

Self-Organized Terrorist- Counterterrorist Adaptive Coevolutions, Part I:

A Conceptual Design

Andrew Ilachinski




4825 Mark Center Drive • Alexandria, Virginia 22311-1850

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE FEB 2005		2. REPORT TYPE		3. DATES COVERED 00-00-2005 to 00-00-2005	
4. TITLE AND SUBTITLE Self-Organized Terrorist-Counterterrorist Adaptive Coevolutions, Part I: A Conceptual Design				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) CNA Analysis & Solutions,Center for Naval Analyses ,4825 Mark Center Drive,Alexandria,VA,22311				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 304	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Approved for distribution:

February 2005

A handwritten signature in dark ink, reading "Gregory M. Swider". The signature is written in a cursive, slightly slanted style.

Dr. Gregory M. Swider
Director, Tactical Analysis Team
Operations Evaluation Group

This document represents the best opinion of CNA at the time of issue.
It does not necessarily represent the opinion of the Department of the Navy.

Approved for Public Release; Distribution Unlimited. Specific authority: N00014-00-D-0700.
For copies of this document call: CNA Document Control and Distribution Section at 703-824-2123.

[Osama bin Laden is a] “...product of a new social structure; a new social feeling in the Muslim world. Where you have strong hostility not only against America, but also against many Arab and Muslim regimes who are allying to America... and that's why, if bin Laden was not there, you would have another bin Laden. You would have another name, with the same character, with the same role, of bin Laden...
...That's why we call it a phenomena, not a person.”¹

“Osama bin Laden’s most brilliant stroke may well have been to allow the global Salafi Jihad network to evolve spontaneously and naturally, and not interfere too much with its evolution, except to guide it through incentives because of his control of resources. The system developed into a small-world network with robustness and flexibility and became more militant and global for both internal and external reasons.”²

1.Extract from an interview with Saad al Fagih (a London-based Saudi Arabian exile who heads the *Movement for Islamic Reform in Arabia*), on Public Broadcasting Service's *Frontline*, 1999.

2.M. Sageman, *Understanding Terror Networks*, University of Pennsylvania Press, 2004; page 172.

Contents

Summary	1
Introduction	5
Purpose	9
Background.	10
Approach	11
Semiotic Agents	12
Modeling terrorist-counterterrorist coevolutions . . .	15
Issues, problems and questions	17
Organization of paper	22
Complex networks: basic concepts	25
Introduction	25
Formalism	31
Research problems	32
Basic terminology	34
Mathematical representations	38
Graph visualization.	40
Complex networks: zoology	47
Graph space	47
Zoology of graphs.	49
In what region(s) of graph space do TNs live?	51
Random graphs	52
Erdos-Renyi random graphs	52
Small-world random graphs	55
Scale-free random graphs	57
Clustered scale-free networks	61
Triad-creation	61
State-activation/deactivation	62
Search in complex networks.	63
Navigability Problem	64
Local Search	64

Evolving searchable nets	67
Network search and information theory	72
Dynamic graphs	75
Structurally dynamic cellular automata	76
Complex networks: metrics	83
Overview	83
Characteristic Path Length	84
Clustering Coefficient	86
Degree centrality	88
Group degree	90
Link Degree	90
Link Density	91
Eigenvector centrality	92
Information Centrality.	93
Closeness centrality	94
Group closeness.	95
Betweenness centrality	96
Computation times	97
Flow betweenness.	98
e-betweenness	98
Efficiency centrality	100
Node information centrality	101
Group information centrality	101
Graph efficiency centrality	101
Comparison with other centrality measures.	102
Structural Holes	102
Effective Network Size	104
Efficiency	107
Community structure	107
Finding community structures	108
Example 1: social networks	109
Example 2: terrorist networks.	110
Cluster finding algorithms	111
Quantifying community structure	113
Vulnerability	115
Network Reliability	117
Error tolerance versus vulnerability to attack	118
Attack Strategies	121

Cascade attacks	121
Using ordered set theory to break Al-Qaeda cells. . .	123
SOTCAC: conceptual design	127
Modeling ontology	128
SOTCAC's ontology	131
Design overview	133
What is, and is not, being simulated	138
Terrorist network.	141
T-agents	141
T-agent types	143
Recruit	143
Infiltrator.	145
Recruiter	145
Trainer	145
Mission Operative	145
Leader	146
Support Agents	146
T-agent characteristics	147
Primitive	147
Composite	149
Dynamic	149
Social network maps	150
LoTo-map	150
Ego-map	152
Trusted Prior (TP) Contacts	154
Current Contacts (CCs).	155
Terminated Contacts (TCs).	155
Potential Contacts (PCs)	156
T-agent personality	156
Physical domain	158
Information domain.	159
T-agent actions	160
Movement	160
Acquiring manpower: recruiting	160
Acquiring skills: training	161
Acquiring resources	161
Communications	161
Defection.	162

Promotion	163
Cells	163
Size.	164
Adhesion	164
Coordination Strength	165
Cohesion	166
Social network links	166
Invisible links	167
Communication.	167
Type	168
Content	169
Vulnerability	170
Adaptive topology.	171
SDCA.	171
EINStein's rules.	172
Social network interaction rules	177
SOTCAC's link rules	184
Motivations.	185
Constraints	187
Counterterrorist network	188
Functions	188
INTEL Assets	190
Physical Agents	190
Virtual Resources	191
CTN Actions.	192
CTN Beliefs	198
OperativeID-belief vector	199
Composition-belief matrix	200
Structure-belief matrix	201
Activity-belief vector	202
CTN inference personality	202
Interpretation of INTEL data.	203
Valuation of INTEL data	204
Fusion of INTEL reports	204
Tactical action plan logic	207
Conclusion	215
Appendix 1: Social network analysis	221

Example	224
Appendix 2: Mapping Al-Qaeda	229
Trusted Prior Contacts	230
Meeting Ties	232
Direct & Indirect Links	234
Observations	236
Lessons	238
Appendix 3: Social network analysis and SOTCAC-related develop- ment resources	241
AGNA	242
aiSee	242
Combinatorica/Mathematica	242
Combinatorica Graph Editor.	243
GraphPlot	243
Graphviz	243
JUNG	244
LEDA	244
Maple/Networks Package	244
NetDraw	245
NetMiner	245
Pajek	246
PIGALE	246
SNA/RSCE	247
UCINET	247
Tom Sawyer Software	247
Visone	248
Self-Organized Networks.	248
International Network for Social Network Analysis	249
Appendix 4: World Wide Web URL links to resources related to ter- rorism, nonlinearity and complex adaptive systems	251
References	255
Bibliography	277
List of figures	287
List of tables	293

Summary

“The enemy we face is a loose coalition of semi-independent terrorist cells, each with a well-defined mission and a high degree of adaptability and flexibility in carrying out that mission. Al Qaeda does not rely on immediate direction from a central authority yet still maintains effective coordination...and hence has been far less susceptible to intrusion or destruction. It adapts its methods to accomplish its goals.”—T. Irene Sanders¹

This paper—the first of a projected series of papers—examines the proposition that *terrorist networks*, such as Al Qaeda, are *complex adaptive systems*; that is, they consist of widely dispersed, autonomous cells that obey a decentralized command and control hierarchy; their mission operatives are highly adaptive and mobile; their cells are strongly compartmentalized, structurally robust, and largely impervious to (unfocused) local attack; and, though the networks, as a whole, are typically covert and amorphous, they can also rapidly coalesce into tightly organized local swarms. This implies that, in principle, terrorist networks, as dynamical systems, ought to be amenable to the same methodological course of study as any other complex adaptive system (such as a natural ecology, a biological immune system, or the human brain). In particular, fundamental insights into the behavior of terrorist networks—including an understanding of *how they form*, *how they evolve*, *how they adapt* (to changing internal and external contexts), and *what their innate strengths and vulnerabilities are*—may be gleaned by studying the patterns that emerge from a multiagent-based simulation of their dynamics.

This paper has two primary goals: (1) to review existing analytical and modeling tools that are applicable to the study of dynamic networks (including *mathematical graph theory*, *social network modeling*, *complex network theory*, *graph visualization*, and *multiagent-based modeling*), and outline how these tools may be leveraged to help understand the dynamics of terrorist networks, and (2) to introduce the conceptual

1. T. Irene Sanders is executive director and founder of the *Washington Center for Complexity & Public Policy*, Washington DC. The quote appears in the article, “To Fight Terror, We Can't Think Straight,” *Washington Post*, May 5, 2002.

design of a new multiagent-based toolkit, called SOTCAC (*Self-Organized Terrorist-Counterterrorist Addaptive Coevolutions*). SOTCAC uses autonomous, intelligent agents to represent the components of *coevolving* terrorist and counterterrorist networks

Graph theory provides a mathematical formalism with which to represent arbitrary relationships among the components of a complex system; as well as a computational aid for discovering latent patterns embedded within those relationships. *Social network analysis* studies patterns of relationships (such communication, information and resource flow, etc.) among individuals and organizations, and is particularly adept at revealing otherwise “hidden” patterns inside of a network; for example, information-flow bottlenecks and other vulnerabilities of, say, a business organization, that are not obvious from a wire-diagram of its members.

While social network analysis has traditionally confined its attention to the study of relatively small (and static) networks, the emerging interdisciplinary research field of *network science* studies the statistical properties of very large, complex networks, focusing on the relationship between the structure and function of evolving networks. As the size and complexity of a network of interest increases (beyond that of a manageably small size of at most, say, a few dozen nodes that can all be easily visualized at once), there is a growing need to develop algorithms to graphically render the structure of complex networks. Novel *graph visualization* techniques have recently been developed that facilitate the visualization of multidimensional feature spaces and the mapping of conceptual spaces into physical space.

Finally, SOTCAC builds upon the multiagent-modeling technologies underlying both the EINSTEIN and SCUDHunt simulations, recently developed at the *Center for Naval Analyses* (CNA).² Agent simulations provide a powerful, generative modeling environment for performing exploratory analyses on self-organized emergent behavior in complex adaptive systems.

2. EINSTEIN and SCUDHunt are agent-based simulations of *ground combat* and *shared situational awareness* in a wargame context, respectively. Those portions of these models that bear directly on SOTCAC’s design are described in the main text of this paper. For additional details, see A. Ilachinski, *EINSTEIN*, CRM 2239, 2000, and P. Perla, *et al.*, *Using Gaming and Agent Technology to Explore Joint Command and Control Issues*, CRM 7164, 2002.

SOTCAC uses adaptive agents to describe the self-organized, emergent behavior of terrorist networks—conceived as *complex adaptive systems*—on three interrelated dynamical levels:



1. **Dynamics on networks**, in which notional terrorist agents process and interpret information, search and acquire resources, and adapt to other agents' actions;
2. **Dynamics of networks**, in which the terrorist network *itself* is a fully dynamic, adaptive entity; and whose agents build, maintain, and modify the network's local (and therefore, collectively, its global) topology; and
3. **Dynamics between networks**, in which the terrorist network and counterterrorist network mutually *coevolve*; the terrorist network's "goal" is to achieve the critical infrastructure (of manpower, weapons, financial resources, and logistics) required to strike, while the counterterrorist network's mission is to prevent the terrorist network from achieving its goal.

All actions within SOTCAC take place within two dynamically reciprocal domains: (1) a *physical domain*, which represents a notionally real space in which both terrorist and counterterrorist agents are free to move about in and interact, and (2) an *information domain*, which represents the abstract space that contains the terrorist network's evolving social network (as well as an associated network that represents the counterterrorist's "best guess" as to what the terrorist network's topology looks like at a given time. Figure 1 shows a schematic of some of SOTCAC's main components.

Although much is known about the statistical properties and behaviors of *static* networks (learned mostly from graph theory, social network analysis and complex network science), the deeper question of how the local properties of *evolving* networks are related to emergent global structure and behavior remains essentially unexplored.

It is this need for having an analytical and/or modeling tool that addresses these fundamental dynamical relationships—and an especially urgent need by the defense intelligence community, as it struggles to "understand" the new enemy; i.e., the *complex adaptive terrorist network*—that has led directly to SOTCAC's development.

Figure 1. A schematic overview of some of SOTCAC's components discussed later in this paper (T=Terrorist, CT=Counterterrorist)



• T-agents

- Recruits
- Recruiters
- Infiltrators
- Trainers
- Mission operatives
- Leaders
- Support agents

• Properties

- Ability
- Move/sense range
- Rank
- Resources
- Value
- Visibility
- Processing load

• CT-agents

- *Physical* (HUMINT)
 - Move/sense range
 - T-agent detect range
 - T-agent detect prob
- *Virtual* (Notional)
 - COMINT
 - ELINT

• CTN-Belief Matrix

- OperativeID
- Composition
- Structure
- Activity

• Ego-map

- Trusted prior contacts
- Current contacts
- Terminated contacts
- Potential contacts

• Motivations

- Mission-centric (resource acquisition)
- Topological
- Behavioral

• Personality

- Allegiance
- Independence
- Discovery risk aversion
- Cell-cell mix proclivity
- Disconnect intolerance

• Actions/functions

- Collect INTEL about TN's activity
- Assimilate CT-agent-filtered data
- Generate "inferred" map of TN
- Issue move orders to CT-agents
- Issue attack orders
- Capture, insertion, and/or reconnaissance
- Eavesdrop and/or intrude on links

• CTN inference personality

- Interpretation of INTEL data
- (CT-agent-specific) valuation of data
- Fusion of INTEL reports (~SCUDHunt)
- Tactical action plan logic

Introduction


“While nothing is easier than to denounce the evildoer, nothing is more difficult than to understand him.”—Fyodor Dostoyevsky

The terrorist attacks on September 11, 2001 [1,2] unveiled to the world the face of a powerful new adversary of an entirely different character from the “conventional” enemies—and their tank- and mortar-driven armies—of old. The new adversary is not bounded by the geography of national borders; it is not driven by state-centric political motives; it does not fight with traditional armed forces; it is not organized according to traditional military rank and command and control hierarchies; and it operates, always, in stealth and remains largely invisible, until it is ready to strike. This new adversary is the *terrorist network* (TN).

Figure 2 shows the social network of ties among the 9/11 hijackers, as reconstructed using information made public shortly after the terrorist attacks (Krebs [3]). While this “snapshot” of the 9/11 social network is neither complete, nor completely accurate, it nonetheless illustrates the TN’s manifestly intelligent design, one that is highly tailored to efficiently carry out the terrorist operation. This network does not result from a willful imposition of order, from a top mastermind on down the chain of command, but is instead a self-organized, emergent entity that is generated by a combination of Al Qaeda’s mission requirements and the (decentralized and essentially autonomous) initiative of its agents [4].

TNs are adaptive, amorphous, decentralized, dynamic, evolving and self-organized; i.e., they constitute a perfect textbook example of a *complex adaptive system* (CAS) [5,6,7] (see table 1): they consist of widely dispersed, autonomous cells that obey a decentralized command and control structure; cell operatives are highly adaptive, compartmentalized, mobile; cells are structurally robust, and largely impervious to (unfocused) local attack; and the network, as a whole, typically functions in a localized, intensely swarm-like fashion.

Figure 2. Social network of ties among 9/11 hijackers (Krebs [3])



This strongly suggests that insights into the operation of TNs—as well as into ways of intruding into them and/or disrupting their operations—may be gleaned by studying patterns and behaviors that emerge from a multiagent-based simulation of their dynamics [8,9,10]. Because TNs extend their roots just as deeply into an abstract “information space” (whose activity is mostly hidden from outside view) as they do into the physical world, identifying their strengths and weaknesses require the use of a radically different set of metrics from the ones that have been developed specifically for assessing an enemy’s offensive capability (i.e., its “order-of-battle”).

Table 1. Terrorist networks as complex adaptive systems; compare to table 1.3 in [11]

General Property of Complex Systems	Description of Relevance to Terrorist Networks
<i>Nonlinearity</i>	Terrorist networks are composed of a large number of nonlinearly interacting parts: sources of nonlinearity include feedback loops in the command & control chain, adaptation to counterterrorist actions, and elements of chance
<i>Networked Dynamics</i>	Terrorist networks consist of many components (including foot-soldiers, mission trainers, weapons and logistical specialists, cell leaders, financiers, etc.), all of which are linked by a dynamic social network of ties
<i>Nonreductionist</i>	The terrorist network's ability to project force and carry out acts of terror cannot be understood as a simple aggregate function of the "fighting ability" of individual terrorists
<i>Adaptive</i>	In order to survive, terrorist networks continually adapt to both internal dissension and external counterterrorist forces; the most powerful terrorist networks have the additional capability to adapt both their tactics and strategies as their "enemies" adjust their own policies and counterterrorist strategies
<i>Bounded Rationality</i>	Individual terrorists have neither infinite resources nor operate in an environment with infinite information; they are constrained to choose their actions quickly, locally and use bounded (i.e., sub-optimal) information
<i>Emergence</i>	Neither the function nor topology of terrorist networks is scripted; rather both unfold, or emerge, out of a continual feedback of local interactions and adaptation to external influences; for example, the 9/11 terrorist cells emerged, spontaneously, from a friendship between Muhammad Atta and two other foreign students in Germany [12]; Al-Qaeda's evolving structure owes more to circumstance than design
<i>Hierarchical Structure</i>	As is true of conventional military forces, the core organizational structure of terrorist networks is loosely hierarchical: there are leaders, an inner circle of close advisors, lieutenants who run day-to-day activities, technical experts, and recruits.
<i>Decentralization</i>	There is no single geographical base of operations or reference, nor is there a single master "oracle" that dictates the actions of terrorist cells or individual agents; 9/11 cells were formed not via the "direction" of a central node, but by the loose, informal ties between a growing network of extremist Muslims sharing a common cause
<i>Self-organization</i>	Though terrorist network activity appears chaotic, locally, it displays manifest long-range order (both locally and globally); for example, just in terms of its finances, Al Qaeda is built upon multiply redundant and distributed financial channels, and both agents and cells are encouraged to nurture autonomous revenue channels [13]
<i>Nonequilibrium Order</i>	Due to its fundamentally amorphous nature, a terrorist network is almost never at equilibrium (either locally or with the external environment); rather, it is best characterized as a system of perpetual "unfolding," as it struggles toward achieving its goal in the face of surrounding counterterrorist forces
<i>Micro::Macro Feedback</i>	There is a continual feedback between the behavior of (low-level) foot-soldier terrorists and the (higher-level) cell leaders; as well as between cell leaders and the leader of the terrorist organization
<i>Autopoiesis</i>	While the identity of terrorists and their leaders, as well as the composition of terrorist cells, changes over time, the viability and function of the of the TN, as a whole, persists; since part of what drives the TN's evolution over time is its own, internal, set of (counterintuitively, self-sustaining) disruptive forces, the process is demonstrably autopoietic; (see quote on the top of the next page)

[Al Qaeda] *"...is directed from the bottom up as much as the top down. The typical pattern before September 11th was of local Al Qaeda cells initiating reconnaissance of potential targets, planning and then going back to the al Qaeda leadership for approval and possible funding. The foot soldiers are self-initiating and self-sustaining."* [14]

Rather than counting the number of soldiers in the enemy's army, and the firepower of its tanks and bombers, new complexity- and network-based measures must be defined to characterize a TN's order-of-battle. Such measures might include, for example, properties of a network's local and global structure, measures of how well nodes (and groups of nodes) communicate with (and facilitate communication with) one another, mathematical relationships between network topology and information flow, relationships between cell autonomy and group activity, the degree of cooperation and task-specific coordination within cells, the "value" that a given agent (or any subnet of the TN) represents in the context of the TN's overall mission, and the dynamical rules according to which TNs form, evolve, and adapt.

