial knowledge: 10%
- One agent in each cell (a cell leader) is connected to at least one other cell.

Blue Team:

- 5 agents, fully interconnected.
- employ wiretaps to collect data on the Red Team
- attempt to create a consensus view of the Red Team via the MetaMatrix representation.
- employ one of the destabilization strategies outlined above to eliminate actors in the Red Team
- employ conduct impact assessment to estimate the damage inflicted on the Red Team.

Virtual Experiment I: Effects of Wiretap Strategy on Network Learning

The goal of this experiment is to learn about the effectiveness of wiretapping strategies. In a three-cell design, the performance of the Blue Team (measured by the Hamming Distance (see [4.6.4](#)) is evaluated at every time period.

The experimental cells correspond to the different wiretapping strategies available in NetWatch:

- **Random:** Any message in the message stream has an equal probability of being intercepted. This is similar to the Echelon signal intelligence gathering.

- **Targeted:** Only messages addressed to or from a specific person can be intercepted, similar to the way standard police wiretaps are done.
- **Centrality-Based Roving:** Every 50 time periods, the most central person is picked for surveillance. This strategy is modeled after wiretaps allowed under the PATRIOT Act in the investigation of terrorist organizations.

Table 2: Virtual Experiment I: Effects of Wiretap Strategy on Network Learning

Wiretap Types	RANDOM; TARGETED; CENTRALITY ROVING
Red Team	50 agents organized into cells; mean cell size=6
Blue Team	5 agents, fully interconnected

Results

Results of the first virtual experiment are plotted in figure [4](#).

RANDOM

Random wiretap strategy provided a consistent, almost linear increase in the structural knowledge acquired by Blue Team. An important caveat of this strategy is that the simulation detailed in this paper only includes agents belonging to one of two teams, and no "innocent bystanders" - thus resulting in a fairly high signal-to-noise ratio in Echelon-style signal intelligence.

In the real world the signal-to-noise ratio of such strategy would be extremely low, resulting in much lower performance of the system.

TARGETED

Targeted wiretap strategy is similar in way of functioning to wiretaps employed in police use. The use of a targeted wiretap allows the Blue Team to get nearly complete knowledge of a small segment of the network. However, due to the cellular structure of covert networks, this strategy has a limited usefulness due to the fact that a large number of operatives in the covert network will not communicate over a monitored channel and thus will not appear "on the radar" of law enforcement.

It is a possibility that extensive use of such narrowly targeted investigative techniques resulted in the rise of cellular organizations as a counter-measure to police activity. Also, an initial choice of the monitoring target may result in drastically different performance: if a fringe cell of the organization is picked as a target, the main body of the organization may not ever be discovered.

ROVING

Roving wiretap strategy has the highest performance potential. As figure 5 shows, during the course of the experiment, the Blue Team fully discovers five cells comprising the main body of the organization and one fringe cell. The learning is accomplished in steps, focusing on tightly interconnected agents within a cell, and switching to other cells as the Blue Team learns of their existence.

The main caveat of this strategy is its cost: if it was employed in the real world, the discovery process would require six separate court hearings. Also, the performance of this strategy may be jeopardized if the Red Team employs larger cells, which would prevent the Blue Team from fully discovering a cell before it is time to switch. However, this problem can be addressed by employing a switch condition based on rate of learning (a target switch occurs when the cell under monitoring appears fully discovered).