Because a TN's lifeblood is decentralized, self-organized, coordination (of vision, missions, requirements, and structure) among otherwise loosely and covertly connected parts—parts that "live" as much in an incorporeal world of dynamic, distributed information as they do in a physical one—traditional forms of attack against it, such as massing local firepower on selected physical targets, are doomed to fail. New methods of assault, not to mention that of understanding, are needed to combat this new enemy. As Dostoyevsky's eloquent aphorism suggests (at the top of page 5), achieving a true "understanding" of terrorist networks poses a formidable challenge. What one needs to understand and combat a network, *is another network*:

"Governments wishing to counter netwar terrorism will need to adopt organizational designs and strategies like those of their adversaries. This does not mean mirroring the opponent, but rather learning to draw on the same design principles of network forms... Netwar: an emerging mode of conflict (and crime) at societal levels, short of traditional military warfare, in which the protagonists use network forms of organization and related doctrines, strategies, and technologies attuned to the information age." [15]

Purpose

The purpose of this report is to compel the reader to appreciate the importance of the assertions made in the quote at the bottom of the last page; in particular, to appreciate the critical need for the intelligence and analytical research communities to develop network-based counterterrorist tactics and strategies, as well as the methodology by which the efficacy of these tactics and strategies is assessed, is emphasized throughout.

The report both (1) outlines how existing analytical tools—borrowed mostly from the mathematical graph and social network theory research communities—may be leveraged to model the dynamics of terrorist networks, and (2) presents a conceptual roadmap toward developing a new multiagent-based toolkit—called SOTCAC (Self-Organized Terrorist-Counterterrorist Addaptive Coevolutions)—that builds upon the core agent modeling technologies introduced, and tested, in the EINSTEIN [16] and SCUDHunt-Agent [17] simulations.

SOTCAC uses autonomous, intelligent agents to represent the components of *coevolving* terrorist and counterterrorist networks, and includes interactions among notional terrorist and counterterrorist intelligence agents, terrorist cells, training, logistical, and miscellaneous support networks, weapons and financial resource networks, and physical terrorist targets. It is designed to help analysts understand the emergent “fitness landscape” of the new enemy, modeled as a complex adaptive system.

Three major theses underlying SOTCAC’s design:

1. That TNs are, fundamentally, *self-organized, emergent “virtual multicellular organisms”* that live as much in the physical domain as they do in an abstract information space.
2. That the topology and function of TNs *coevolve with their enemy*.
3. That the best approach to understanding a TN—how it forms, grows, and adapts to changing environment or overt attack—is an interdisciplinary one that combines the precepts and methodologies of *systems theory, complex network science, social network analysis, mathematical graph theory, and multiagent-based modeling*.

A number of recent papers discuss terrorist networks, in general—and Al Qaeda, in particular—from the point of view of complex systems theory; see Beech [12], Fellman and Wright [18], Goolsby [19], Kaplan [20], Marion and Uhl-Bien [21,22], Mesjasz [23], Raab [24], and Sageman [25]. An excellent overview of the formation, coordination, development and adaptation of small groups as complex systems is given by Arrow, *et al.* [26]. Chapter nine of Rosenau’s monograph, *Distant Proximities* [27], cogently summarizes the dynamical effect that the micro behaviors of human institutions have on emergent, transnational global complexity and stability (including terrorism).

Background

The primary focus of a recently completed CNA project—*An Intelligent-Agent Based Conceptual Laboratory for Exploring Self-Organized Emergent Behavior in Land Combat* [28,29]—was to harness the tools of complex adaptive systems (CAS) theory to develop an “artificial-life”-like simulation of land combat called EINSTein. EINSTein was conceived with the premise that modern combat possesses the key characteristics of complex systems, and pioneered the application of multiagent-based modeling techniques to the fundamental understanding of the dynamics of warfare.

Even a cursory examination of the basic organizational and dynamical properties of TNs shows that they encompass the critical hallmarks of CASs (and, arguably, do so to an even greater degree than does land combat): TNs consist of a large number of “parts” (i.e. cells and individual terrorists) that are widely dispersed; TNs typically operate under a decentralized command and control; individual cells are autonomous, mobile and highly adaptive; attacks proceed in swarm-like fashion; and TNs are, as a whole, robust and strongly impervious to infiltration and random attack. Thus, it is reasonable to expect that many of the same mathematical and agent-based simulation techniques used to develop EINSTein, may also be harnessed to develop tools to better understand the behavior of terrorist networks.

While purely topological and other static characteristics of social (and terrorist) networks have been well studied, much less is known about the self-organized emergent dynamical properties of such systems as they interact (and coevolve) with an “outside” environment.

Intuitively, if a TN is viewed as a CAS, one expects a rich coupling between local interactions and global behavior; the full set of requisite conceptual and methodological tools needed for exploring this coupling have yet to be developed, and is the focus and goal of this study. This paper reviews the core ideas and mathematical techniques needed to understand dynamic social networks, particularly as they relate to terrorist nets and coevolutions between terrorist and counterterrorist organizations.


The final section of this paper introduces a conceptual model, called SOTCAC (*Self-Organized Terrorist-Counterterrorist Addaptive Coevolutions*), that uses EINSTEIn-like agents to represent the components of terrorist and counterterrorist networks (agents, cells, financial and logistical resources, weapons, targets, and so on).

Approach

The approach consists of first generalizing the class of “agents” as used by EINSTEIn (i.e., interpreted as information-processing entities that live in a notional “physical” space, namely a battlefield) to a class of agents that live in a more abstract information space (i.e. a *mathematical graph* whose vertices possess an inner “semantic space”) and assimilate, process and adapt to various forms of information (such as a multidimensional feature space that describes the properties of the Al Qaeda network, as derived from intelligence sources).

Figure 3 shows a schematic of how graphs—which are mathematical constructs consisting of abstract nodes and links [30]—can act as both analytical and conceptual intermediaries between a dynamic data space and multiagent-based models, such as SOTCAC.

Figure 3. Schematic of how mathematical graphs can be used to capture arbitrary relationships among objects, and serve as conceptual anchors of multiagent-based models



The typical “data space” of a physical system consists of a set of *objects*, the objects’ *properties*, and the dynamical *relationships* among objects; relationships may further be classified according to *type* and *strength*. A graph may be used to represent the topology of the system: objects are associated with nodes (or vertices); object properties with vertex labels; and relations with links. If the system is dynamic, one may still use a graph to represent the system, but the components of the graph assume an active role: *agents* replace vertices, and become free to *sense, think, move, and/or act*; a dynamic space of internal states replaces static vertex labels; agent $\leftrightarrow$ agent communications replace link labels; and interaction strength replaces a graph’s visual representation of the strength of a relationship between objects.

Semiotic Agents

We christen the generalized, dynamical graphical entities introduced above as *semiotic agents*, to distinguish them from “software agents,” used in more conventional, and simplified, agent models than is being proposed for SOTCAC. Semiotic agents generalize local decision making to networked, local/semi-global information sharing and collaboration. They are able to sense, interpret, understand and communicate the meaning of signs and symbols. “Living” within a multidimensional space (that contains both physical and abstract information components), semiotic agents are motivated by, and behave according to, various social, cultural, financial, and strategic influences, integrate bits of *a-priori* isolated and/or unrelated data, are able to discover resources by exploiting latent social connections, and strengthen or weaken existing connections to other members of the network and/or forge new links. The coevolving web of semiotic agents, as a *whole*, thus essentially forms a self-organized, “self aware” dynamic social network in which hidden meanings and patterns are allowed to emerge naturally on their own.

Potential applications of this approach range from understanding the fundamental dynamical characteristics and behaviors of terrorist networks, to developing dynamic models of terrorist networks as aids to intelligence analysts, to adaptive data-retrieval and data-mining, to semi-automated knowledge discovery and pattern recognition in intelligence databases.

While there are prototype systems that use certain aspects of this approach (see, for example, [31]-[36]), the author is unaware of any existing research that successfully integrates *dynamic graph theory*, *social network modeling* and *visualization*, *multiagent-based simulation* and *latent pattern recognition* and *knowledge discovery*.

For example, while the SNA community has developed tools for correlating *static* properties of graphs (such as connectivity, maximal degree and degree distributions, spanning trees, and so on)³ to its “*dynamics*”—the latter of which (in a social dynamics setting) typically involves such metrics as “who plays the central role in an organization,” or “who monitors communication flow”)—few, if any, tools exist for exploring networks that *evolve* and *adapt* to (internal and external) stimuli.


Still fewer studies have analyzed the dynamics of *coevolving* networks, such as the one being proposed in this paper. Moreover, conventional SNA typically confines its focus to relatively small networks, whose nodes have minimal (or no) internal states, and gives little or no attention to how such networks scale with size and/or complexity of inner dynamics.

EINSTEIN and SOTCAC both describe the emergent behaviors of systems of interacting agents, each of which has access to only a limited set of features of its immediate environment, adapts only to local stimuli, but also communicates, and interacts with other agents. Moreover, both programs depend, fundamentally, on *conflict* as a major source of dynamics. In EINSTEIN, the conflict is explicitly between a “red” and “blue” force; in SOTCAC, the conflict is between a terrorist network (equivalent to EINSTEIN’s “red force”) and a counterterrorist net (equivalent to EINSTEIN’s “blue force”).

The major difference between EINSTEIN and SOTCAC, is that while all of EINSTEIN’s agents live and act entirely in a physical (albeit only notionally physical) space, SOTCAC’s agents extend roots into both *physical* and abstract, *information*, spaces (see figure 4); the information space contains the terrorist network’s dynamic, adaptive web of social links.

3. See discussion in the section **Complex networks: metrics** and **Appendix 1: social network analysis**.

Figure 4. Schematic illustration of SOTCAC's coupled *information* and *physical* spaces; in contrast, the dynamics of the EINSTEIN combat model are confined solely to the physical domain



Thus, while in EINSTEIN, the agent action-space is effectively restricted to *movement* and *combat* (combined with, on a more superficial level, *communication* and *targeting*), in SOTCAC, the central component of an agent's action-space consists of the dynamic restructuring of its local "ego-centered map" of various kinds of social contacts (along with the information and resources that those contacts represent, or serve as indirect links to). SOTCAC's agents also move about on a notional 2D "activity field," but physical movement plays a subordinate role to the action that takes place in the information domain

Where network analysis has traditionally been used to explore the dynamics of human relationships in business organizations [37]-[40] and social communities and systems [41]-[45], the approach taken in this paper is to develop a broader, semiotic-agent-based methodology to help understand self-organized, emergent dynamical behavior in general information and communication networks. This approach represents a fundamental shift in focus away from studying relationships of *static* properties (such as correlating a business firm's financial "success" with its executive organization) to exploring the generative causes (and local and global consequences) of adaptive structural changes in evolving networks, tasked with performing possibly conflicting functions.

Modeling terrorist-counterterrorist *coevolutions*

SOTCAC has two opposing sides—*terrorists* and *counterterrorists*—that interact and coevolve. The main conceptual difference between EINSTEIN and SOTCAC (aside from the obviously different issues they address: one combat, the other dynamic terrorist networks), is that while EINSTEIN's agents confront each other in battle *directly* and *complementarily* (as the two forces are essentially "copies" of one another, albeit with different offensive/defensive capabilities and agent behaviors), in SOTCAC, the two ostensibly opposed sides have dynamically related but substantively different objects of focus:

- *Terrorists are interested in acquiring whatever manpower and material resources are required to commit acts of terror; a major aspect of this task is to create and maintain a robust social network of ties with other terrorists that simultaneously "optimizes" two general classes of conflicting constraints:*
 - *maximize coordination capability and/or information/resource flow through open channels, and*
 - *minimize the risk of discovery, intrusion, and/or destruction by counterterrorist forces.*
- *Counterterrorists are interested in either preventing this from happening or prolonging the time it takes the terrorists to marshall the requisite resources for launching their first strike.*

Both terrorist-agents (or T-agents) and counterterrorist-agents (or CT-agents) “live” in two coupled dynamical spaces (as illustrated in figure 4): (1) a *physical* space, and (2) an *information* space. The physical space consists of an N -by- M sized “box,” and defines the region in which T-agents and CT-agents maneuver and interact. Initially, the motion of all agents is random, but gradually assumes a unique agent-specific character as individual agents acquire experience and build and restructure their social networks (in the case of T-agents) or build their database of accumulated INTEL-discovery (in the case of CT-agents). The information space contains the TN’s evolving social structure, and is where the counterterrorist network (CTN) fuses INTEL-reports and updates its *beliefs* regarding the structure and actions of the TN.

While both T-agents and CT-agents have a relatively large palette of possible “actions” to choose from in any given context, all actions fall into one of three general classes: (1) *physical movement* (which applies to both types of agents), (2) *creating, or deleting, communication channels* (which applies, at least in the prototype version of SOTCAC, to T-agents only), and, for CT-agents, (3) *selective targeting of TN-components* (including both agents and links), for detection, intrusion, and/or capture.

Creating and deleting communication channels is the social-network-equivalent of “moving” in a physical space. Each T-agent sits at the center of its own unique “ego-centered” local map of the larger TN (which we will later just call the T-agent’s *ego-map*); it does not know the entire structure, only those parts that are accessible to it within its sensor range. The sensor range ($=R$) on a network is the maximum value of a generalized distance function, where “distance” is defined to be the minimum number of links that must be traversed along a path that connects one point in the network and another. For example, if the sensor range $R=3$, a T-agent’s local view extends only to those agents that are at most three “steps” removed. Since the TN’s social structures are dynamic, local map are not fixed.

The degree to which a given agent is motivated to accomplish specific tasks (within the physical domain) changes with time and context, and most strongly depends on its *type*. For example, while a *recruit*—who has not yet joined the TN—is also not yet explicitly motivated to

do anything, a *recruiter* (who is already a T-agent) is motivated to find new recruits (and the weights of his *movement*-personality vector are thus valued to favor chance encounters with possible recruits); and *mission operatives* are motivated to acquire the skills, weapons and other miscellaneous resources necessary to accomplish their assigned missions. A CT-agent may be “called” over by another CT-agent (who has detected a high-valued T-agent) to assist in “capturing” that T-agent.

Issues, problems and questions

The primary focus of this report is the understanding of TNs as complex, adaptive, self-organized, dynamic graphs. However, implicit in the ensuing discussion is that TNs are but a single—albeit singularly important—exemplar of a vastly larger general class of dynamic graphs, characterized by nodes that possess an interior dynamical space, communicate and exchange information with other nodes, and adaptively forge new links and/or sever existing ones. Consequently, the study of the dynamics of TNs shares (and is plagued by) many of the same outstanding issues and problems facing other interdisciplinary studies of complex networks.

The three main sets of *generic* research questions, that apply to any class of complex networks, are [46]-[49]:

1. “What properties are appropriate for characterizing the topology and behavior of real networks?”; “How can these properties be measured?”
2. “How do these observable properties arise, dynamically, within real networks?”; “How are they interrelated?”; “Can they be modeled in simulated networks?”
3. “What is the relationship between an evolving network’s local properties (i.e., its nodes and links—or, more generally—its agents), and its emergent global structure and behavior?”

While none of these questions have been fully explored or answered (particularly in the case of *terrorist* networks, about whose topology and properties reliable information is obviously difficult, if not impossible, to obtain), the first two sets of questions have garnered the widest attention. Sections two and three of this report review several recent models of network formation and structure.

The third question echoes the central— and, as yet, *unsolved*—problem in both *natural biology* [50] and *artificial life* research [51], which is to understand the relationship between an organism’s *genotype* (or the primitive instructions that are genetically encoded by the organism’s chromosome), and an organism’s *phenotype*, or its emergent, macroscopic form (which includes both its physical morphology and how it interacts with other organisms).

While considerable progress has been made toward understanding the genotype-phenotype relationship for biological systems in recent years, the more general study of the effects of local dynamics and topology on emergent behavior in large complex networks is in its infancy.⁴ It is fair to say that complex network theory—as a nascent “science”—is at the present time roughly where the mathematical theory of *chaos* was in the early 1970s [58]-[60], and *complexity* theory was in the middle 1980s [61]-[63]. Many of the necessary analytical tools and theoretical methodologies have yet to be developed. It is fervently hoped that the SOTCAC model described later in this paper will prove to be useful not just for studying the dynamics of terrorist networks (for which it is being explicitly designed), but for helping usher in an entirely new class of general-purpose multiagent-based dynamic graph models that can be used to explore the fundamental properties of complex adaptive, *evolving* networks [64,65].

4. For example, while the recently completed *Human Genome Project* (HGP) succeeded in identifying the approximately 25K genes in human DNA [52], the associated problem of understanding how the information that is encoded in those genes generates bodily forms and organs remains mostly a mystery. Two reasons why the problem is so difficult is that whatever is the true form of the mapping from genotype to phenotype, it is almost certainly both *nonlinear* [53, 54] and *nonlocal* [55]. Our current understanding of the genotype-phenotype problem in complex networks lags even farther behind. This is mainly because it is only very recently that, with the advent of telecommunications and computer networks, and the World Wide Web, the graph research community has shifted its focus away from analyzing small networks to studying the global properties of extremely large complex networks [56]. Moreover, just the problem of *visualizing* the structure of large complex networks (that contain more than, say a few thousand nodes), has spawned a major research effort to develop novel 3D graphical rendering techniques [57]. On the one hand, because the efficacy of these techniques obviously depends on how closely rendered structures correlate, visually, with the innate topological and statistical measures that characterize a given network, considerable effort is being put into developing new global (i.e. *phenotypic*) measures. On the other hand, because this work is so new, and it is not entirely clear how to best describe a network’s global “fitness,” the deeper issue of how local dynamics and topology are related to global network behavior is largely unexplored.

What follows is a sampling of complex-network-theory-inspired issues and questions that relate specifically to the analysis of terrorist networks; almost all of the problems listed below are examined, to varying degrees, in the remaining sections of this paper:

- *“What are the environmental factors (cultural, familial, financial, legal, political, and/or religious) that influence the dynamical formation of terrorist networks?”*
- *“How do TNs adapt to changing local and global environments?”*
- *“What environmental factors can be influenced, and in what way, so that the terrorist network’s ability to sustain itself is maximally disrupted?”*

Goolsby [66] argues that Al Qaeda’s network evolved from local insurgencies. Although the goals, needs, and ideologies of local groups are typically varied, Al Qaeda is able to sustain itself by tapping into these groups, transforming and/or realigning their ideologies as necessary, and subsuming them into a more global cause. Paradoxically, it is the *diversity* of these local, and only loosely associated, groups that provides Al Qaeda the links necessary for it to sustain its global network. The problem is to find ways of disrupting a terrorist network’s activities—beyond the simple, conventional, targeting of its existing links and structures—but attacking the core of its formative dynamics.

- *“What kinds of TN topologies emerge as a function of network evolution rules and/or strategies?”*

Multiagent-based models are particularly relevant for this issue, in that they provide a powerful “laboratory” in which to explore alternative TN evolutions in graph space, and to map out the superspace of all possible network landscapes.

- *“What are the appropriate metrics to characterize the communication and information flow within a TN?”; “How can an evolving network’s efficiency of information flow be tracked over time?”*

- “What is the relationship between a TN’s leadership structure, cell topology, ability to marshal resources (and other emergent global behavior), and the local rules that define how the TN’s agents interact and create/delete social links?”
- “What are the vulnerabilities of a TN to infiltration and/or attack, and how are they correlated with its agent dynamics and/or dynamic topology?”
- “What is the ‘best’ counterattack strategy (or strategies) that a TN can use to minimize the effects of whatever ‘disruptive’ attacks are used against it?”
- “What are the effects of “missing data” and/or missing “intelligence?”; “What analytical, computational and/or visualization tools are appropriate for inferring and/or predicting network structure given only an incomplete or uncertain ‘snapshot’ of network topology?”
- “What are the (possibly time-dependent) criteria by which critical nodes/links and/or local components of a TN may be identified?”; “Are there observable topological and/or information-flow patterns that can be used to predict TN activity (and/or provide insight into its otherwise covert operations)?”
- “What set of topological (and/or emergent dynamical) features characterize the subspace of all possible network structures to which GA-bred TNs evolve?” Genetic algorithms (or other evolutionary programming techniques) can be used to explore space of possible networks vis-a-vis various fitness measures. Sample “terrorist mission” fitness measures might include simultaneously *maximizing* intra-network cell communication ability and cell autonomy and *minimizing* vulnerability (from discovery, intrusion & disruption).

Genetic algorithms may also be used to search for viable counterterrorist strategies to use against TNs. For example, a typical counterterrorist “mission” might consist of using a specified set of intelligence agents, with specified constraints on their abilities, to delay as long as possible the TN’s ability to marshal whatever resources it needs to conduct its own mission.

- “What are the emergent properties of the coevolution of two or more interacting (that is, adversarial) networks?”

Given that one of the basic lessons of systems behavior from complexity theory is that a “network” (thought of as an the “enemy”) is best attacked by another network (rather than by a traditional top-down “hierarchy” ruled by hard-wired command and control structures), an obvious question is, “*What network properties are most conducive to a successful attack (as well as defense)?*”

- “*What is the dynamical relationship between a network’s ‘error tolerance’ (i.e. imperviousness to removal or random malfunctioning of a network’s nodes) and its ‘vulnerability to attack’ (that is, vulnerability to a focused disruption, by an outside force, of its existing communication links and/or nodes)?*”
- “*What graphical visualization tools are best suited for rendering the dynamic topologies of TNs?*”;
- “*What metrics are most useful for relating physical distance (such as that between nodes of a graph embedded in a conventional Euclidean space) and conceptual distance between, say, the ‘information’ content (defined in some suitable manner) of individual nodes?*”
- “*What network metrics (of both form and function) are necessary and sufficient to identify the critical components of a TN?*” “*What are the appropriate measures of ‘criticality’?*”
- “*What observable features of the TN ought the intelligence community focus its data collection resources on in order to maximize the probability of detecting the TN’s most valuable components?*”; “*If the counterterrorist organization must decide between destroying node A or B, what are the criteria by which an ‘optimal’ decision may be made?*”
- “*How can a TN be surreptitiously ‘probed’ (or indirectly stimulated) into revealing information about its structure and activity?* “
- “*What actions can a counterterrorist organization take to either prolong the period of time the TN requires to self-organize and marshal the resources it deems necessary to initiate an attack, or to prevent such an attack from ever taking place?*”
- “*How does a TN’s vulnerability to attack depend on its developmental timeline?*”; “*Is there a phase during which an intervention by the counterterrorist organization to intervene “optimally” disrupts (and/or destroys key parts of) the TN, and/or the TN’s ability to adapt?*”

Organization of paper

The remainder of this paper is divided into four relatively self-contained sections (the third section uses some of the formalism introduced in the second, and some details of SOTCAC discussed in the fourth section assume the reader has read portions of the third), and four appendices (that contain additional background material and links to resources:

- **Complex networks: *overview*.** The first section provides the framework for the ensuing discussion by introducing graphs (as conceptual modeling tools), the basic nomenclature, formalism and mathematical representation of complex networks, and an overview of graph visualization techniques.
- **Complex networks: *zoology*.** The second section contains a primer on random graphs (including discussions of the properties of, and constructive models for, Erdos-Renyi random graphs, small-world random graphs, scale-free random graphs, and clustered networks), local search and the general problem of “navigability” in complex networks, and an introduction to dynamic graphs. Dynamics graphs, which lie at the heart of the SOTCAC model, are graphs whose topology evolves according to dynamical rules that are functions of local information. A simple, but concrete example, of a dynamic graph is given, called *structurally dynamic cellular automata*
- **Complex networks: *metrics*.** The third section surveys the local and global topological properties of complex networks, emphasizing those metrics that have proven to be particularly useful in studies of human social networks. Properties discussed include the characteristic path length, clustering coefficient, degree, link and information centralities, betweenness, and several different measures of network efficiency. The section also contains discussions of structural hole theory and how topological metrics may be used, in practice, to understand emergent dynamical behavior; the latter topic is discussed in the context of network reliability and vulnerability to attack. The section concludes with an example of how ordered set theory, combined with a knowledge of basic structural properties of real networks, can be used to “break” Al-Qaeda terrorist cells.

- **SOTCAC: a conceptual model.** The fourth, and last main, section of this report introduces SOTCAC. SOTCAC is a conceptual model that uses autonomous, intelligent agents to represent the components of coevolving terrorist and counterterrorist networks, and includes interactions among notional terrorist and counterterrorist intelligence agents, terrorist cells, training, logistical, and miscellaneous support networks, weapons and financial resource networks, and physical terrorist targets. SOTCAC serves as the core “logical engine” within which the $TN \leftrightarrow CTN$ coevolution takes place, adjudicates the communications between, and interactions among, all terrorist and counterterrorist agents, and provides a visualization of the emerging graphical structures.
- **Appendix 1: Social network analysis.** The first appendix introduces social network analysis by applying network communication metrics, as defined in the main text for abstract networks, to a simple case study of human interactions.
- **Appendix 2: Mapping Al-Qaeda.** The second appendix summarizes a case study that applies the same social network analysis methodology that has traditionally been applied to understanding organizational structures for business firms, to mapping the structure (if not dynamics) of the Al Qaeda terrorist network. The case study uses only public information gleaned from major newspapers.
- **Appendix 3: Social network analysis and SOTCAC development resources.** The third appendix contains brief descriptions of (and WWW URL links to) some of the mathematical analysis, modeling and simulation toolkits that the author has tested for suitability for use in developing SOTCAC. The list includes toolkits that are both freely available, or shareware, programs that are designed mainly for academic research, and full-featured commercial development packages that run under multiple operating systems.
- **Appendix 4: World Wide Web resources.** The fourth appendix contains World Wide Web URL links to resources related to terrorism, nonlinearity and complex adaptive systems.

Complex networks: *basic concepts*

“As a net is made up of a series of ties, so everything in this world is connected by a series of ties. If anyone thinks that the mesh of a net is an independent, isolated thing, he is mistaken. It is called a net because it is made up of a series of interconnected meshes, and each mesh has its place and responsibility in relation to other meshes.”

—Buddha

Introduction

In 1967, Stanley Milgram, a social psychologist at Harvard, was interested in learning how many acquaintances it would take to connect two randomly selected individuals from the population [67]. To estimate this number, he sent out hundreds of letters to people in Nebraska, asking them to participate in his social contact study along with instructions about what to do regarding the identity of a selected “target” person (a stock broker in Boston).⁵

Since (at that time) no study of this kind had ever been performed before, Milgram had no idea what number of steps would be required to establish a link, or even if the number of letters he sent out was sufficient to establish any links at all. He was surprised, therefore, when the first letter found its way to the “target” within only a few days, having passed only through two intermediate links. In the end, Milgram discovered that 42 out of 160 letters made it to the “target,” some having passed through a dozen intermediaries. The median number of intermediate “nodes” was 5.5!

5. Those willing to participate in Milgram’s study were asked to follow these steps [68]: (1) add their name to the roster at the bottom of the letter; (2) detach a postcard (with their name and location), and send it to Milgram; (3) mail the study folder to the “target” if they *already know the target*; or (4) if they *do not* know the “target,” mail the study folder to some personal acquaintance of their’s who they believe may know the target.

What Milgram discovered is a basic property of an important class of large networks that is both ubiquitous and (and, up until fairly recently) mysterious, in the sense that there has heretofore not been a generative “explanation” for why this property exists. It has come to be known as the *small worlds* property.⁶ A “small world” network is a network that may be very large, but has at least one relatively short path (and, typically, many short paths) between any two of its nodes. This property is colloquially known as the “six degrees of separation” concept (after Milgram’s experiment), and has recently been extended and popularized by Watts [71,72].⁷

The small worlds property, alone, *does not* imply that a network has evolved according to a particular organizing principle. Indeed, large networks consisting of completely random links will possess this property, as do many different kinds of real-world networks, ranging from networks of chemical interactions in cells, to the coauthorship and citation networks of scientists, to the distance of separation of actors in Hollywood, to the foodweb of predator-prey interactions, to the neural network of the *Caenorhabditis elegans* worm [75].⁸ However, what is noteworthy about the small worlds property is that—due partly to its ubiquity and partly to its counterintuitive-seeming nature—it has served as a catalyst, stimulating further research into the properties of complex networks; research that in recent years has resulted in deeper insights.

-
6. The existence of a “small-worlds” property (though not by this name) was speculated upon about 30 years before Milgram’s work, in a 1929 Hungarian short-story [69]. See also [70].
 7. The familiar colloquialism “Six Degrees of Separation” is due to a play by that name, written by John Guare [73]. In mathematical circles, there is a related “Erdos Number,” that is assigned to all mathematicians, and which measures the minimum number of steps required to trace an individual’s academic papers to some paper authored by the prolific Hungarian mathematician, Paul Erdos. See the *Erdős Number Project*: <http://www.oakland.edu/enp/>. Watts has also recently reproduced Milgram’s experiment using email [74].
 8. See, for example, Albert-Barabasi [76], Buchanan [47], Newman [46] and Strogatz [77].

Kleinberg [78] points out that Milgram’s experiment actually reveals not one, but *two* fundamentally surprising results, both of which play important roles during the ensuing discussion:

1. That short paths *exist* in acquaintanceship networks (i.e., the aforementioned small-worlds property), and
2. That individuals belonging to these networks are able to *determine* what these short paths are.

While the first result is arguably counterintuitive, it is also demonstrably real and existential in nature; the second result is subtle, and is really an algorithmic assertion: it implies that individuals possessing only local knowledge (i.e., the locations of their direct acquaintances) can nonetheless, collectively, construct a short path between two points in the network. We will return to this important observation later in this paper, where we use it as a conceptual stepping stone for suggesting ways in which terrorist-agents (thought of as components of a terrorist-acquaintance “social network”) rely on decentralized search algorithms to acquire the resources required for performing their missions.

Traditionally, the study of complex networks has been under the purview of mathematical graph theory. While early studies focused on relatively simple structures (mainly those that could be readily described mathematically), attention shifted in the 1950s to the study of random graphs as a basis for understanding large networks with no apparent design. Pioneering studies of this type were first performed by Erdos and Renyi [79].⁹ While *random graph theory*, as it has come to be known, established itself as the prevailing conceptual and modeling methodology during the ensuing decades, there has been growing interest in recent years to explore alternative models of complex networks in order to better understand the evolution, structure and dynamics of complex systems.

9. Random graphs are generated by randomly adding links between a fixed number of nodes, n , so that each of the $n(n-1)/2$ possible links has the same probability, p , of being added. See pages 52-55.

Examples of complex systems taking the form of networks include the nervous system, the internet connectivity and WWW, interdependencies of business organizations and the global economy, food chains, epidemics, social networks of family, acquaintances and relations with other people, and—as we shall argue at length later in this paper—*terrorist organizations*.

Many (if not all) of these systems also fall under the broad rubric of *complex adaptive systems theory* (see pages 101-114 of [11]). Therefore, all of the interdisciplinary tools and (still evolving) methodology developed for studying these systems, from the point of view of *self-organized emergence*, apply.

However, the focus of this paper is on those aspects of the behavior of complex systems that derive explicitly from a given system's *topology*; i.e., the focus is less on the details of how parts interact locally (though that is obviously still important), and more on how interacting parts are *organized*, and on how that “organization” changes in time, as the system adapts to both internal and external stimuli.

In essence, the abstract, but fundamental, question being asked is:

“To what extent does a complex system's global behavior depend on the network of local interactions of its constituent parts?”

The more specific question in the context of using the theory of complex networks to obtain a deeper understanding of the dynamics of terrorism, is:

“To what extent is the behavior of terrorist organizations, and their interaction with counter- (and anti-) terrorist organizations, driven by their internal structure?”

If the *topology* of a system is an integral component of a given complex systems' dynamics, and it *does not* exhibit the properties expected of random graphs, new methods are clearly needed to identify and describe the underlying organizing principles. Before we review the theory and apply it to the study of terrorist networks, we will first introduce some basic terms and present some illustrative examples.

Formally, a *network* (or *graph*) is a set of abstract objects, which we will interchangeably call *nodes*, *vertices* or (if they are also endowed with an internal space) *agents*, with connections between them, called *edges* (or *links*).

The “nodes” of a network can be simple (for example, they can be used as nothing more than alphanumeric labels of individual components of a system) or more complex (by possessing a dynamic state that changes as a function of the states of other nodes in their local neighborhood and/or consisting of entire networks that evolve inside of them).


Likewise, “links” can be simple abstract “lines” that represent some form of connection between simple nodes, or more complex, and be directed, weighed, and/or consist of multiple edges. We will eventually make heavy use of these, and even more complicated, variants of a basic network.

Figure 5 shows four examples of large, complex networks. Though the examples come from very different contexts, they reveal some properties that appear to be shared, along with others that are unique to a particular example.

It is the chief task of complex network theory to develop the metrics and tools necessary to discriminate among the many different kinds of graphs that can exist. Figure 5-a, taken from [80], contains a map of *interacting proteins in yeast*, and is color coded according to level of importance for as cell’s survival: *red* = essential proteins (i.e., their removal cause the cell to die); *orange* = important proteins (i.e., their removal slow cell growth); and *green/yellow* = less essential proteins (or those whose importance is unknown). The figure reveals, suggestively, that highly linked proteins—or the *hubs*—tend to be the most critical for a cell’s survival.

Figure 5-b shows a social map of the CNN news agency, as generated by *Tom Sawyer Software’s* graph visualization toolkit (see **Appendix 3**). Note the relative sparseness of the graph, but with a high degree of local clustering. Figure 5-c shows a graph of internet industry partnership alliances among 250 companies, as recorded during the period from 1998 to 2001, and mapped using *OrgNet.Com’s* social network analysis toolkit *Inflow* [81].

Figure 5. Examples of large complex networks; see text for explanation



Two nodes are linked if the companies they represent had a joint venture, strategic alliance or other partnership during those years. Observe that most companies appear to have few partnerships, but a few have very many. The industry as a whole is dominated by several *hubs*. The network shows that, prior to 2002, the best positioned companies (where “best positioned” means companies that effectively serve as bridges to connect the unconnected) are *Microsoft*, *AOL-Time Warner* and *IBM*.

Finally, figure 5-d, created by the *Internet Mapping Project* [82], is a snapshot of the internet links as recorded in December 1998.

As we commented earlier, while systems such as those appearing in figure 5 have been traditionally modeled as random graphs, it is

increasingly recognized that the topology of many real networks is *not random*, and that their evolution is governed by robust, possibly universal, organizing principles. For example, an important feature that is common in social networks is *clustering*, which represents a local clique of friends and acquaintances, all of whose members know each other. That some model other than a random graph is necessary to describe such social networks was first pointed out by Watts and Strogatz [83], who observed that clustering in most (if not all) real networks is typically much larger than exists in a random network with an equal number of nodes and links.

The next section introduces basic nomenclature and concepts that will be used throughout the ensuing discussion, and reviews some recent advances in the field of complex networks that are particularly relevant to the focus of this study.

Formalism

A significant fraction of complex network research uses the formalism of mathematical graph theory, a subject that goes back to the *Konigsberg Bridges* problem devised by the great mathematician Leonhard Euler in the late 1700s.¹⁰ This section covers only those aspects of graph theory that are needed to describe the coevolutionary terrorist $\leftrightarrow$ counterterrorist model—called SOTCAC—that is introduced in a later section.

There are many excellent texts that develop the theory more fully than is possible to do here. See, for example, the texts by Berge [30], Diestel [85], and Harary [86]. The classic reference on random graph theory is by Bollobas [87]. In-depth reviews of complex networks are given by Albert and Barabasi [88], Dorogovtsev and Mendes [89], Newman [46], and Strogatz [77].

For insightful discussions of virtually any topic in graph theory the reader may wish to consult CRC's *Handbook of Graph Theory* [90]. Appendix 3 lists several readily-available academic and commercial software packages that provide analytical tools for studying and visualizing graphs.

10. The *Konigsberg Bridges* problem consists of finding the shortest route around the city's bridges in such a way that each bridge is crossed only once; see, for example, [84].

Research problems


There are three fundamental (and partially overlapping) research problems that define current studies of complex networks [46]:

1. *The characterization, and measurement, of local and global properties of real-world networks;*
2. *The development of models to simulate the evolution of real-world networks; and*
3. *Understanding the relationship between the structure of complex networks and the processes that take place on those networks.*

Unfortunately, complex networks are inherently difficult to study. Strogatz [77] suggests six reasons why this is so: (1) *structural complexity*: topologies are usually very complicated, as already evidenced by figures 5-a and 5-d; (2) *structural dynamics*: the topology of networks can change over time, such as happens to the World Wide Web as its many pages and links change very minute; (3) *topological diversity*: the links between nodes can assume many different forms, be directional, and/or contain additional tags or weights (the synapses of the brain, for example, alter their strength over time and can be either inhibitory or excitatory); (4) *dynamical complexity*: the (typically coupled) states of the nodes may change according to nonlinear dynamical functions and thus display complicated behavior; (5) *nodal diversity*: complex networks typically contain a complicated heterogeneous mix of nodal forms and functions (for example, figure 6 shows only a small portion of the molecular interaction map for the regulatory network responsible for mammalian cell cycles¹¹); and (6) *nodal-topological coupling*: the nodal states and overall topology of complex networks are typically also coupled and coevolving (for example, neural dynamics in the brain, which depend on the spatial synaptic patterns, do not only alter the neural states themselves but alter the synapses; i.e. nodal and synaptic patterns coevolve).