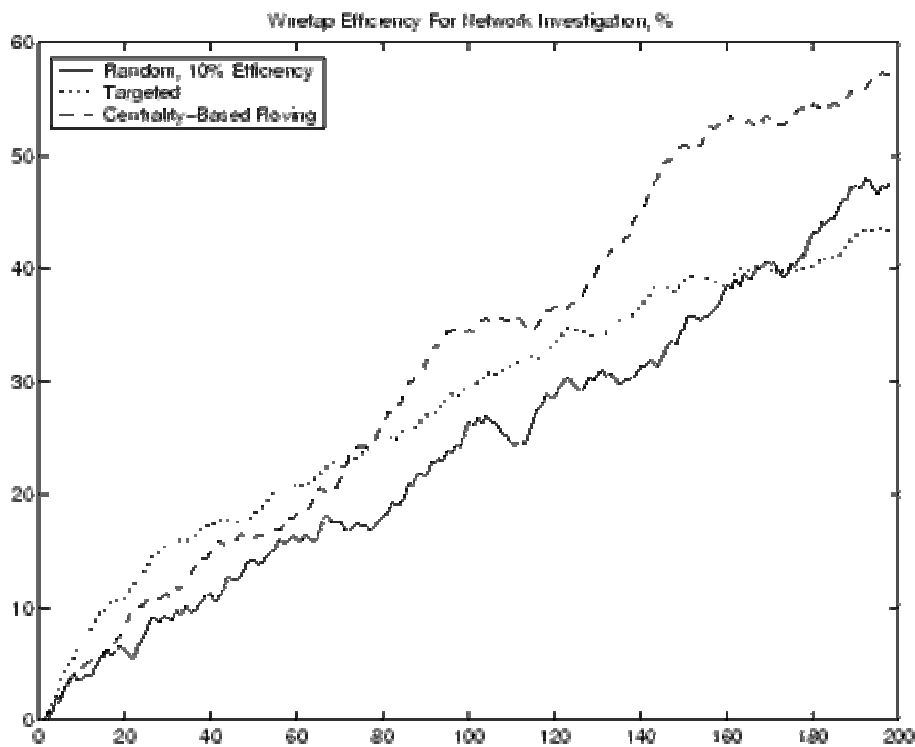


Figure 4: Network Discovery by Wiretap Strategy




Figure 5: Cell Discovery with Roving Wiretaps (click the figure to zoom)


Virtual Experiment II: Effects of Network Learning on Network Destabilization

The goal of this experiment is to find out the interdependence between intelligence gathering techniques and effectiveness of network destabilization. The experiment is built as a 3x6+1 design, combining the wiretapping strategies (see [4.6.2](#)) and network disruption strategies (see [4.6.3](#)). We also ran a baseline case with the Blue Team absent.

Table 3: Virtual Experiment II: Effects of Network Learning on Network Destabilization

	RANDOM	TARGETED	ROVING
NONE	N/A	N/A	N/A
Random	RAND/RAND	RAND/TARGET	RAND/ROVING
Degree Centrality	DCENT/RAND	DCENT/TARGET	DCENT/ROVING
Betweenness Centrality	BCENT/RAND	BCENT/TARGET	BCENT/ROVING
Cognitive Load	CL/RAND	CL/TARGET	CL/ROVING
Task Accuracy	TA/RAND	TA/TARGET	TA/ROVING
Knowledge	K/RAND	K/TARGET	K/ROVING

Network Recovery After Isolation Operations

Figure 6: Recovery of a Cellular Network After Disconnection of a Gatekeeper Node (drag the mouse in the image window to rotate and examine; click "Animate" button to step through the demonstration) 

One of the most notable results of this experiment has been the discovery of the mechanism that the network uses to recover after removal of one of its key members. Figure [6](#) demonstrates this process on a small cellular network.

1. In the original configuration, the network consists of two fully interconnected cells and one gatekeeper agent. As the organization goes about its business, information flows from cell 1 (agents 30-35) to cell 2 (agents 37-42) through the gatekeeper agent 36. In the process of passing knowledge and requests, small amounts of referential data (such as "Agent X knows fact Y") is passed from one cell to another and stored by the agents. Relevancy or immediate usefulness of this information is low because of the fact that all needed information can be easily obtained from querying the gatekeeper agent.

2. Agent 36 is identified by the Blue Team as being important to the network because of its degree centrality, betweenness centrality and large amount of messages that it processes (cognitive load).
3. The Blue Team proceeds to remove Agent 36, disconnecting the two cells. Because of the cellular structure of the organization, this makes information transfer between cells impossible and the performance of the organization is greatly degraded.
4. As information becomes unavailable from the central source, agents use the referential data accumulated in previous transactions to attempt to find information in the other cell. For the connection from Agent X to Agent Y to succeed, both agents have to have knowledge of each other. However, referential data is asymmetric and thus most of these connection attempts fail.
5. One of the connection attempts (Agent 31 to Agent 39) succeeds, thus opening a single pathway between cells.
6. Referential information about Agent 39 spreads throughout Cell 1, and more connections between agents in Cell 1 and Agent 39 are created. Within a short period of time, Agent 39 emerges as the new gatekeeper between the two cells.