11. Figure 6 is taken from the second of four panels that appear between pages 2708 and 2711 in [91]. Different colors encode different interactions: *black* = binding interactions; *red* = gene expression; *green* = enzyme actions; and *blue* = stimulations and inhibitions.

Figure 6. A small portion of the molecular interaction map for the regulatory network responsible for mammalian cell cycles [91]




While almost all existing networks models (such as random graphs, small world graphs, and scale-free networks; see below) are probabilistic, and are designed so that certain desired properties (that are typically selected to reflect some desired real-world behaviors) emerge as the graphs evolve, the SOTCAC model introduced later in this paper is unique in two ways: (1) the nodes of its associated (terrorist and counterterrorist) graphs are fully autonomous agents, intelligent enough to use local information to “decide” their own manner of evolution; and (2) part of the information accessible to them, and is that forms the basis of their actions, is the set of graph metrics that describe their local topology. SOTCAC therefore represents not just a multiagent-based model of terrorist $\leftrightarrow$ counterterrorist net dynamics, but may be used as a more general model of intelligent graph evolution.

Basic terminology

A *graph*, $G(N, M)$, is a finite, nonempty set $V(G) = \{v_1, \dots, v_N\}$ of *vertices* (or *nodes*) together with (a possibly empty) set $E(G) = \{e_1, \dots, e_M\}$ of *edges* (or *links*). Each edge $e = e(i, j) \in E(G)$, connects vertex “i” to “j.” If the edges represent ordered pairs of vertices, then G is a *directed* graph: each $e(i, j)$ defines an edge directed from i to j , and is denoted by $e(i, j) = \vec{ij}$; graphically, directed links will be depicted using arrows, with tails anchored on the starting vertex and arrow-head pointed toward the terminal vertex. If, on the other hand, $e(i, j) = e(j, i)$ for all “i” and “j,” then G is an *undirected* graph and all edges are denoted simply by $e(i, j) = ij = ji$. Two vertices, i and j , are said to be *adjacent* if $e(i, j) \in E(G)$. An example of a small graph with eight nodes and ten links is shown in figure 7.

Figure 7. Graph consisting of eight nodes ($V = \{1, 2, \dots, 8\}$) and ten links ($E = \{(1, 2), (1, 7), (1, 8), (2, 4), (2, 8), (3, 4), (3, 6), (4, 5), (4, 6), (7, 8)\}$)



A *simple graph* is one that contains only single links between any two nodes. If all the links in a graph are used to denote only the presence, but not strength, of a connection between the nodes, the graph is called *unweighted*; if, on the other hand, links are assigned a strength of a connection, the graph is called a *weighted graph*. As we will see in a later section, SOTCAC generalizes the notion of a weighted graph by incorporating both multiple links (with each link representing different domains of possible connectivity) and endowing each link with *strength-vector* (the components of which are metrics that define a particular link, such as type, strength, duration, and vulnerability to eavesdropping and/or disruption). Finally, a *sparse* graph is one in which the number of links is significantly less than the maximal pos-

sible number for the graph; i.e., a graph $G(N,M)$ is sparse if $M \ll N(N-1)/2$.

The *order* of a graph G is the number of vertices of G : $|G| = |V(G)| = N$. The *size* of G is the total number of edges in G : $\|G\| = |E(G)| = M$. In principle, the order and size of G may both be infinite, but we will mostly consider cases in which they are both finite. Clearly, for every M , $0 \leq M \leq \binom{N}{2}$, there is a graph $G(N,M)$. The graph of order N and size $\binom{N}{2}$ is called a *complete N -graph*; it is denoted by K_N .

A graph $G(N,M)$ is *labeled* if all of its N vertices are associated with N distinct labels in a one-to-one manner. An *edge-labeled* graph is defined in an analogous fashion.


Two graphs, G_1 and G_2 , are *isomorphic*—which we write as $G_1 \cong G_2$, if there exists a bijection $V_1 \rightarrow \psi(V_1) = V_2$ which preserves adjacency (i.e. such that $e(i,j) \in E(G_1)$ if and only if $e(\psi(i), \psi(j)) \in E(G_2)$). Isomorphic graphs necessarily must have the same order and size. It is easy to see that isomorphism is an equivalence relation on graphs: it divides the collection of all graphs into equivalence classes, for which any two member graphs must have the same structure. Since all graphs in the same class are in this sense topologically equivalent, a representative member of each class, devoid of explicit vertex (or edge) labels, defines a unique unlabeled graph.

Figure 8, for example, shows all of the undirected and unlabeled graphs of order 4.

Suppose p is a permutation of the N labels, $\{1, 2, \dots, N\}$ and G is a labeled graph of order N . Let $p(G)$ represent the graph obtained from G by relabeling each vertex, i , by $p(i)$. Then the set $P(G)$, consisting of all permutations such that G and $p(G)$ are identical graphs is called the *automorphism group* of G . Since any graph G can be labeled in any of $n!$ distinct ways, the order of the automorphism group, $|P(G)|$, tells us how many of these labellings will be topologically equivalent: i.e. the number of ways of labeling an unlabeled graph is given by:

$$\alpha(G) = \frac{n!}{|P(G)|}. \quad (1)$$

Figure 8. All undirected, unlabeled graphs G of order 4



It is easy to show that the number of graphs of order N , $\mathfrak{N}_N = 2^{\binom{N}{2}}$, and the number of graphs of order N and size M , $\mathfrak{N}_{N,M} = \binom{\binom{N}{2}}{M}$.

$\tilde{G} = (\tilde{V}, \tilde{E})$ is a subgraph of $G = (V, E)$, written $\tilde{G} \subset G$, if $\tilde{E} \subset E$ and $\tilde{V} \subset V$. If $\tilde{G}$ contains all edges of G that join two vertices in $\tilde{V}$, then $\tilde{G}$ is the subgraph induced (or spanned) by $\tilde{V}$, and is denoted by $G[\tilde{V}]$. A subgraph $\tilde{G}$ is an induced subgraph if $\tilde{G} = G[V(\tilde{G})]$. $\tilde{G}$ is a spanning subgraph of G if $\tilde{V} = V$.

The set of vertices adjacent to a vertex $i \in V(G)$ is denoted by $\Gamma(i)$. The *degree* of i , $\deg(i)$, is defined to be the number of vertices adjacent to (i.e. the *neighbors* of) i : $\deg(i) = |\Gamma(i)|$. Note that, for undirected graphs, the sum of the degrees is exactly twice the number of edges: $\sum_{v \in V(G)} \deg(v) = 2M$. We will denote the minimum degree of G by $d(G)$, and the maximum degree by $D(G)$. In general, of course, $\delta(G) \leq \bar{d} \leq \Delta(G)$, where $\bar{d} = \frac{1}{|V(G)|} \sum_{v \in V(G)} \deg(v)$ is G 's average degree. If $d(G) = D(G) = r$ (i.e. each vertex of G has the same degree $\deg = r$) then G is an *r -regular graph*. A circuit, for example, can be alternatively defined as a *2-regular connected graph*.

A *walk* from the vertex v_1 to vertex v_i is a finite sequence of edges of the form $\{(v_1; v_2), (v_2; v_3), \dots, (v_{i-1}; v_i)\}$. In general, a given vertex or edge can be visited any number of times. A walk in which each vertex is distinct is called a *path*. A path of length N is therefore a graph con-

sisting of the distinct sequence of vertices $\{v_1, \dots, v_N\}$ and edges $\{(v_1; v_2), \dots, (v_{i-1}; v_i)\}$. If the edge (v_N, v_1) is added, the path becomes a circuit of length N . If we substitute the *arrows* (or *arcs*) $v_{i-1} \rightarrow v_i$ for the lines $(v_{i-1}; v_i)$, the circuit becomes a *cycle* of length N . If a graph G contains at least one cycle which itself contains all of the vertices of G , then G is said to be *Hamiltonian*.

G is *connected* if for every pair of distinct vertices in G there exists at least one path along existing edges. The connected (undirected) graph of order N which has the smallest size is called a *tree*. Any G that is described by any *one* of the following four properties is a tree: (1) G contains no circuits and has $N - 1$ edges; (2) G is connected and has $N - 1$ edges; (3) any two vertices of G are connected by exactly one path; and (4) G contains no circuits and the addition of any new edge creates exactly one circuit.

A *spanning tree* in G is an edge-subgraph of G which has $N - 1$ edges and contains no circuits (i.e., a subgraph of G that contains all nodes in G and is also a tree). It is easy to show that a graph is connected if and only if it has a spanning tree.

The *minimum spanning tree* of a weighted graph is the spanning tree that minimizes the total weights on the spanning tree of the graph. (If G is unweighted, then any spanning tree is a minimum spanning tree.) The minimum spanning tree is not difficult to find, computationally; polynomial time algorithms include those by Prim and Kruskal [92].

The *distance* in G between nodes i and j , $Dist(i, j)$, is equal to the length of the shortest path between i and j .¹² If no path exists (such as, for example, in the case when i and j are on two disconnected components of G), $Dist(i, j) \equiv \infty$.

The maximal distance between a node i and any other node in G is called the *eccentricity* of i , $Ecc(i)$. The *diameter* of G , $Diam(G)$, is equal to the greatest distance between any two nodes in G (i.e., the maximal eccentricity value of all nodes; the minimal value of eccentricity is G 's *radius*).

12. The shortest path between any two nodes in a graph is called a *geodesic* [84].

Mathematical representations

Although many properties of a graph may be identified by visual inspection alone (although even the seemingly “simplest” properties, such as connectivity, may often be difficult or even impossible to discern for large graphs), any kind of formal analysis requires the use of a symbolic mathematical representation.

A commonly used representation of a graph, G , is its *adjacency matrix*, $A(G)$. $A(G)$ of an order N graph with $V(G) = \{v_1, \dots, v_N\}$ is an N -by- N matrix $[a_{ij}]$ with entries $a_{ij} = 1$ if $(v_i, v_j) \in E(G)$ and $a_{ij} = 0$ otherwise. Therefore, for loopless graphs, $A(G)$ is a symmetric ($a_{ij} = a_{ji}$) $(0,1)$ -matrix with *trace* $Tr(A) = \sum a_{ii} = 0$ (i.e. A 's diagonal entries are all equal to zero). Note that this is an exact correspondence: i.e., the set of all N -by- N matrices satisfying these properties represents the class of all graphs G of order N .

The entries of the n^{th} power of $A(G)$ have a simple interpretation: $[a_{ij}^n]$, $n \geq 1$ is equal to the number of different walks from v_i to v_j of length n in G . In particular, $[a_{ij}^2]$, $i \neq j$ is the number of v_i-v_j paths of length two and $[a_{ii}^2] = d(v_i)$. Since there are $\binom{N}{2} = N(N-1)/2$ possible links among N nodes, and the maximal length of a path between any two nodes is $N-1$, the average number of independent paths between nodes in G is given by $\binom{N}{2}^{-1} = \sum_{i,j} \sum_{k \leq N} [a_{ij}^k]$.

While the rows and columns of A obviously depend on a particular choice of vertex labels, the generic structural properties of G must remain invariant under a permutation of rows and columns. Information about a graph's structure can often be extracted, analytically, from its spectrum: the *spectrum* of a graph G ,

$$Sp(G) = \begin{pmatrix} \lambda_1 & \lambda_2 & \dots & \lambda_s \\ m(\lambda_1) & m(\lambda_2) & \dots & m(\lambda_s) \end{pmatrix}, \quad (2)$$

is the set of eigenvalues, λ_i of $A(G)$, together with their multiplicities, $m(\lambda_i)$. Since $A(G)$ is real and symmetric, it follows that the eigenvalues must also be real. For example, the spectrum of the complete graph

of order n , $Sp(K_n) = \begin{pmatrix} n-1 & -1 \\ 1 & n-1 \end{pmatrix}$. Two non-isomorphic graphs, G_1 and G_2 , are *cospectral* if $Sp(G_1) = Sp(G_2)$.

The relationship between the topological properties of graphs and their spectra is an intensely studied field [93]. For some questions, the spectrum yields a wealth of knowledge; for other questions, it adds little to what can be obtained using other methods. Some useful structural information about G can also be obtained from the *characteristic polynomial* of $A(G)$:

$$\chi_G(\lambda) \equiv \det(A - \lambda \cdot I) = \sum_{i=0}^N a_i \lambda^{N-i}. \quad (3)$$

Since the coefficients a_i can be interpreted as sums of principal minors of $A(G)$, it is easy to check that (i) $a_1 = 0$, (ii) a_2 = the number of *edges* in G , and (iii) a_3 = twice the number of *triangles* in G . Considerably more powerful versions of this result can be proven; see, for example, chapter five in [94].

Another way to describe a graph is by its *adjacency list*, $L(G)$. $L(G)$ consists of a list of links $\{(v_i, v_j), (v_i, v_k), \dots\}$ incident to vertex i , for each $i \in V(G)$.

A third way to describe G is via its *incidence matrix*, $B(G)$. If G has vertex set $V(G) = \{v_1, \dots, v_N\}$ and edge set $E(G) = \{e_1, \dots, e_M\}$, then $B(G)$ is a N -by- M matrix $[b_{ij}]$ for which the $(ij)^{\text{th}}$ entry is defined by:

$$b_{ij} = \begin{cases} +1 & \text{if } v_i \text{ is the initial vertex of edge } e_j, \\ -1 & \text{if } v_i \text{ is the terminal vertex of edge } e_j, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

$B(G)$ can be used to define the *complexity* of a graph G , $k(G)$, equal to the number of *spanning trees* of G :

$$\kappa(G) = \det(J + BB^t) = \det(J + D - A), \quad (5)$$

where J is an N -by- N matrix all of whose entries are equal to one, D is an N -by- N diagonal matrix in which $[d_{ii}] = d(v_i)$, $\det(x)$ is the *determinant* of x , and B^t is the *transpose* of B . (Note that the relationship between G 's adjacency and incidence matrices is revealed by $k(G)$ via the relation $BB^t = D - A$.) [95]

Graph visualization

The most important kinds of visualization tools for the purposes of this study are those that optimize maps between *physical distance* (such as that between nodes of a graph embedded in Euclidean space) and *conceptual distance* between, say, the semantic content of individual nodes. A prototypical example is WEBSOM [96], which is a neural-network-based aid for, among other things, displaying self-organized maps of voluminous text data in which the physical distance between two documents is maximally correlated with semantic distance (as defined by keywords and/or subject matter).¹³

Intuitively, it is easy to visualize a graph, geometrically, in a two or three dimensional Euclidean space: use *points* to represent the graph's nodes and draw *lines* between nodes wherever they are linked (or *arrows* in the case of directed graphs). However, in practice, finding the best (or even merely “good”) way to render a given graph—so that the rendering simultaneously respects the graph's topology, allows direct visual inspection (with minimal distractions, such as multiple line crossings, for example), and is aesthetically appealing—can be an extremely difficult (if not impossible) problem to solve, particularly for large complex graphs.

A visual representation of a graph is useful to the researcher only to the extent that it conveys meaningful information. A “good” visualization aids in understanding the system, and often helps to reveal otherwise latent and/or hidden patterns; a “poor” visualization is confusing at best, and misleading at worst.

13. WEBSOM (websom.hut.fi/websom/) is but one example of more general *physical to conceptual space* mapping techniques, most of which use some variant of a “spring relaxation” algorithm: a network is modeled as a system of springs the nodes of which adjust their position by responding to attractive and repulsive forces due to other nodes (and the details of which depend on the “conceptual” space being probed); this is, in principle, very similar to the system of “personality weights” that EINSTEIN's combat agents use to maneuver on a battlefield. The system—or *graph*—relaxes over time to minimize the stress on all springs, thus evolving toward a state in which two nodes are close (in the graph), if their distance in the conceptual space is small.

There is an entire subfield of graph theory that is devoted exclusively to developing algorithms to automate visualization. See, for example, the texts by Battista, *et al.* [97], Kaufmann and Wagner [98], and Mutzel and Junger [99]. Texts on computational geometry, such as by O'Rourke [100], provide the mathematical background used by many graph drawing algorithms.

Some of the criteria used by graph drawing algorithms are purely aesthetic. That is, they attempt to satisfy as many (mostly intuitive) measures of readability as possible. These measures include minimizing the number of line crossings, minimizing the area of the graph, minimizing the maximum length of links, minimizing the variance of link size, maximizing the symmetry of the whole graph, among many others. Not all possible criteria can (or ought to be) applied to a given graph; a subset of properties may directly contradict one another and/or be aesthetically incompatible in other ways, and it has been shown that for certain aesthetic criteria the rendering problem is *NP-complete* [101] (meaning that its solution, in general, requires a time that grows exponentially with the size of the graph).

The most useful graph drawing algorithms, at least as far as the topics covered by this paper are concerned, are those that make some attempt at mapping the information that is encoded symbolically in the graph to the physical space in which the graph is being embedded. Some of the most common renderings are: (1) *circular*, (2) *ranked*, (3) *radial*, (4) *rooted*, and (5) *force-directed*.

The simplest graph rendering is a *circular embedding*, in which the positions of the nodes are all evenly spaced on a unit circle. While it does not take into account any aesthetic criteria for the arrangement, one virtue of a circular embedding is that each link is unambiguously represented (since no two nodes are colinear).

A *ranked embedding* is one in which nodes are placed on evenly spaced lines. This embedding is useful when the graph consists of loosely connected subsets of nodes, or nodes are selectively ranked according to a desired measure. For example, one might wish to rank the nodes according to their distance from a selected clique of nodes within the graph. A *radial embedding* generalizes a ranked embedding by arranging the nodes of a graph on a series of concentric circles

(around a selected node) in such a way that nodes of the same rank appear on the same circle. Nodes of the same rank that share a common neighbor of smaller rank are also clustered together.

A *rooted embedding*, which is used primarily for hierarchical structures, starts from an arbitrary (or selected) “root” node and distributes all other nodes along parallel lines. Nodes are ranked according to their distance from the root node, and each line contains nodes of the same rank.

So-called *force-directed* (or *spring-embedding*) methods are based on drawing an analogy between the links in a graph and forces acting between physical objects, to render graphs. Typical algorithms of this kind consist of two parts: (1) defining a *model for the force system* that will be used to assign an interaction strength between nodes of a graph, and (2) choosing a technique for finding an equilibrium state of the force system. The final rendering is achieved by allowing all the nodes of the graph to relax to their equilibrium positions that minimize tension throughout the (physically realized representation) of the abstract graph; i.e. a position in which the sum of the notional forces on each node is zero.

The simplest force-directed methods use a combination of *spring* and *electrical* forces: links are modeled as *springs* (with forces f_{ij} exerted on node j by the spring between i and j), and nodes are charged particles that attract or repel one another with force g_{ij} (whose exact form depends on the criteria that are appropriate for a given problem). The total force on node j is then defined by [97]:

$$F(j) = \sum_{i|(i,j) \in E(G)} f_{ij} + \sum_{i \in V(G)} g_{ij} . \quad (6)$$


Typically, f_{ij} follows *Hooke's* law (i.e., f_{ij} is proportional to the difference between the distance between i and j , D_{ij} , and the zero-length of the spring, $l_{ij,0}$), and g_{ij} obeys an *inverse-square* law. If the graph is to be rendered in a two-dimension Euclidean space, for example,

$$\bar{F}(j) = \sum_{i|(i,j) \in E(G)} s_{ij} \cdot [D_{ij} - l_{ij,0}] \cdot \frac{(\bar{x}_j - \bar{x}_i)}{D_{ij}} + \sum_{i \in V(G)} \frac{e_{ij}}{D_{ij}^2} \cdot \frac{(\bar{x}_j - \bar{x}_i)}{D_{ij}}, \quad (7)$$

where $\bar{x}_i$ and $\bar{x}_j$ are the positions of the i^{th} and j^{th} nodes, respectively; s_{ij} is the *stiffness coefficient* of the notional spring between i and j (so that larger values tend to more strongly compel D_{ij} to be near $l_{ij,0}$); and e_{ij} represents the *strength of the electrical attraction* (or *repulsion*) between i and j . The “aesthetics” of the rendering are controlled by the three parameters $l_{ij,0}$, s_{ij} , and e_{ij} .

Figure 9 shows an example of applying each of these embedding algorithms to a random graph consisting of 25 nodes and 49 links. The graph has no distinguishing features, and is chosen for illustrative purposes only.

Figure 9. Sample renderings of the *same* (order 25 and size 49) random graph using the five visualization algorithms discussed in the text




The *ranked* embedding ranks the nodes according to their distances from the first four nodes (drawn at left-hand-side of the graph shown in the figure; note also that the *labels* do not appear). Both the *radial* and *rooted* embeddings use the *center node* of the graph as the root, the distance to which is used to rank the other nodes. The “center” consists of all nodes with minimum *eccentricity*, where the eccentricity of a node i is defined as the maximum distance from i to any other node in the graph; intuitively, the center consists of that part of a graph that is closest to the rest of the graph. Finally, the force-directed embedding shown here uses only the spring force; i.e., the notional electrical forces between nodes are ignored.

Of the five embeddings introduced here,¹⁴ force-directed methods come closest to embodying what might loosely be called an “objective

aesthetic”; that is, an aesthetic that derives from an explicit map between abstract topology and physical (i.e., graph rendering) space. If employed wisely, in a way that respects those aspects of the relationships of a system that are of greatest interest to the researcher, force-directed algorithms can also be used as powerful pattern-recognition tools to help identify otherwise invisible structures and substructures within a graph; particularly when the graph is large and/or complex.

Figure 10. An example of graph visualization using spring-embedding



Consider, as a real-world example of the utility of spring embedding, the graphs shown in figure 10. Figure 10-a shows a circular embedding of the social ties of a homeless woman [106].¹⁵ The points on the circle represent various social contacts and the lines of various thickness represent the strength of ties. The existence of any latent patterns or relationships is not immediately apparent. Figure 10-b shows the same network, but is rendered using *spring-embedding*. The existence of subgroups is now obvious.

Force-directed methods are particularly adept at revealing hidden structures when the spring (or electrical) forces are not just (otherwise arbitrary) analogues of physical forces—defined by, and used solely for, generating geometrically pleasing renderings—but include

14. Many more graph embedding algorithms are available in the literature: see [97]-[99] and [102]-[105].

15. The node labeled “Respondent” is the homeless woman that was the focus of the study.

components that depend on the *information* content of the (nodes and/or links of the) graph.

For example, the “force” between nodes may be generalized to include a term that is a function of how “similar” the nodes are with respect to some selected subset of an internal feature space; or a function that depends on the degree of mutual interest, cooperation and/or motivation to achieve certain goals (if the nodes are “agents” in a social setting); or a component may be added that depends on various local and/or global information processing metrics, such as *centrality*, *betweenness*, *efficiency*, and *symmetry* (network metrics are defined in a later section.) Using such generalized information-driven, force-directed renderings, otherwise latent relationships among a system’s “hidden variables” that are *a priori* cloaked behind a jumble of unsorted, overlapping links, may be revealed.


Complex networks: zoology

“Far away in the heavenly abode of the great god Indra, there is a wonderful net that has been hung by some cunning artificer in such a manner that it stretches out infinitely in all directions...the artificer has hung a single glittering jewel in each ‘eye’ of the net, and since the net itself is infinite in all dimensions, the jewels are infinite in number...If we now arbitrarily select one of these jewels for inspection and look closely at it, we will discover that in its polished surface there are reflected all the other jewels in the net, infinite in number. Not only that, but each of the jewels reflected in this one jewel is also reflecting all the other jewels, so that there is an infinite reflecting process occurring.”—Avatamsaka Sutra (F.H. Cook)

Graph space

Consider the set of all possible *undirected graphs*; that is, graphs containing only symmetric links between originating and terminating nodes. Figure 11 shows, schematically, the spectrum of graphs belonging to this set, in which the number of nodes, N , is fixed, and the graphs are arranged—from left to right—according to their number of links, M .

Figure 11. Schematic of the spectrum of all possible graphs, $G(N,M)$, of size N and order M



The spectrum starts with the totally disconnected graph on the left, and ends with the totally connected graph on the right. The red-col-

ored graph, $G(N,M)$, that appears in the dotted box near the center of the spectrum, represents some particular graph, with N nodes and $0 \leq M \leq N(N-1)/2$ links.

Implicit in the figure, is the fact that some *set of graphs*—the member graphs of which are not necessarily adjacent in the spectrum (since the differences among them are likely to depend on a significantly greater number of features than simply the number of links they contain)—*best performs some associated global function* (or set of functions). A graph—any *single* graph—is but one abstract exemplar of a typically much larger set of possible graphs, each of which represents the topology and information processing capabilities of a real physical system. The problem is to identify the core set of features that a graph must have in order to best perform its task.


For example, if the graph represents the social network of a company that produces “widgets,” G ’s associated function is to produce widgets; or, if the graph represents the ties of a terrorist organization, G ’s function is perform acts of terror. The (possibly myriad) set of intermediate *local steps* that are required to accomplish these global tasks, both by the real systems and the components of whatever graph is used to model them, are purposefully left unspecified. Indeed, the *raison d’être* for developing a dynamic random graph model of any complex adaptive system—in which intelligent agents are substituted for conventionally static nodes—is to provide the *graph* with a sufficient set of local decision rules, along with an evolutionary algorithm that allows the graph to learn, so that it may “discover” the steps necessary to accomplish its goal.

The two fundamental problems for the research analyst can be stated formally as:

Given a graph G that is a model of a physical system tasked with performing a global function F , and that G performs that function with “efficiency” $0 \leq e_G(F) \leq 1$, (1) identify the set of primitive local and/or global topological and information-processing properties of G , $\{p_i(G)\}$, that are most relevant to describing G ’s behavior (in the context of G performing function F), and (2) determine the set of values of G ’s primitive properties that maximize the value of $e_G(F)$.

Figure 11 already shows a decomposition of graph space according to one global metric; namely, the number of links, M , in a graph. Unfortunately, knowing only M yields little or no information about a graph; in the sense that M alone is generally a poor predictor of the efficiency with which G will perform function F . Thus a much broader set of both local and global properties of graphs is needed to understand the relationship between their structure and function.

Figure 12. An overview of the zoology of graphs, partitioned according to *heterogeneity*, *randomness*, and *modularity* (taken, partially, from [107]); see text for details



Zoology of graphs

A schematic of the zoology of graphs, in which graphs are characterized according to each of three measures, appears in figure 12; the figure is taken partly from a recent review of the properties of information flow in complex networks by Sole and Valverde [107]. A given graph is specified—and thereby physically positioned in the three-dimensional “zoo”—according to the degree to which it is either *heterogeneous*, *random* (in the Erdos-Renyi sense), and/or *modular*.

While even this broader partitioning is still far from complete, it illustrates some basic ways in which graphs may be distinguished, both on the local and global levels. We will make frequent references to this figure throughout the ensuing discussion, as we introduce the metrics that are depicted schematically in it. Examples of graphs that are *least* random, *least* modular and *least* heterogeneous, located at the bottom left of the graph space in the figure, are lattice meshes and r -regular trees. The nodes of these graphs are all alike, in that they all share essentially the same local neighborhood structure.

As more and more random links are injected into these homogenous nets, via the Erdos-Renyi RG model, the local neighborhood structures obviously assume a random structure; but there is also a global regularity that emerges: the *degree distribution*, $P(k)$, is strongly peaked at degree, $k = \langle k \rangle$, and decays exponentially with k .

As RGs become more and more modular (and we move toward the lower right of the graph space in figure 12), the graphs occupying this volume of the space are characterized by a power-law behavior of their *clustering coefficient*, $C(k) \sim k^{-\delta}$. $C(k)$ (defined more precisely below), measures the density of connections in the local neighborhood of a node.

An important set of graphs, called *scale free* networks (which were mentioned earlier during our discussion of RGs), is located near the top left portion of the graph space in figure 12, and is characterized by a power-law behavior of the degree distribution: $P(k) \sim k^{-\gamma}$. A scale free network is *extremely inhomogeneous*: while the majority of nodes typically have only one or two links, a few nodes have a large number of links (thus also assuring that the network is fully connected).

The colored graph that appears in figure 12 at the bottom right of the plot of $\log P(k)$ vs $\log k$ (taken from [108]) illustrates this property by coloring *red* the five nodes with the highest number of links, and coloring their closest neighbors *green*. While in the sample Erdos-Renyi RG (i.e. the colored graph that appears in the center of the bottom row of figure 12) only 27% of the nodes are nearest neighbors of the five most connected nodes, in the scale-free network more than 60% are. (Note that both networks contain the same number of nodes and links, 130 and 430, respectively.) This also illustrates the important roles that *hubs* play in scale-free networks; a property noted earlier.

In what region(s) of graph space do TNs live?

A natural question to ask about terrorist networks—viewed as mathematical graphs—is “*What region of the abstract graph space do they occupy?*” While figure 11 frames the question, suggestively, by pointing to a (red dashed-line) region, the size and location of that region (or multiple regions, as there is no *a priori* reason to expect the space of viable terrorist structures to constitute a connected subspace in the space of all graphs) is unknown. It is with this basic problem in mind that SOTCAC is being designed.

Reserving a more detailed discussion of the dynamics of terrorist net structure and formation until a later section (see **SOTCAC**), we will here only comment that it is precisely because terrorist organizations *self-organize*—rather than come into existence, all at once, according to the decree of a *de facto* master central plan—that we can use dynamic graph models to help us identify the relevant emergent features of terrorist networks. Terrorist nets evolve according to their (possibly changing) operational needs, the motivations, skills and personality of their members, and various internal and external constraints (including their coevolution with antiterrorist and counterterrorist organizations).

One can guess that terrorist nets will be neither completely random nor completely regular; but to exhibit local and global regularities that reflect the *ethnic*, *geo*-, *political*-, and *religious*-arenas in which they are spawned and within which they grow, adapt, and evolve. We can expect terrorist nets to be strongly modular, and partly hierarchical. And, to the extent that the Barabasi-Albert RG model has shown “preferential-link” attachment as a necessary ingredient of real-world scale-free network formation, we can expect large terrorist networks to be scale-free; though the degree to which a terrorist net’s major hubs constitute a true vulnerability of the system as a whole depends on the specific network.

In **Appendix 2: Al-Qaeda** we show how social network analysis can be applied to publicly available information to generate snapshots of part of the Al Qaeda network’s structure prior to September 11, 2001.

One of the most important conceptual problems that must be solved, before actually developing any multiagent-based, dynamic graph model of terrorist $\leftrightarrow$ counterterrorist network coevolutions, is to determine what the appropriate set of *measurable* topological properties is by which the behavior of terrorist networks may be described and characterized mathematically.

Random graphs

The theory of random graphs (RGs) was developed principally by Erdos and Renyi in the late 1950s, and offers a mathematical framework for posing questions about the general topological structures of computational systems and the expected behavior of certain random dynamical systems [79].

Although there are many alternative models of RG evolution (three of the most popular of which are introduced briefly in this section), and different models exhibit unique properties, the essential idea is to think of a graph as a *living organism* that develops by acquiring more and more links. The problem for the researcher is to explore the manner in which the evolving graph's local and global topological features emerge (usually in a dramatically abrupt and unexpected fashion).

RG theory is important in the context of this paper for two reasons: (1) because it provides the basis for describing the *evolution*—and not just the structure—of graphs, and (2) because it provides the foundation on which the more general *agent-mediated evolution* that lies at the core of SOTCAC itself depends. While RGs typically evolve probabilistically, according to tunable probabilities of adding nodes and links, SOTCAC's networks evolve according to local rewiring decisions made by “intelligent” nodes, using information that consists, in part, of the decisions made by other neighboring nodes, and, in part, of local topology (see **SOTCAC: a conceptual model**; page 121)).

Erdos-Renyi random graphs

In their classic 1960 paper, Erdos and Renyi consider a sample space in which each realization of the $\binom{N}{2}$ order- N labeled graphs $G(N, M)$ of size M are equiprobable. An alternative, and equivalent, definition

of an RG (called the *binomial model*), uses a sample space in which the probability of $G(N,M)$, P_G is given by:

$$P(G) = p^M (1-p)^{\binom{N}{2}-M}, \quad (8)$$

where p is an independent probability, $0 \leq p \leq 1$, and is typically taken to be a function of N .


The first model can be compared to the *microcanonical ensemble* in physics, in that it depends on the direct enumeration of possible configurations of a given graph. The alternative model, with its tunable probability, resembles the *canonical ensemble*. As for physical systems, these two models give the same results in the limit as $N \rightarrow \infty$. Depending on the problem, the sample space can be generalized to include loops, multiple edges or even ensembles of graphs embedded in other topologies. In all cases, the fundamental problem of RG theory is to determine the point at which (as defined either by $M(N)$ or $p(N)$, depending on the sample space being used) a particular property of a graph almost certainly appears. The main contribution of Erdos and Renyi, apart from formulating the RG “problem” itself, is the demonstration that many important structural properties appear quite suddenly, usually at well-defined critical points (i.e., at well-defined values of either $M_c(N)$ or $p_c(N)$).

In order to show the emergence of some specific property, Q , of a graph G , a probability space of graphs must first be defined, in which the probability that G has this property is positive. Erdos and Renyi used the criterion that almost every G has Q if the probability of having Q approaches 1 as $N \rightarrow \infty$.

In fact, as more and more edges are added to G (or the probability of adding an edge is systematically increased in the binomial model), G undergoes a *succession* of abrupt structural changes. For example, if $M = cN$ in Erdos’s and Renyi’s original formulation (or $p(N) = 2c/N$), then a “phase transition” occurs at $c = 1/2$: for $0 < c < 1/2$, G consists entirely of isolated points or small trees of order at most $\sim N \log(N)$; at $c = 1/2$, these small trees are suddenly replaced by a single tree spanning $\sim N^{2/3}$ vertices. Another dramatic structural change occurs

when $M = (1/2)cN\log(N)$ (or $p(N) = c\log(N)/N$): while G is almost always disconnected for $0 < c < 1$, for $c > 1$, G is not only almost always connected but is Hamiltonian as well.

Figure 13. A schematic illustration of several important epochs during the evolution of a random graph, in which links are added with a probability $p(N) \propto N^{-x}$ (N is the fixed number of nodes in the graph; see text for discussion;



A somewhat more complete summary of the sequence of changes that takes place during G 's "evolution" appears, schematically, in figure 13.¹⁶ The chart traces several distinct epochs in the evolution of an Erdos and Renyi RG, in which the link probability scales as $p(N) \propto N^{-x}$, where x is a tunable parameter.

The figure shows that for x greater than $3/2$, almost all graphs consist only of isolated nodes and links. As x passes through $3/2$, suddenly trees or order 3 appear. Trees of order 4 appear as x passes through $4/3$. As x gets closer and closer to 1, trees of increasing order start appearing. However, as long as $x < 1$, the (possibly disconnected) graph consists only of trees; in particular, no cycles are yet present.

As x passes through 1, cycles appear spontaneously; in the sense that the asymptotic probability for having a cycle of all orders approaches *one* (as $N \rightarrow \infty$). Complete subgraphs of order 4 appear for the first time only as x passes through $2/3$; of order 5 as x passes through $1/2$; and so on. As x approaches 0, the asymptotic probability that almost all RGs of order N approaches *one*.

Since the individual degrees of an Erdos-Renyi RG are distributed binomially, it follows that a random variable x_k , that represents the number of nodes with degree d , is distributed asymptotically according to the *Poisson distribution* [87]:

16. Figure 13 is adapted from the illustration that appears on page 11 of [88].

$$\text{Prob}[x_r = d] \sim \text{Poisson}(\lambda_k) = \frac{\lambda_k^d}{d!} e^{-\lambda_k}, \quad (9)$$

where $\lambda_k = \binom{n-1}{k} p^k (1-p)^{n-1-k}$ and d is any integer.

Small-world random graphs

In their seminal 1998 paper, Watts and Strogatz [83] formalize the notion of a *small-world graph* by showing that the topology of many real networks is neither completely regular nor completely random (as are Erdos-Renyi graphs).¹⁷

Like regular graphs, small-world networks exhibit a high clustering coefficient ($=C$);¹⁸ and, like random graphs, small-world networks have a small characteristic path length, L (= average distance between two arbitrary nodes in the graph). More precisely, small-world graphs of order N have diameter $\Theta(\log N)$, which means that between any two nodes there exists a path of size $\Theta(\log N)$.

Watts and Strogatz [109] propose a two-step algorithm for generating random small-world networks:

1. *Start with an ordered graph, and*
2. *Randomize it.*

The algorithm works, essentially, by randomly rewiring a circulant (i.e., $2K$ -regular) “seed” network. A circulant network is a graph such that each of its N nodes is linked to $2K$ of its nearest neighbors (or to K of each neighbors on its right and left sides when the graph is drawn on a circle).

17. The specific examples cited by Watts and Strogatz include collaboration social-networks of film actors, the neural network of a neotode (the *C. Elegans*; neotodes are a form of microscopic worm, and are among the most numerous multicellular animals on earth), and the US electric power grid.

18. The *clustering coefficient* measures the average “cliquishness” of a node within the graph (or subgraph), or the degree to which a graph is modular (i.e., is organized in a hierarchical fashion); see **Complex networks: metrics**.

The Watts-Strogatz algorithm proceeds from a circulant network by selecting a node i and a link $e(i,j)$ connecting it to another node j . The link $e(i,j)$ is then rewired, with probability p , by replacing j with a random node s (with multiple edges being forbidden). This probabilistic rewiring is repeated for each node in the graph. A second round of rewirings is performed in the same manner, but on links between second nearest neighbors; followed by a third round on nodes a distance $D = 3$ apart; and so on, for a total of K rounds. In effect, the algorithm provides an interpolation between a regular lattice (when $p = 0$) and a random graph (when $p = 1$), without changing the average number of neighbors.

In practice, the algorithm described above can generate disconnected graphs. To prevent this from happening, Watts and Strogatz later added the constraint $1 \ll \ln(n) \ll 2K \ll n$, where $2K \gg \ln(n)$ ensures that the resulting RG is connected [87].