Information flows easily between the two cells and organizational performance is restored to levels similar to these before the removal of Agent 36.

A priority in research on destabilization of covert networks has been finding key individuals - the removal of which will separate a cellular network into subparts. However, the recovery process we demonstrate illustrates that even if the Blue Team achieves separation of the covert network into disconnected cells, the network will use its latent resources and quickly recover from damage.

Thus, the goal of network destabilization techniques should be to cause permanent damage to the covert network and not allow it to recover from the attack.

This finding also prompts us to study not just the effectiveness of removal of certain individuals, but to look at the performance of these measures over periods of time.

Effects of Destabilization on Task Performance




Figure 7: Effectiveness of Destabilization Strategies on Organizational Performance

- **Random** case consisted of random removal strategy combined with one of the three wiretapping strategies. The baseline increases with accuracy of network discovery because of the fact that learning more of the network gives the random strategy better access to targets (including these of high importance). In the absence of false positives, the baseline strategies perform better than many complex heuristics.
- **Degree Centrality**-based removal of agents performs well only with targeted wiretapping. This is because targeted wiretapping allows the Blue Team to completely discover a cell and all of its neighborhood, and thus be very confident in removal of a most central person. However, the network does recover reasonably quickly from the hit, thus decreasing the effectiveness of this strategy.
- **Betweenness Centrality**-based removal performs well with roving wiretaps - because of the fact that they allow quick location of key gatekeepers in the network.
- **Cognitive Load**-based removal requires significant knowledge of the network, and thus only performs well under targeted and roving wiretap strategies.

- Quite surprisingly, removal of **Well-Performing** individuals did not result in significant damage to the terrorist network. Most likely this is caused by the structure of the tasks that the agents perform and equivalence of individuals in equivalent network positions.
- **Knowledge-based** removal proved to be the most effective strategy of the lot, and the only one to outperform the baseline strategies. The key to knowledge-based removal, though, is that its effectiveness hinges on speed of response more than the accuracy of structural knowledge. The experiments show that the highest efficiency of this strategy is achieved if Red Team experts are found and removed before they transmit much of their knowledge to the rest of the group, thus favoring quick decisive action to lengthy deliberation and heuristic learning that are characteristic of roving wiretaps.

Effects of Destabilization on Knowledge Diffusion




Figure 8: Effectiveness of Destabilization Strategies on Knowledge Diffusion

The **Knowledge-based** isolation strategy showed the most profound effect on diffusion of knowledge throughout the Red Team network (see figure 8) After an expert has been removed in timestep 50, the knowledge diffusion is sharply capped and does not grow further.

Other isolation strategies had a damping effect on the diffusion of knowledge, but did not result in permanent damage and merely slowed down the diffusion process. The most successful of these were based on **Betweenness Centrality** and **Cognitive Load**, mirroring the success of these strategies in capping the task performance of the organization.

Nodes high in betweenness were generally in the position of gatekeepers acting as liaisons between pendant cells and the main body of the network. In the real world, the pendant cells would likely be operatives or suicide squads preparing for immediate action. Thus, isolation of high betweenness individuals may play a role in prevention of terrorist acts, but is unlikely to permanently degrade the network's ability to prepare or execute other operations.

Cognitive load-based isolation strategy concentrates its attention on individuals with high dynamic characteristics, such as amount of information processed. However, the experts within the network engaged in dissemination of their knowledge to their close associates, as well as shedding tasks and as a result did not exhibit high cognitive demand. This behavior is similar to task-shedding behavior shown by Carley and Svoboda [32].

Discussion

Covert networks exhibit a number of qualities that hinder effective study and modeling. We find that multi-agent network models as described in this paper have a number of advantages over traditional methods in study of such networks.