Figure 14 shows an example that uses a “seed” circulant graph consisting of 16 nodes and $K = 3$.

Figure 14. An example of using the Watts-Strogatz algorithm to “rewire” an $(N=16; K=3)$ circulant graph into a small-worlds graph; adapted from [71]





Figure 15 shows, on a semi-logarithmic scale, the typical behavior of C and L as the values of p range from about 10^{-4} to 1 [71]. Initially, as might be intuitively expected, C and L both decrease as the amount of rewiring increases, albeit at different rates: as long-range short-cuts are added, the ordered lattice only slowly loses its innate clustering,

while L drops off at a much faster rate. The most important structural changes occur during this early phase; by the time the first ten nodes are rewired, a graph is generally already indistinguishable from a random graph.

Figure 15. Typical decays of characteristic path length ($=L$) and clustering coefficient ($=C$) using the Watts-Strogatz small-world random graph model; after [110]



The shaded region in figure 15 represents the “small-worlds” regime in which graphs have large C , but small L . These are the graphs that have both a high degree of local clustering, and many shortcuts between clusters; they possess the short average path lengths that Milgram [68] found in his social networking experiments and associated with other “six degrees of separation” phenomena [72].

Scale-free random graphs

Scale free networks are characterized by a power-law behavior of their degree distribution: $P(k) \sim k^{-\gamma}$, where $P(k)$ is the probability that a node has degree k , and γ is the power-law scaling exponent. A scale free network is thus *extremely inhomogeneous*.

The discovery that many real-world graphs have a power-law degree distribution (by Barabasi [48], Faloutsos [111], and others) is of particular fascination to physicists, because power laws play a fundamen-

tal role in statistical mechanics. It is well known, for example, that a physical system that is poised exactly at a phase transition between an ordered and random state exhibits “self-similar”—or scale-free—properties; i.e., it will appear the same whatever scale it is examined at. Mathematically, if we take f as our functional “probe” of the scale-free system, this means that at the phase transition, f must be such that $f(ax)=bf(x)$, where a is a scale parameter, and b is a constant. The only solution to this equation is a power-law: $f(x) \sim x^{-\gamma}$. Graphs with a power-law degree distribution are for this reason called *scale free*.

While the majority of a scale-free network’s nodes typically have only one or two links, a few nodes have a large number of links. Such networks are called “scale free” because the ratio of the number of very connected nodes to the number of nodes in the rest of the network remains constant as the network changes in size. Compare this to, say, Erdos-Renyi RGs, in which one expects to have few, if any, strongly connected nodes (or to have, at most, so few as to be statistically insignificant); most nodes in an Erdos-Renyi RG typically have a number of links near a small, average value. Moreover, as a Erdos-Renyi RG evolves, the relative number of very connected nodes decreases.

Note that while *both* Erdos-Renyi graphs and scale-free networks can possess the small-worlds property, only in scale-free networks is it probable that one or more small-worlds-like “short” paths will pass through one of the network’s highly-connected hubs. This has significant ramifications both for how the system behaves and for how the system tolerates outside intrusion and/or a targeted attack.

Randomly connected nets tend to degrade steadily and predictably, slowly losing their connectivity as nodes (or links) are either randomly or selectively removed; eventually they break apart into multiple disconnected subgraphs. While scale-free networks may not *initially* suffer any ill effects from random node removal—since it is statistically unlikely that any key (i.e., highly connected) nodes will be removed by a *random* attack—if the attack on the network is intelligently focused on its most important hubs (which are, typically, relatively few in number), the results may be catastrophic. One needs to remove only a few of a scale-free network’s most vital hubs to cripple the entire system.

Among the many real-world networks that have been shown to be scale-free in the literature are power grids, the connectivity of the internet (in which the nodes are computers and the links are the various physical and wireless connections and routers) and WWW (in which the nodes are web pages and links are the hyperlinks connecting them), social networks of collaboration among researchers (in which the nodes are scientists and links exist only between coauthors), research citation networks (in which links are citations among papers), metabolic networks and protein regulatory networks, and the dispersal networks of sexually transmitted diseases [80]. The crippling effects suffered by the US airline system, financial markets and telecommunications networks as a direct result of the 9/11 terror attacks testify to the vulnerability of critical hubs.

Scale-free behavior does *not* appear in either Erdos-Renyi random graphs or in the Watts-Strogatz small-world model. Both models actually yield $P(k)$ distributions that are highly peaked around the average degree [46]. Barabasi and Albert [88] have proposed an alternative algorithm to evolve random scale free networks that relies on two critical ingredients:

1. *Growth*, so that new nodes and links are continually added to the system (much in the same way as, for example, the WWW grows exponentially in time by continuous creation of new web pages and hyperlinks among new and old pages); and
2. *Preferential link attachment*, so that new nodes are more likely to be connected to existing nodes with high connectivity (i.e., new links are assigned via a “rich-gets-richer” rule). As an example of preferential attachment, consider again the WWW, in which new webpages are more likely to include links to sites (and documents) that are already well known and have high connectivity.

The basic algorithm proceeds by starting at time $t=0$ with a small initial graph $G=G(N_0, M_0)$. At each time step t , one adds a new node n_t to which are assigned m new links (to already existing nodes).

The probability that n_t is added to a node $n'_{t-1} \in V_{t-1}$ (where V_{t-1} is the set of nodes in the graph at time $t-1$) is proportional to the degree of the *old*-node candidates:

$$\text{Prob}[(n_t, n'_{t-1})] = \frac{\deg(n'_{t-1})}{\sum_{v \in V(G)} \deg(v)} \cdot 19 \quad (10)$$

As the Barabasi-Albert graph, G_p , evolves, the number of nodes at time t is $|V_t| = |V_0| + t$ and the number of links grows as $M_0 + mt$. Barabasi and Albert [88] show that G_t has an asymptotic degree distribution that obeys a power law with $\gamma = 2.9 \pm 0.1$ and is independent of t .²⁰

Crucitti, *et al.* [113] use a measure called *efficiency*²¹ to analyze the error and attack tolerance of random networks, and find that scale-free networks display both a high degree of (local and global) error tolerance and an extreme vulnerability to attacks that target crucial nodes.

Mossa, *et al.* [114], consider another generalized model of scale-free networks that is particularly relevant for studying the growth of real-world networks such as terrorist networks; their model assumes that nodes can process information about only a *subset* of the existing

-
19. Barabasi's and Albert's preferential link attachment rule is manifestly *linear*: a new node is, say, twice as likely to be connected to an existing node that has twice as many connections as its neighbor. Krapivsky and Redner [49] show that using link-creation probabilities that scale nonlinearly with the degree of *old*-node candidates results in networks whose degree distributions deviate from the generic power-law behavior observed in scale-free networks. Moreover, links are often added to real networks between existing nodes, or nodes and links can disappear. Dorogovstev and Mendes [64] show that if one adds to the dynamics elements of both *creation* and *deletion* (so that nodes and links can both appear and disappear as a network evolves, as they do in the SOTCAC model), then the presence of such events can modify the power-law exponent, γ , so that it can assume essentially any value between one and infinity.
 20. Barabasi, Albert and Jeong [112] use a mean-field theory approach to argue that $\gamma \sim 3$. Having demonstrated that growth and preferential attachment are sufficient to generate scale-free networks, Barabasi and Albert also show that *uniform* attachment does not lead to a power-law degree distribution, but rather to an *exponential* decay: $\text{Prob}(k) \sim \exp(-bk)$ [88].
 21. *Efficiency* is introduced in a later section (see discussion in **Complex networks: metrics** that begins on page 94), and is used by SOTCAC.

nodes in the network. The authors find that the distribution of the number of incoming links to a node decays as a power law with an exponential truncation controlled by two factors: (1) system size, and (2) the subset of the network that is accessible to the node.

The Watts-Strogatz small-worlds model and the Barabasi-Albert scale-free network model are both important because they emphasize that nonuniform structure in the connection topology of complex networks. Unlike Erdos-Renyi RGs, in which nodes and links are essentially homogenous, the small-worlds and scale-free models begin by assuming, and respecting, the fact that some nodes and links in real networks are fundamentally more or less important than—and exert unequal influence over— other nodes and links.

Clustered scale-free networks

While the Barabasi-Albert scale-free network model reliably generates graphs with small average characteristic path lengths and degree distributions obeying a power law, it lacks the high clustering found in many real social networks [115]; in other words, the Barabasi-Albert does *not* yield small-world networks. However, two alternative algorithms have recently been introduced that effectively produce hybrid scale-free/small-world networks; that is, scale-free random networks that exhibit both small average path lengths and strong clustering:

- *Triad-creation*, and
- *State-activation/deactivation*.

Triad-creation

Holme and Kim [116] modifies the basic Barabasi-Albert algorithm by adding a “triad creation” rule:

If a link between n_1 and n_2 was added during the previous step of preferential attachment, then add a link from n_1 to a randomly chosen neighbor j of n_2 . This forms a triad; i.e., a local clique of three vertices linked to each other. If there are no available nodes to link with (within the set $G(n_2)$) then perform preferential attachment as before.

For each new node, the triad creation rule is applied with a probability P_t and a preferential attachment rule is applied with probability

with the probability $1 - P_t$. The average number of triad creation trials per added node is $N_t = (m - 1) * P_t$ is used as a control parameter in the *clustered scale-free* network model. The virtue of the model is that while it introduces a desired (and tunable, via the control parameter N_t) level of clustering, it preserves the scale-free degree distribution, and other properties, of the original Barabasi-Albert model. (Note that the clustered scale-free model reduces to the original Barabasi-Albert model in the limiting case of $N_t = 0$).

State-activation/deactivation

Klemm and Eguiluz [117] have recently introduced another variation of the Barabasi-Albert algorithm that produces scale-free degree distributions and maintains the small-world properties of small average distances and strong clustering.

The algorithm depends on an assignment of a binary state variable (that represents, say, an *active* or *inactive* state) to each of the nodes of a developing network. A random graph is “seeded” with a completely connected graph of M *active* nodes. The three-step Klemm-Eguiluz algorithm is then iterated for t steps:

1. *Growth*: A new node with M links is attached to the network.
2. *Preferential Attachment*: Each of the M links of the new node connects either to one of the active nodes (with probability $1 - \mu$) or to a random inactive node (with probability μ). If the latter, the random node is selected according to Barabasi-Albert’s preferential link attachment (i.e., the probability that a node obtains a link is proportional to the node’s degree; see equation 10).
3. *Node activation and deactivation*: One of the active nodes is *deactivated*, with probability $P_{deac} = k_i^{-1} / \sum_j k_j^{-1}$. The new node is made *active*.

The Klemm-Eguiluz algorithm generates scale-free random networks with degree distribution $P(k) = 2M^2 k^{-3}$ for $k \geq M$ and average connectivity $\langle k \rangle = 2k / N = 2m$ [117]. Moreover, by varying μ (between the values zero and one), the algorithm permits one to study the transition between cases with large path length and clustering [118] and

cases with small path length and clustering (note that for $\mu=1$, the Klemm-Eguiluz algorithm reduces to the Barabasi-Albert model). Klemm and Eguiluz [117] show that only a few long-range connections (which arise as soon as μ becomes nonzero) are needed to induce a small-world transition. In fact, when their algorithm is used with any $0 < \mu \ll 1$ (say, $\mu = 0.1$), the resulting graphs exhibit all three basic properties of real-world complex networks: *scale-free degree distributions*, *small characteristic path lengths*, and a *high degree of clustering*.

Search in complex networks

As mentioned in the introduction to this section, social networks possess not one but *two* surprising fundamental properties. The first property is the ubiquity of short paths between any two nodes (i.e., the “small-worlds property”), which is discussed earlier and which is reproduced using the Watts-Strogatz algorithm.²² The second property is the fact that individuals belonging to these networks are able to *determine* what these short paths are. What is surprising is that short paths not only exist, but that individuals can find them using only *local information*. Since one can easily imagine networks in which short paths exist (and can be recognized as such, *globally*), but in which there is insufficient local information to help guide searches for them, the understanding of why real social networks are easily navigable cannot be derived simply from knowing that they possess the small-worlds property.²³

In Milgram’s experiment, individuals found short paths to other distant people—at least, occasionally—by pursuing a greedy strategy whereby they passed the message to whichever one of their immediate acquaintances they felt was nearest to the intended target recipient. Contrast this to all Erdos-Renyi-type random graphs: although they may also (by chance) contain short paths, since all search algorithms are effectively equivalent to performing random walks on the network, it is impossible to construct a local algorithm.

22. See section **Small-world random graphs**; page 49.

23. The distinction between possessing short-paths and being able to find them was first made by Kleinberg [78] and (in the context of searching the World Wide Web) by Albert, Jeong and Barabasi [112].

Navigability Problem

The social network *navigability* problem consists of ascertaining the properties a network must possess to ensure that arbitrary pairs of agents in the network are able to find short paths of acquaintances that link them using only local information. A related question is, *given* a network, and a resource R located somewhere on that net, what is the “optimal” local search strategy to use to find R , starting from an arbitrary node?²⁴ These already difficult problems become even more difficult in the event that the underlying topology of the network being searched itself changes as a function of the search. For example, the structure of the World Wide Web constantly changes as its pages and links adapt to searches. A popular method used by Web search engines (such as *Google*, *Yahoo*, and others) is to first create a distilled local map of the Web, which is indexed—or “crawled”—by a search engine spider. Although this strategy is certainly effective—particularly so for networks that exact a heavy communication cost for real-time searches (the WWW is a prime example)—it is only loosely local, since it uses a local snapshot of global information.

Local Search

Kleinberg [119] was the first to study the navigability problem in the context of small-world graphs. Introducing an infinite family of network models that generalizes the Watts-Strogatz model, Kleinberg shows that—for *one of these models*—there exists a local, decentralized algorithm for finding short paths with high probability.

24. We introduce and briefly discuss the *navigability problem* in this section for two reasons: (1) the general problem of how to best “search” a network currently constitutes an important, and open, general research area in complex network theory; and (2) the success of SOTCAC’s notional terrorist-agents at finding mission-required resources scattered throughout the terrorist network (the state of which they have only a local understanding of), depends on their ability to “solve” the navigability problem on a dynamic topology; and therefore draws on the observations and algorithms that have been introduced in the recent literature. On an even more basic level, the ultimate goal of complex network theory is to understand the relationship between a given network’s *structure* and the *function* (or *functions*) that it performs. Viewing “navigability” as a basic network function, the fundamental question becomes, “*What properties must a net possess in order for it to be locally navigable?*”

Kleinberg's model combines a two-dimensional regular lattice with a number of long-range links. The distance between nodes i and j ($=\Delta_{ij}$) is taken to be the number of steps that separate them on the underlying lattice (i.e., independent of the set long-range links that are to be added). The set of long-range links is constructed randomly, with the probability that node s is linked with another node t , $\text{Prob}(I_{st} = 1) \propto (\Delta_{st})^{-r}$, where r is a tunable parameter. Kleinberg interprets his model by associating the geometry of the underlying lattice with the agents' "social space" (which therefore defines where they "live" and who they "know"); the parameter " r " roughly measures the extent of one's long-range network, with larger values yielding "long-range" contacts that are more tightly clustered in the vicinity of a given agent. number of outgoing links per node is fixed. While the underlying lattice obviously oversimplifies the complexity of real social spaces, the model nonetheless captures an essential ingredient of the navigation problem: *how does one minimize the routing time of a message, given only local information about the network's metric?*

Now, consider two arbitrary nodes in this network, a and b . The problem is to send a message from a to b that requires the minimum number of steps. Kleinberg assumes that each agent, after receiving the message, knows only three local facts: (1) the set of his own local contacts (on the underlying lattice); (2) the location (on the lattice) of the intended target, b ; and (3) the locations and long-range-contacts of the all agents that have previously received (and passed on) the message. In particular, it is assumed that no agent, at any time, has any knowledge of the long-range contacts of agents that have not yet received the message. With these assumptions and constraints, Kleinberg's search algorithm may be stated, simply, as follows:

Each node sends a message to one of its neighbors (either local or long-range) that is closer to the destination, in terms of lattice distance, Δ .

Kleinberg's major finding is that *only for $r = 2$* does this decentralized routing algorithm yield an expected delivery time of $\Theta(\log^2 N)$ on lattices of order N ; for all other values of r , the delivery time grows as N^b , for some $b > 0$.²⁵ (Similar results are found for cases in which the underlying lattice assumes a different topology; for example, when the underlying graph is a tree [119]).

Adamic, *et al.* [120], have recently introduced an alternative locally guided search algorithm, that exploits the statistical features of scale-free networks (see **Scale-free random graphs**, page 51). Aimed at optimizing search times on the World Wide Web and other, related, peer-to-peer distributed filesharing systems that do not have a central server (such as *Gnutella*²⁶ and *Freenet*²⁷), the algorithm generalizes the more conventional “breadth-first” search (in which each node, after receiving a query of the form, “*Do you have the information I need?*”, either answers “*yes*” and stops the search, or forwards the query to each of its neighbors) by having each node that does not have the required information send back the reply, “*No, but I have k neighbors.*” The original sender node sifts through its responses, chooses the neighbor with the highest k value, and passes the responsibility of finding the target to that neighbor (who then proceeds to take the same steps).

The algorithm is obviously tuned to work well on scale-free-like networks, because a significant fraction of the nodes in such networks are made up of neighbors of high-degree nodes. Therefore, on average, only a few iteration steps of the query must be processed until a node with a neighbor that contains the required information is found. Adamic, *et al.* have tested their algorithm numerically on a variety of real and simulated random networks, and have found that the behavior roughly matches the expected scaling behavior [121].

Unfortunately, neither of the two algorithms discussed above addresses the problem of how real networks *evolve* navigable struc-

25. This surprising result may be understood, heuristically, by noting that the value of r represents a trade-off between relying on long-range contacts and their utility in reducing the expected wait time. As r increases, Kleinberg’s decentralized algorithm makes increasingly frequent use of long-range contacts, but, at the same time, “long-range” contacts (which become more and more tightly clustered on the underlying lattice) are also less likely to improve on a routing that only makes use of the underlying lattice’s original links. The value at which this trade-off is “optimal” is the inverse-square distribution; that is, at $r = 2$. More generally, on d -dimensional underlying lattices, Kleinberg’s algorithm performs optimally at $r = d$ [119].

26. <http://www.gnutella.com/>.

27. <http://freenet.sourceforge.net/>.

tures. Kleinberg's result only characterizes the structure that a *static* network must have in order to be locally navigable (and assumes that the target node's location is known *a priori*). Adamic, *et al.*'s algorithm, while useful for decentralized file-sharing systems, merely exploits the scale-free property that peer-to-peer networks must already possess, *prior* to applying the search algorithm. In reality, of course, a social network's structure is never fixed, but is constantly changing as its members rewire their local neighborhood in an effort to "optimize" their (ego-centered) searches. Therefore, the deeper navigability problem requires an answer to two interrelated questions:

- "What local search strategies are optimal for a given network?," and
- "How are navigable topologies self-organized?"