Standard network analysis tools use graph-theoretic measures of node criticality to determine which nodes are most important in a covert network. However, Borgatti [21] showed that traditional network analysis methods such as centrality-based measures and cutpoint algorithms often fall short of the goal of separating the network into discrete components.

Optimization techniques such as the Key Player algorithm [21] have been proposed as an alternative to traditional network methods for finding critical nodes in the system. These algorithms can provide deterministic and optimal solutions to the problem of finding important nodes in cellular networks. However, the deterministic algorithms do not degrade gracefully when the data on the shape of the network is incomplete or inaccurate - which should be expected for covert networks.

Due to the dynamism of covert networks and their ability to self-heal using low-priority links and referential data (as illustrated above), it is also unclear that separation of a covert network into components is the desired outcome of an isolation event. In our virtual experiments, the covert network was able to recover quickly from isolation of individuals in a gatekeeper position, while removal of key experts resulted in permanent damage. Thus, destabilization strategies that only consider the connections between agents are not sufficient. One needs to take into account properties such as knowledge and resource distribution, as well as to analyze the dynamic processes that reshape the network as it is being acted upon.

This work also pushes the boundaries of multi-agent simulation. Traditional A-Life based multi-agent simulations have been effective at modelling many dynamic phenomena in organization theory. A-Life based simulations are built upon the concept of agents located on a grid of pre-specified shape and size. Interactions are based on concept of proximity, as defined by the agent neighborhood on the grid. Unfortunately, the size and shape of both the grid and the neighborhood have significant effects on agent behavior and are difficult to justify from the

theoretic point of view. In this paper, we have placed the agents in dynamically evolving networks rather than a grid. The notion of grid proximity is replaced with that of a graph proximity, which allows for multiple overlapping networks as well as irregular-shaped neighborhoods, thus allowing a much higher fidelity of simulated organizational structures.

Despite overcoming the limitations of traditional social network analysis and A-Life models, the NetWatch model has a number of limitations. First of all, we study behavior of a covert network of a limited size and clear boundaries. This does not take into account the fact that most covert networks actively recruit new members, as well as the fact that covert networks are often tightly integrated into the fabric of the society they operate in, which makes the boundaries of such networks difficult to determine ("fuzzy" [14]).

Addition of recruiting to the model also raises new important questions. How can we stem the growth of terrorist organizations? Is it possible to integrate spies into the covert network and use them to obtain human intelligence? Fuzziness of boundaries not only increases the cost and difficulty of surveillance, but also the cost of false positives, thus changing the heuristic criteria for target selection.

We also realize that a ternary classification task does not capture all of the variety of tasks done in terrorist organizations. This task was used in the initial iteration of the model to link our findings to prior models of organizational behavior (such as CONSTRUCT and ORGAHEAD). Future work might consider more complex tasks, such as attacks, operational support and resource channelling.

Lastly, in our model the anti-terrorist units ("Blue Team") are represented as a small cooperative network. In fact, there are a large number of agencies involved in anti-terrorism activity and intelligence gathering, and their network is neither fully connected nor fully cooperative because of legal and organizational boundaries.

While the NetWatch model in its current iteration is limited, it has produced a number of important findings. First of all, it has shown that finding the most central or structurally important individual in a covert network is not equivalent to success. It has also shown an emergent self-healing behavior of dynamic networks, which corresponds to anecdotal data on the effect of anti-terrorist actions.

The multi-agent network simulation paradigm allows us to build a very detailed model of the activity of the organization, from realistic modelling of the task to simulation of multiple overlapping networks. It also allows us to test the behaviors of two competing organizations (the Red Team and the Blue Team) in a dynamic co-evolutionary environment that is uniquely suited for simulation of large complex organizations.

Conclusion

In this paper, we have also shown that multi-agent network simulation is capable of producing high-fidelity models of complex organizational structures. Further, using this approach we have

detailed a number of strategies for obtaining information about covert networks and destabilizing them through selective isolation of individuals. Our model has shown that cellular networks do exhibit the property of containing structurally equivalent roles and individuals, and isolation of such individuals is a valid strategy for anti-terrorist operations.