Evolving searchable nets

A decentralized search algorithm that is both directly applicable to a certain class of searchable networks and offers a model to explain how this particular class of searchable social networks may arise naturally (constrained only by a set plausible assumptions of human social structures), has recently been introduced by Watts, Dodds, and Newman [122].

Earlier work by Killworth and Bernard [70,123] suggests that people tend to navigate networks by searching for any social characteristics that they perceive the target and their acquaintances to *have in common* (such as social status, occupation, or location). While in random networks, the overlap between a node's neighbors and a node's neighbor's neighbors is typically very small, social networks are very different: many of one's friends' friends are also one's friends. While the Watts-Strogatz small-worlds model shows that it takes only a few random links between people outside their local cliques to generate a sufficient number of small-paths to spawn a small-world network, the model did not explain how people are able to select just the right subset of acquaintances to form all the necessary links.

Building upon these ideas, Watts, *et al.* introduce a class of networks, the nodes (or agents) of which are all endowed with a set of social characteristics, and are assigned to groups that are hierarchically chained to larger groups. For example, an individual may be a


member of certain work group, that belongs to a certain organization, which in turn is part of a larger business network, and so on. The result is a tree-like hierarchy (which serves more as a conceptual scaffolding rather than a model of the actual network; see discussion below) that lends itself naturally to define a “social distance.”

Watts, *et al.* base their hierarchical network model on six propositions about real social networks [122]:²⁸

1. *Individuals are defined not only by their network ties, but by unique identities.* Groups are defined as collections of individuals all sharing some well-defined set of social characteristics. (Note that this proposition entails a bit of dynamic self-reference, insofar as the “characteristics” are derived from how individuals view themselves and how they perceive others in the context of their mutual association (vis-a-vis the group to which they all belong.)
2. *Individuals are grouped in (and therefore partition the world into) a hierarchy of categories* (see figure 16). The top layer of the hierarchy contains the entire network, and each successively deeper, and multitudinous, layer represents an increasingly finer distinction of features and groups. The bottom layer represents the set of individuals, each of whom effectively defines their own group.
3. *Social interaction (i.e., acquaintanceship) is a function of group membership: individuals are increasingly likely to know each other the closer they are within the hierarchy.* Watts, et al., assume that the probability that two individuals, A and B , are acquainted, $P(A,B)$, decreases exponentially with decreasing group similarity: $P(A,B) = k \exp\{-\alpha \cdot x_{A,B}\}$, where α is a tunable parameter,²⁹ k is a normalization constant, and $x_{A,B}$ is the “height” of the lowest level of the hierarchy at which the two people are linked.

28. Many of these properties are reminiscent of *holons*, introduced by Koestler in the 1970s [124]: “No man is an island- he is a holon. A Janus-faced entity who, looking inward, sees himself as a self-contained unique whole, looking outward as a dependent part. His self-assertive tendency is the dynamic manifestation of his unique wholeness, his autonomy and independence as a holon. Its equally universal antagonist, the integrative tendency, expresses his dependence on the larger whole to which he belongs: his part-ness.”

Figure 16. Schematic illustration of a single chain in Watts, *et al.*'s hierarchical social network model; in this example, $b (=2)$ is the *branching ratio*, $g (=6)$ is the *group size*, the *height* of the hierarchy is equal to *five*, and x_{ij} is the “similarity” between groups i and j ($=$ height of the lowest common ancestor level)



4. Each feature of social identity defines a separate hierarchy. This proposition supposes that social worlds are formed using multiple independent attributes, in the sense that proximity in one hierarchy (that is, say, defined by occupation) does not imply proximity in another (for example, defined by location). Watts, *et al.*, thus represent each node i by an H -dimensional *identity vector*, $\vec{v}_i$, where the h^{th} component is i 's position in the h^{th} hierarchy (or dimension).
5. The social distance between two nodes A and B , $R_{\text{Social}}(A,B)$, is the minimum “ultrametric” distance between the two nodes in all hierarchies:

$$R_{\text{Social}} = \text{Minimum}_h \{x_{A,B}^h\}. \quad (11)$$

29. The parameter α is best described, heuristically, as a tunable homophily index; i.e., it measures the degree to which nodes associate with other nodes that are “like” them. When $\exp(-\alpha) \ll 1$, all links are very short and individuals only associate with individuals that are most like them (that is, belong to the same bottom-dwelling group in the hierarchy); when $\exp(-\alpha) = b$, individuals are equally likely to associate with any other individual and the result is a random graph in which any two nodes are as similar or dissimilar as any other two randomly selected nodes. Homophily plays an important role as tunable parameter in SOTCAC, where it is used by notional terrorist-agents to tune the way in which they “rewire” their local communication nets.

Social distance thus effectively measures the degree to which two nodes (or agents) perceive themselves as being similar, where “similarity” is defined according to the social features that characterize the multidimensional hierarchies. R_{Social} captures the intuitive property that for two people to be socially “close” it is sufficient that they are similar in a single social dimension (such as when two individuals working for two different firms collaborate on the same project). However, R_{Social} is not a true metric function since it violates the triangle inequality: $R_{Social}(A,C)$ can be *large* even if $R_{Social}(A,B)$ and $R_{Social}(B,C)$ are both *small*.

6. *It is assumed that each individual node knows only the identity vectors of (1) itself, (2) its friends, and (3) the target node.* Moreover, it is assumed that individuals can estimate the social distance between their friends and the target.

Thus, nodes possess two complementary kinds of information: (i) *social distance*, which encodes (a processed form of) global information, but is not a true distance; and (ii) *network distance*, which is a true distance, but of which they only have a local view. It is important to note that neither kind of information—by itself—is sufficient for navigating the social hierarchy. However, Watts, *et al.*, show that the same greedy algorithm as used by the participants of Milgram’s experiment [68], provides an efficient means of directing messages throughout the network:


Each node i along the message chain forwards the message to its neighbor j who i perceives to be closer to—in terms of social distance, R_{Social} —the target node t .

Let P_{target} to be the probability that a random message—originating from a random node and targeting another random node—reaches the target. Watts, *et al.*, then define a *searchable network* as one in which $P_{Target} \geq r$, for any desired r .

They report two significant findings [122,125]:

- Searchable networks make up a broad region of the (a,H) parameter space; and, moreover, this regions corresponds to the choices in the model parameters that are the most sociologically reasonable; and

Figure 17. Comparison between the number of message chains, $n(L)$ of length L , as observed in Milgrim's original "small-worlds" experiment and numerical results derived from Watts, *et al.*'s hierarchical social network model using model parameters that are consistent with Milgrim's experiment; after [122]



- The distribution of message path lengths, calculated for parameter space values that are consistent with Milgrim's experiment (as well as subsequent variants), closely matches the empirically observed data (see figure 17).

As Watts, *et al.*, point out in their paper, while their model is sociologically based, its utility is considerably more general, as it applies to any decentralized network whose nodes are organized according to the same two primitive discriminatory components used in defining a social hierarchy; namely, an *identity* (that derives from a multifaceted feature vector) and *similarity* (which is defined by a suitably generalized multidimensional social distance).

For the purposes of this paper, Watts, *et al.*'s model provides several important ingredients for SOTCAC's conceptual design. As discussed in a later section,³⁰ one of the major tasks of SOTCAC's notional terrorist-agents (or *T-agents*) is to acquire resources that are distributed across the terrorist-link network. Since *T-agents* are innately *local* entities, all issues having to do with local search—such as determining the best search strategies, adaptively maintaining local searchability by

30. See **SOTCAC: a conceptual model**, page 121.

rewiring a node's local network structure, and assessing "closeness" and "similarity"—are direct analogues of the corresponding sets of issues as explored by Watts, *et al.*, using their hierarchical social network model. Their use of the ultrametric distance as a measure of "similarity" in a multidimensional social hierarchy is particularly relevant, since all of SOTCAC's *T-agents* are also defined by a multidimensional "personality" vector, the components of which describe such features as network *value* (or *rank*, as a terrorist), *function* (as a member of a given terrorist cell), and various social and/or communicative propensities and motivations. Watts, *et al.*'s, social distance—after a suitable "translation" into the context of an analogous, but considerably more complicated model—thus provides a natural starting point for incorporating a well-defined social metric into SOTCAC's dynamic terrorist network.

Network search and information theory


Sneppen, *et al.* [126,127] examine the navigability problem—that is, of how to best navigate a network using only local knowledge—by using *information theory* to describe the constraints a network's structure imposes on global communications. A byproduct of this work is that it implicitly provides an information-theoretic metric for identifying certain "key" nodes of a network (such as nodes that provide the best access to the system, and those that are the best "hidden"). SOTCAC needs metrics of this kind (many more will be introduced in the next section) because, in order to "discover" (and to determine the efficacy of) tactics and strategies to use against the notional terrorist network, SOTCAC's *counterterrorist network* needs to base its actions on well-defined measures of "value" that individual nodes (and groups of nodes) represent to the terrorist network. *Search information*, defined here, is one such objective measure.

Following Sneppen, *et al.* [126], suppose an agent is at node i and wants to send a message to node b (see figure 18). Assuming that the message follows the shortest path, $path(i,b)$ —or, in the event that there are multiple shortest paths, along a randomly selected path in some set, $\{path(i,b)\}$ —the probability that a given shortest path is followed is given by:

$$P\{path(i,b)\} = \frac{1}{deg(i)} \cdot \prod_{j \in path(i,b)} \frac{1}{deg(j)-1}, \quad (12)$$

where $deg(x)$ is the degree of the x^{th} node and the product is taken over all j nodes on the path excluding the start ($=i$) and end points ($=b$). The expression follows immediately from observing that, absent any information about where the target is, an agent first randomly selects one of the $deg(i)$ links available to it at node i , and then one of $deg(j)-1$ links at each subsequent node (taking into account the information gained by following the path up until that point; i.e., having no need to travel back along a previously traversed link).

Figure 18. Network schematic for calculating *search information*, *target entropy* and *road entropy* (after [126]); see text for details



The information that the agent needs to identify one of the degenerate paths between i and b is what Sneppen, *et al.* [127] call the *search information*, $S(i \rightarrow b)$:

$$S(i \rightarrow b) = -\log_2 \left[\sum_{path(i,b)} P\{path(i,b)\} \right], \quad (13)$$

where $\log_2$ is the logarithm base *two*, and the sum is over all degenerate paths between nodes i and b . A large value of S means that the agent needs to ask a large number of *yes/no* questions to find b ; a small S means that there are a large number of degenerate paths, and the agent will likely be able to easily find b .

Search information's practical utility is in providing an objective measure by which key nodes of a network may be identified. For example, suppose one is interested in finding the node that provides the best access to the entire network? From an information theoretic point of view, the node that has best access, $i_{best-access}$, is the one from which the total search information is minimal:

$$i_{best-access} = \underset{\{\text{nodes } i\}}{\text{Minimum}} \left\{ \mathcal{Q}_i = \sum_b S(i \rightarrow b) \right\}. \quad (14)$$

Similarly, if the problem is to identify the part of a network that is the most difficult to *find*—i.e., the best place to *hide*, b_{hide} —the answer is to determine the node that maximizes the search information:

$$b_{hide} = \underset{\{\text{nodes } b\}}{\text{Maximum}} \left\{ \mathcal{K}_b = \sum_i S(i \rightarrow b) \right\}. \quad (15)$$

A third measure related to search information is *predictability* [126]: the ability to predict local message traffic both to and across nodes in a network. More precisely, from the point of view of an agent at node i , the problem is to predict the neighboring node from which the next message will arrive. Without any advance information, all $\deg(i)$ incoming links are a priori equally likely and i requires on the order of $\log_2[\deg(i)]$ *yes/no* questions to guess the active link. On the other hand, one can expect fewer *yes/no* questions to be necessary if the agent has access to some information regarding local message traffic.

Sneppen, *et al.* [126,127] define two measure to quantify predictability: *target entropy* ($=T_i$) and *road entropy* ($=R_i$). Target entropy is defined as the entropy of the messages that are *targeted toward* a specific node i (see figure 18-a); road entropy is defined as the entropy of all message *across* a given node (see figure 18-b):

$$\begin{cases} T_i = \sum_{j=1}^{\deg(i)} c_{ij} \log_2(c_{ij}) \\ R_i = \sum_{j=1}^{\deg(i)} b_{ij} \log_2(b_{ij}) \end{cases}, \quad (16)$$

where both sums are over all neighbors of the j^{th} node, c_{ij} is the fraction of message that are targeted toward i that pass through j , and b_{ij} is the fraction of messages that pass through i that *also pass* through j .³¹ Heuristically, as the number of alternative pathways increases in a network, the search information is minimized, and both T and R will be large. Values of target entropy tends to be small for networks in which message traffic is concentrated into hierarchies; since, in such cases, most nodes are able to easily predict where the next message will come from. Similarly, values of road entropy tend to be small when the network consists of loosely connected tight clusters (or hubs); since, in this case, relatively many links are important. Sneppen, *et al.* [126] thus relate road entropy to an edge attack in a network, and target entropy to node attack [46].

Dynamic graphs

A “dynamic graph” is a network whose topology changes in time according to a set of nontrivial local and/or global “rewiring” rules. While all of the random graph models discussed thus far in this section are generated by “evolving” the structure (and sometimes the nodal composition) of some initial “seed” graph—and therefore also technically model *dynamic* graphs—the dynamical “rules” in all of these models are, physically speaking, trivial, and amount to little more than defining probabilistic additions or deletions of nodes and/or links.

The structure changing rules in RGs are generally not functions either of local topology or of any information that may be encoded within a graph’s transient local topology. The sole exception is the Barabasi-Albert scale-free model’s “preferential link attachment” rule, but even this rule merely constrains what remains a simple probabilistic evolution; the rule is introduced only to *refine* what remains a probabilistic growth, and whose aim is to achieve a power-law degree distribution in a random network. None of the random graph models discussed thus far depend in any way on a more general set of tunable

31. The quantity b_{ij} is proportional to a well-studied network metric called betweenness, that—roughly speaking—measures the extent to which a node mediates, or plays the role of “information broker” between, any two other nodes in a social network. It is introduced in the next section.

decision rules that, effectively, allow nodes to *rewire themselves locally*, according to their unique needs and motivations.

SOTCAC, on the other hand, is a *fully dynamic graph*: its global structure evolves according to rewiring decisions made by each of its nodes (which, in turn, base their decisions on local structure and information states). As in most multiagent-based models, probabilities are used by SOTCAC mostly as tunable randomizations of otherwise deterministic decision rules.

Structurally dynamic cellular automata

A conceptual precursor of SOTCAC is a model called *Structurally Dynamic Cellular Automata* (SDCA) [128]. In SDCA, the topological structure of a network is explicitly, and dynamically, coupled to the primitive information contained at each of the network's nodes. The coupling treats geometry and information on an approximately equal footing: the structure of the graph is altered locally as a function of individual neighborhood information states and geometries, while the topology is used propagate local information states according to conventional *cellular automata* (CA) rules.³²

In addition to providing a natural framework for analyzing the generation, transmission and interaction of topological disturbances in graphs, the SDCA model was the first model that allowed researchers to study properties of *self organizing geometry* (which is distinct and more general than, say, the analogues of geometry one finds in the self organized space-time patterns of cellular automata defined on fixed lattices). Applications range from simulations of crystal growth, to studying pattern formation of random cellular structures, to

32. CA are a general class of spatially and temporally discrete, deterministic mathematical systems characterized by local interaction and an inherently parallel form of evolution. First introduced by the mathematician John von Neumann in the 1950s as simple models of biological self-reproduction, CA are prototypical models for complex systems and processes consisting of a large number of simple, locally interacting homogenous components. CA are fascinating because very simple rules often yield highly complex emergent behaviors. In their most common form, CA “live” on *one*-, *two*-, or higher-dimensional Euclidean lattices; the lattices, themselves, however, *do not* evolve. See [94] for additional details.

endowing neural networks with a dynamic synaptic plasticity, to describing fluids in which the particle flow is coupled to a dynamical geometry (i.e. chemical self-assembly). Majercik [129] has studied SDCA as generalized models of computation, and shows that SDCA are actually more efficient “computers” than conventional CA; moreover, he defines a class of CA-universal SDCA that is capable of simulating any conventional CA of the same dimension. More recently, O’Sullivan [130] and Saidani [131,132] have introduced graph-based CA models: the former is applied to the study of urban dynamics, while the latter is introduced in the context of self-reconfigurable robots.

CA are usually defined on two-dimensional lattices, the sites of which are populated with discrete-valued dynamic elements, s_i , evolving under certain local transition functions. More generally, an N -node lattice may be characterized as an order- N undirected graph G , with adjacency matrix:

$$a_{ij} = \begin{cases} 1 & \text{if } i \text{ and } j \text{ are linked} \\ 0 & \text{otherwise} \end{cases}. \quad (17)$$

Using the graph metric function $Dist(i,j) = \text{Min}_P[\text{Number of links between } i \text{ and } j \text{ on path } P]$, a general CA transition rule f (i.e., the local algorithm by which the value of the i^{th} node at time $t-1$, s_i^{t-1} , is “updated” to the value s_i^t , at time t) that depends on a given node’s r neighbors may be written as follows:

$$\sigma_i^t = f \left[\left\{ \sigma_j^{t-1} \right\} \middle| j \in S_r^{t-1}(i) \right], \quad (18)$$

where $S_r^{t-1}(i) = \{j | Dist(i, j) \leq r\}$ is the radius- r graph-sphere at time $t-1$, centered on site i .


In the simplest case of elementary one-dimensional CA, for example, for which $r=1$, the graph-sphere about the i^{th} node consists of the node’s immediate *left* ($=s_{i-1}$) and *right* ($=s_{i+1}$) neighbors, and a typical function (for binary valued s ’s) is, say, $\sigma_i^t = \sigma_{i-1}^{t-1} \oplus_2 \sigma_{i+1}^{t-1}$, where $\oplus_2$ is addition *modulo-2*.³³ Of course, in a conventional CA, the underlying graph does not change in time, and so $S_r^t(i) = S_r^{t-1}(i)$ for all i , r and t .

SDCA models are defined by endowing G ’s adjacency matrix with a generalized form of (a conventional CA’s) transition rule:

$$a_{ij}^t = g \left[\left(\left\{ a_{kk'}^{t-1}, \sigma_k^{t-1} \right\} \right) \middle| k, k' \in S_r^{t-1}(i) \cup S_r^{t-1}(j) \right]. \quad (19)$$

In words, the value of a_{ij} , and therefore the decision to either *link*, or *sever an existing connection between*, nodes i and j at time t is a function of both (1) the local topology of the graph (as defined by the value of the graph's adjacency matrix at time $t-1$, $a_{kk'}^{t-1}$, and where “local” includes any nodes k within distance r from either i or j on the graph as it appeared at time $t-1$), and (2) the values of neighboring nodes (in a manner that is entirely analogous to what is usually done in conventional CA).