However, while isolation of well-connected individuals may disconnect the covert network into separate components, the network exhibits an emergent healing behavior that negates the drop in performance precipitated by loss of a central agent, and reconnects the components within a short period of time. Nevertheless, isolation of individuals in other structurally equivalent roles, such as experts, was shown to cause the most amount of permanent damage to the cellular network.

The main implication of this finding is that in order to accurately map organizational networks (whether they are covert networks or corporate structures), one must not limit the data collection and evaluation to mapping pure social networks. Diversity of data and inclusion of knowledge, task and resource data enables significant increase in mapping performance.

The theory of structural equivalence, presented by Lorrain and White [16], argues that agents that occupy the same role in a network should be structurally equivalent to each other (i.e., exhibit a similar pattern of linkages with neighboring nodes), and thus be substitutable for one another. In this sense, nodes that are structurally equivalent can be used to "replace" each other.

Dynamic recovery of cellular networks presents a new twist by noting the impact of such replacement on the overall performance of the network. When an actor is isolated, the impact of that isolation will depend on whether or not there are other actors that are structurally equivalent and can come to take the place of the isolated actor.

Thus, destabilization strategies that locate actors that have alters who are approximately structurally equivalent are less effective because, in using nearly similar connections to reconnect groups, they are increasing their similarity with the isolated actor and preserving the cellular nature of the network. Social networks with sets of key actors who are nearly structurally equivalent in this "actor space" cannot be disrupted by removing only one of those actors. In the pure social network, experts or those high in cognitive load, are unlikely to appear as being structurally equivalent to each other. As such, there is no obvious structural role that can be drawn on to fill the communication gaps.

Bibliography

1

Z. Lin, K. M. Carley, *Designing Stress Resistant Organizations: Computational Theorizing and Crisis Applications* (Boston, MA: Kluwer, 2003).

2

P. Lawrence, J. Lorsch, "Differentiation and Integration in Complex Organizations," *Administrative Science Quarterly* 12 (1967): 1-47. Available in JSTOR: <http://www.jstor.org/view/00018392/di995417/99p0086p/0> [March 17, 2005].

3

H. Baligh, R. M. Burton, B. Obel, "Devising Expert Systems in Organization Theory: The Organizational Consultant." In Michael Masuch (ed.), *Organization, Management, and Expert Systems* (Berlin: Walter De Gruyter, 1990), 35-57.

H. Aldrich, *Organizations Evolving* (London: Sage Publications, 1999).

K. M. Carley, "Inhibiting Adaptation." In *Proceedings of the 2002 Command and Control Research and Technology Symposium*, conference held in Naval Postgraduate School, Monterey, CA (Vienna, VA: Evidence Based Research, 2002).

K. M. Carley, Y. Ren, "Tradeoffs Between Performance and Adaptability for C3i Architectures." In *Proceedings of the 2000 International Symposium on Command and Control Research and Technology Symposium*.

K. M. Carley, J.-S. Lee, D. Krackhardt, "Destabilizing Networks," *Connections* 24, 3 (2001): 31-34.

R. Gunaratna, *Inside Al Qaeda: Global Network of Terror* (New York: Columbia University Press, 2002).

R. Goolsby, "Combating Terrorist Networks: An Evolutionary Approach." In *Proceedings of the 8th International Command and Control Research and Technology Symposium*, conference held at National Defense War College, Washington, DC (Vienna, VA: Evidence Based Research, 2003). Available: http://www.dodccrp.org/events/2003/8th_ICCRTS/pdf/044.pdf [March 17, 2005].

International Crisis Group, "Indonesia Backgrounder: How The Jemaah Islamiyah Terrorist Network Operates," *Asia Report* 43 (December 11, 2002). Available: <http://www.icg.org/home/index.cfm?id=3011&l=1> [March 17, 2005].

D. Ronfeldt, J. Arquilla, "Networks, Netwars and the Fight for the Future," *First Monday* 6, 10. Available: http://www.firstmonday.org/issues/issue6_10/ronfeldt/ [March 17, 2005].

L. P. Gerlach, V. H.Hine, *People, Power, Change: Movements of Social Transformation* (Indianapolis: Bobbs-Merrill, 1970).

B. H. Erickson, "Secret Societies and Social Structure," *Social Forces* 60, 1 (1981): 188-210. Available in JSTOR: <http://links.jstor.org/sici?sici=0037-7732%28198109%2960%3A1%3C188%3ASSASS%3E2.0.CO%3B2-Z> [March 17, 2005].

V. E. Krebs, "Mapping Networks of Terrorist Cells," *Connections* 24, 3 (2002): 43-52. Available: <http://www.sfu.ca/~insna/Connections-Web/Volume24-3/Valdis.Krebs.L2.pdf> [March 17, 2005].

- 16 K. M. Carley, "Dynamic Network Analysis." In Ronald Breiger, K. M. Carley and P. Pattison (eds.), *Dynamic Social Network Modeling and Analysis: Workshop Summary and Papers*, Committee on Human Factors, National Research Council (National Research Council, 2003), 133-145.
- 17 F. Lorrain, H. White, "Structural Equivalence of Individuals in Social Networks," *Journal of Mathematical Sociology* 1 (1971): 49-80.
- 18 D. Krackhardt, K. M. Carley, "A PCANS Model of Structure in Organizations." In *Proceedings of the 1998 International Symposium on Command and Control Research and Technology* (1998), 113-119.
- 19 K. M. Carley, "On the Evolution of Social and Organizational Networks." In *Research in the Sociology of Organizations* 16 [special issue on 'Networks In and Around Organizations'] (1999), 3-30.
- 20 K. M. Carley, "Smart Agents and Organizations of the Future." In Leah Lievrouw and Sonia Livingstone (eds.), *The Handbook of New Media: Social Shaping and Consequences of ICTs* (Thousand Oaks, CA: Sage Publications, 2002), 206-220.
- 21 L. C. Freeman, "Centrality in Social Networks: Conceptual Clarification," *Social Networks* 1 (1979): 215-239.
- 22 S. Borgatti, "The Key Player Problem." In *Proceedings of CASOS 2002 Conference*, Pittsburgh, PA (2002).
- 23 E. J. Bienenstock, P. Bonacich, "Balancing Efficiency and Vulnerability in Social Networks." In R. Breiger and K. M. Carley (eds.), *Summary of the NRC Workshop on Social Network Modeling and Analysis* (National Research Council, forthcoming).
- 24 J. Johnson, L. Palinkas, J. Boster, *Informal Social Roles and the Evolution and Stability of Social Networks* (forthcoming).
- 25 K. M. Carley, "A Theory of Group Stability," *American Sociological Review* 56, 3 (1991): 331-354.
- 26 P. Klerks, "The Network Paradigm Applied to Criminal Organizations: Theoretical Nitpicking or a Relevant Doctrine for Investigators?" *Connections* 24, 3 (2001), 53-65. Available: <http://www.sfu.ca/~insna/Connections-Web/Volume24-3/klerks.pdf> [March 17, 2005].
- 27 H. Simon, "A Behavioral Model of Rational Choice," *Quarterly Journal of Economics* 69 (1955): 99-118.
- H. Simon, "Rational Choice and the Structure of the Environment," *Psychological Review* 63 (1956): 129-138.

28

K. M. Carley, A. Newell, "The nature of the Social Agent," *Journal of Mathematical Sociology* 19, 4 (1994): 221-262.

29

K. M. Carley, M. Prietula, "ACTS Theory: Extending the Model of Bounded Rationality." In Carley and Prietula (eds.) *Computational Organization Theory* (Hillsdale, NJ: L. Erlbaum, 1994).

30

D. Krackhardt, "Assessing the Political Landscape: Structure, Cognition, and Power in Organizations," *Administrative Science Quarterly* 35 (1990): 342-369.

31

K. M. Carley, "Organizational Learning and Personnel Turnover," *Organization Science* 3, 1 (February 1992): 20-46. Available in JSTOR: <http://links.jstor.org/sici?sici=1047-7039%28199202%293%3A1%3C20%3AOLAPT%3E2.0.CO%3B2-J> [March 17, 2005].

32

K. M. Carley, D. M. Svoboda, "Modeling Organizational Adaptation as a Simulated Annealing Process," *Sociological Methods and Research* 25, 1 (1996): 138-168.