Figure 19. Schematic of partitioning a local neighborhood into three disjoint sets: the nodes i and j , a set of nodes denoted by A_{ij} , and another set of nodes denoted by B_{ij} ; see text for details



Because the core of SOTCAC's adaptive topology dynamics shares many basic features of the SDCA model, it is instructive to give an explicit example of an SDCA transition function. Toward this end, consider figure 19, which shows, schematically, a partition of a local neighborhood of an arbitrary graph G into three disjoint sets (that are to be used to define the value-component of the combined value/topology transition function): (1) the nodes i and j , (2) the set of nodes that belong to *both* i and j ($=A_{ij}$), and (3) the set of nodes in G that belong to *either* i or j , but not both ($=B_{ij}$). (Note that both A_{ij} and B_{ij} can also be generalized to sets of nodes that are a distance $Dist(i,j)=d$ apart from i and j .)

33. It is easy to show that, in the more general case where s takes on one of k possible values ($s = 0, 1, 2, \dots, k-1$), the total number of possible elementary one-dimensional CA radius- r rules for which is equal to $k^{k^{2r+1}}$ [94].

For example, carrying out the calculations explicitly for node 22 and pairs (18,23) and (17,23), we find:

$$\begin{cases} \sigma_{22}^t = f[\sigma_{17}^{t-1} + \sigma_{21}^{t-1} + \sigma_{22}^{t-1} + \sigma_{23}^{t-1}] = f[2] = 0, \\ \alpha_{18,23}^{t-1} = \sigma_{18}^{t-1} + \sigma_{23}^{t-1} = 1, \beta_{18,23}^{t-1} = \sigma_{13}^{t-1} + \sigma_{17}^{t-1} + \sigma_{19}^{t-1} + \sigma_{22}^{t-1} + \sigma_{24}^{t-1} = 3 \rightarrow a_{18,23} = 0, \\ \alpha_{17,23}^{t-1} = \sigma_{17}^{t-1} + \sigma_{23}^{t-1} = 1, \beta_{17,23}^{t-1} = \sigma_{12}^{t-1} + \sigma_{16}^{t-1} + \sigma_{18}^{t-1} + \sigma_{22}^{t-1} + \sigma_{24}^{t-1} = 3 \rightarrow a_{17,23} = 1. \end{cases} \quad (20)$$

Figure 21. Four snapshots of an evolution of (an initially regular 35-by-35) Euclidean lattice, the values of whose nodes are randomly seeded (with $\text{prob}(\sigma=0) = \text{prob}(\sigma=1) = 1/2$.)




Figure 20 thus shows a black circle centered on node 22 at $time = 0$, and a blank node at $time = 1$. Similarly, the figure shows that the link between pair (18,23) that exists at $time = 0$ is absent at $time = 1$, and a new link is established between initially unconnected nodes 17 and 23. The updated graph, at $time = 1$, can be drawn when we apply the

above rules to each of the 25 nodes, each of the 20 links (at *time* = 0), and each of the 56 next-nearest neighbors (at *time* = 0). One can continue in this fashion, applying the same rule to each new graph at time t , to evolve the full set of emergent structures.

A typical evolution starting from an initial state in which all sites are randomly assigned $s = 1$ with probability $p = 1/2$ is shown in figure 21. Notice the rapid emergence of complex local connectivity patterns, the appearance of which points to a geometrical self-organization. In general, structural behaviors emerging from random value states under typical rules can be grouped into four broad categories: (1) *decay* (in which initial states decay into structurally much simpler final states in which most links have been destroyed so that the resulting graph consists mostly of a large number of small local subgraphs); (2) *periodicity* (in which periodic, and globally connected, geometries emerge); (3) *complexity* (exemplified by systems that continue to change over the course of a long evolution, and punctuated by complicated local and global topologies); and (4) *dynamic equilibrium* (which is a remarkable set of rules whose action appears to lead to a prolonged series of topological changes, but during which the global effective dimensionality (defined as the average ratio of the number of *next-nearest* to *nearest* neighbors in the graph at time t [128]) stays approximately constant.

Complex networks: *metrics*

“Measure what is measurable, and make measurable what is not so.”

—Galileo Galilei

Complex networks—whether natural, technological, or social—all reflect the rules and conditions under which they evolve; though what those rules and conditions are may not be obvious by direct visual inspection alone.

A useful way to think of network evolutions—mentioned earlier in our discussion of RGs—is to think of networks as living organisms that grow and develop by acquiring nodes and links. Their growth may be constrained by physical space (so that the number of nodes cannot exceed a certain number), by available energy (which may limit the number of neighbors a given node may establish a link with), or by certain rules that preferentially induce (or inhibit) the creation of links between nodes possessing certain features.

Much in the same way as complex adaptive systems theory is concerned mostly with understanding the relationship between the emergent global properties of a given system and the local rules governing its constituent parts, an important goal of complex network theory is to understand the relationship between a network’s emergent *form* (and function) and the set of rules and constraints by which it *evolves*. The first step toward achieving this goal is to develop a set of local and global metrics by which otherwise large, often apparently featureless (at least via only a “visual” inspection), classes of networks may be distinguished.

Overview

This section provides a brief survey of the structural properties of graphs, with an emphasis on those metrics that have proven to be particularly useful for solving social network problems.

Since social network analysis is typically interested in identifying the “most important” individual in a social network, or the “most important” groups and communication links, many of the metrics discussed below may be used to establish implicit hierarchies (at least in terms of context-specific measures of “importance”) by *ranking* various components of a graph. We will later exploit this implicit ranking in defining SOTCAC’s local decision rules.

Tables 2, 3 and 4 (adapted from [133]) list some common metrics for *links*, *local topology*, and *global topology*, respectively.

Table 2. Some typical social network metrics that measure properties of links between nodes

Metric	Description
<i>Capacity</i>	Measures the load capacity of a link to carry information (or general “resource”) of a given type
<i>Duration</i>	Measures the duration of a link; is a link permanent, transient, decay of strength in time, general stability over time?
<i>Frequency</i>	Measures how often a link is active
<i>Multiplexity</i>	Extent to which a link between two nodes represents multiple kinds of relationships
<i>Strength</i>	Measures the intensity of the relationship between nodes; communicative, emotional, degree of sharing, reciprocity
<i>Symmetry</i>	Measures the extent to which a link (and/or the information that it is a conduit for) is bidirectional
<i>Type</i>	Direct/indirect, directed/undirected, weighted/unweighted
<i>Visibility/Vulnerability</i>	Measures the degree to which a link is vulnerable to eavesdropping, jamming, physical disruption or destruction

Characteristic Path Length

The *characteristic path length* of a graph G , $L(G)$, measures the typical separation between the nodes in G [41]:

$$L(G) = \frac{1}{N(N-1)} \sum_{i,j \in V(G)} \text{Dist}(i, j), \quad (21)$$

where $\text{Dist}(i, j)$ is the length of the shortest path between i and j . Because $\text{Dist}(i, j) \equiv \infty$ if i and j lie on disconnected subgraphs of G , $L(G)$

is defined only for connected graphs. (This deficiency is remedied by using another measure called *global efficiency*, which is defined later in this section; see page 94.)

Table 3. Some typical network metrics that measure properties of individual nodes; adapted from [133]

Metric	Description
<i>Betweenness</i>	Measures the extent to which a node mediates, or plays the role of “information broker” between, any two other nodes
<i>Brokerage</i>	Measures a node’s “brokerage” strength; i.e., the degree to which a node manages the information flow between two or more groups that otherwise would not be linked
<i>Centrality</i>	Measures the degree to which a node plays an important (or “central”) role in a network
<i>Closeness</i>	Measures the extent to which a given node is “close to” other nodes in the network; typically defined by averaging over all possible paths to other nodes
<i>Degree</i>	Number of links to other nodes
<i>Diversity</i>	Number of links to different node (where “different” means either that they are not linked to one another and/or otherwise represent agents that have different internal states)
<i>Eccentricity</i>	Measures maximal distance between a given node and any other node in the graph
<i>Effective Network Size</i>	Measure introduced by [134] in his analysis of “structural holes” in a network; based on supposition that links among a node’s neighbors attenuate the effective size of that node’s local network
<i>In-degree</i>	Number of directional links that point <i>toward</i> a given node
<i>Isolation</i>	Measures the degree to which a node is isolated, relative to others in the group to which it belongs
<i>Out-degree</i>	Number of directional links that point <i>away from</i> a given node
<i>Prestige</i>	Measures how strongly a given node is on the receiving end of information flow; it is defined only for directions graphs

In the limiting cases of regular and Erdos-Renyi RGs, it can be shown analytically that, for large N (= the number of nodes in G), $L(G=\text{Regular}) \sim N/2k$ and $L(G=\text{Random}) \sim \ln(N)/\ln(k-1)$, respectively, where k is the average degree [88]. Newman, *et al.* [109], have shown that, for the Watts-Strogatz model, $L_{WS}(N,p) \sim (N/K) f(pKN^d)$, where $f(x)$ is a universal scaling function of the form, $f(x) = \text{const}$ if x much less than 1, and $f(x) \sim \ln(x)/x$ if $x \gg 1$. Albert and Barabasi [88] report that, for

scale-free RG models, the path length scales approximately logarithmically with N : $L \sim a \ln(n - b) + c$ (where a , b , and c are constants).

Table 4. Some typical social network metrics used to describe entire graphs; adapted from [133]

Metric	Description
<i>Clustering Coefficient</i>	Measures the average “cliquishness” of a node within the graph (or subgraph); estimates the degree to which a graph is modular (i.e., is organized in a hierarchical fashion)
<i>Component</i>	The largest connected subset of nodes and links; all nodes within the component graph are connected, either directly or indirectly, and none of the nodes have any connections to parts of the graph outside the component
<i>Connectivity</i>	Measures the extent to which agents are linked to one another by either direct or indirect routes; typically defined using the maximum or average path distance
<i>Density</i>	Ratio of the actual number of links in a network to the total possible number
<i>Inclusiveness</i>	Measures the total number of nodes in a network minus (and sometimes the ratio to) the number of isolated (or minimally connected) nodes
<i>Path Length</i>	Measures the typical separation between nodes
<i>Size</i>	Number of agents and links in a graph
<i>Spanning Tree</i>	A subgraph of a network that is a tree that contains all nodes; of particular interest (for weighted graphs) is the <i>minimum spanning tree</i> (= the spanning tree that minimizes the weights along all links in the network)
<i>Symmetry</i>	Ratio of the number of symmetric to asymmetric links (or the total number of links) in a network
<i>Transitivity</i>	Measures the ratio of the number of transitive triples over the total number of possible transitive triples (= number of paths of length two); x, y, z are <i>transitive</i> if whenever x is linked to y and y is linked to z , z is also linked to x

Clustering Coefficient


Many real complex networks exhibit a high degree of local clustering. For example, social networks are typically *transitive*: if nodes i and j are linked, and nodes j and k are linked, there is a strong likelihood that i and k are also linked. A simple way to measure the degree of transitivity of a graph G , is to calculate the ratio of transitive triplets (i.e. triangles) over the number of connected 3-tuples of nodes.

Watts and Strogatz have introduced a more general measure they call the *clustering coefficient* of G , $Clus(G)$. In order to define it, we first define the local clustering coefficient $Clus_i$, which measures the average “cliquishness” of node i :

$$Clus_i = \frac{\text{Number of edges in } G_i}{\text{Max possible number of edges in } G_i} = \frac{\text{Number of edges in } G_i}{k_i(k_i-1)/2}, \quad (22)$$

where G_i is the subgraph of G consisting of all neighbors of i , and $k_i = \deg(i)$ is the number of i 's neighbors. The maximum number of links that can exist among nodes making up i 's immediate neighborhood is achieved in the case when G_i is completely connected, and is equal to $k_i(k_i-1)/2$. $Clus_i$ is the fraction of those possible links that actually exist. A large value of $Clus_i$ suggests that i is located in a highly connected region of G (more so in the case where G is sparse).

Figure 22. An example of how to calculate the local clustering coefficient $Clus_s$ of the vertex σ . The solid black circles denote the nodes that belong to G_σ ; thick lines indicate links among nodes in G_σ ; since there are seven edges and the degree $k_s=8$, equation 17 yields $Clus_s = 7/28$



G 's *global* clustering coefficient is the average of the local clustering over all nodes:

$$Clus(G) = \frac{1}{N} \sum_{i \in V(G)} Clus_i. \quad (23)$$

In an Erdos-Renyi random graph, the links are distributed randomly; therefore $Clus(G) = p$. Watts and Strogatz [83] were the first to point

out the discrepancy between this behavior and the typically much larger values of the clustering coefficient that are observed in real networks (of the same size and order). Using the measures $L(G)$ and $Clus(G)$, *small-world* graphs are somewhere “in between” regular and random graphs: they are highly clustered (like regular lattices), but have very small characteristic path lengths.

Degree centrality

The degree of a node $i \in V(G)$ is an individual node’s simplest, and most obvious property; it also represents the simplest measure of a graph’s central—or most “important”—nodes, by identifying the most active nodes, in the sense that they have the largest number of links to other nodes in the graph.

Recalling that $\Gamma(i)$ is the set of nodes adjacent to i , the *degree* of i , $deg(i)$, is defined to be the number of nodes adjacent to (i.e. the *neighbors* of) i : $deg(i) = |\Gamma(i)|$. For *undirected* graphs, the sum of the degrees is exactly twice the number of edges: $\sum_{v \in V(G)} deg(v) = 2M$. We will denote the minimum degree of G by $d(G)$, and the maximum degree by $D(G)$. In general, of course, $\delta(G) \leq \bar{d} \leq \Delta(G)$, where $\bar{d} = \frac{1}{|V(G)|} \sum_{v \in V(G)} deg(v)$ is G ’s average degree.

For *directed* graphs, we can define the analogous metrics *in-degree* ($=deg_{in}(i)$) and *out-degree* ($=deg_{out}(i)$), which provide measures of the number of directional links to a node from other nodes and the number of directional links from a node to other nodes, respectively.

To ensure commensurability of measures for graphs of different order, we normalize them with the maximal possible value, $d_{max}=N-1$:

$$d(i) = \frac{deg(i)}{d_{max}} = \frac{1}{N-1} \sum_{j \in V(G)} a_{ij}, \quad d_{in}(i) = \frac{deg_{in}(i)}{N-1}, \quad d_{out}(i) = \frac{deg_{out}(i)}{N-1}, \quad (24)$$

where a_{ij} is G ’s adjacency matrix.


The degree of a given node provides some (albeit limited) information. In social networks, for example, in which communication and information flow plays a prominent role, the nodes having the high-

est degree often serve (and are characterized) as the “centers” of the network. As such, “degree” as a metric can be said to be a local *centrality* metric.

Degree can also be interpreted as the probability that a node will succumb to the influence of (or adopt an innovation offered by) other nodes. Without factoring in other features of a network, degree centrality essentially measures the relative risk (or reward) of accepting whatever is flowing throughout the network. Other kinds of “centrality” metrics are discussed below.

For example, calculating the degree centrality for each of the three graphs shown in figure 23, we find that $d_{Star} = \{1, 1/6, \dots, 1/6\}$, $d_{Circle} = \{2, \dots, 2\}$ and $d_{Line} = \{1/6, 1/3, \dots, 1/3, 1/6\}$.

Figure 23. Sample *star*, *circle* and *line* graphs used for comparing centrality metrics; see text



A property of the graph, as a whole, is the probability density function of the degree, or the *degree distribution*. We have already alluded to its utility in our discussion of the power-law decay of the degree-distribution in Barabasi-Albert’s scale-free random graphs. Recall that, for scale-free networks, the degree distribution, $p(k)$ scales as $p(k) = a k^{-\gamma}$, where k is the degree, and a and γ are constants. Of course, a network’s degree distribution is not confined to being scale-free, as it can obey a variety of forms. For example, degree distributions can be scale-free but with a *cutoff* (do to the finite size and order or real-world networks), they can obey Gaussian statistics [88], or they can decay exponentially [108].

Group degree

A metric that represents an intermediate step between individual degree and degree distribution is *group degree*, deg_{group} introduced by Freeman [135]. The idea is to measure the variability of individual degrees across all the nodes of a graph:

$$deg_{group} = \frac{\sum_{v \in V(G)} [\Delta(G) - deg(v)]}{Maximum\left\{\sum_{v \in V(G)} [\Delta(G) - deg(v)]\right\}} = \frac{\sum_{v \in V(G)} [\Delta(G) - deg(v)]}{(N-1)(N-2)}. \quad (25)$$

The group degree obviously measures the dispersion of degrees in a graph, by comparing each degree with the maximum possible value. Note that the graph for which $(d_{group})_{max} = 1$ is the one in which one node has $deg = N-1$ (i.e., it is linked with all other nodes), and all other nodes have $deg = 1$; i.e., a *star graph*. The group degree attains its minimal possible value, $(d_{group})_{min} = 0$, when $deg(i) = D(G)$ for all i ; i.e. a *regular graph*. Graphs for which $0 \leq d_{group} \leq 1$ harbor intermediate degree of centralization of degree.

Link Degree

Holme, *et al.* [136,137], in the context of surveying attack strategies on a network's nodes and links (see section **Vulnerability**; page 109), generalize the definition of degree, as it naturally and intuitively applies to nodes, to the *links* between two nodes. They suggest four ways to assign a measure of “importance” to a link as a function of the degrees of the nodes it connects.³⁴

$$deg(l_{ij}) = \begin{cases} deg(i) \cdot deg(j), \\ deg(i) + deg(j), \\ Minimum[deg(i), deg(j)], \\ Maximum[deg(i), deg(j)], \end{cases} \quad (26)$$

where $deg(l_{ij})$ is the degree of link l_{ij} connecting nodes i and j .

34. The “link degrees” are used in SOTCAC as primitive measures by which notional agents assign value to emerging social-network links (see **SOTCAC: a conceptual model**; page 121).

Link Density

The *link density* of a graph, $0 \leq \rho_G \leq 1$, provides a coarse measure of the general connectivity of G . It is defined as the ratio of the number of links in G to the maximum possible number of links (which occurs for a complete graph) [42]:

$$\rho_G \equiv \frac{\# \text{ links in } G}{\# \text{ links in a complete graph of } N \text{ nodes}} = \frac{\# \text{ links in } G}{N \cdot (N-1)/2}. \quad (27)$$

One can also define a local, *egocentric*, link density that measures the connectivity around a given node. It is defined by the following expression:

$$\rho_G(\sigma_i) \equiv \frac{\# \text{ links in } (\Gamma(i) - \sigma_i)}{\deg(\sigma_i) \cdot [\deg(\sigma_i) - 1]/2}, \quad (28)$$

where $\Gamma(i)$ is the set of nodes that are directly adjacent to s_i , and $\deg(s_i)$ is s_i 's degree. Thus, $\rho_G(s_i)$ measures only the local density, as “seen” from s_i 's point-of-view; it ignores s_i itself, and focuses only on the links that exist among s_i 's contacts. The average egocentric link density, for the entire graph, $\langle \rho_G \rangle_{Ego}$, is then given by:

$$\langle \rho_G \rangle_{Ego} \equiv N^{-1} \cdot \sum_{\sigma_i \in V(G)} \rho_G(\sigma_i). \quad (29)$$

Figure 24. Sample calculations of global and local link densities (for two selected nodes, s_1 and s_2), computed for a graph consisting of 10 nodes and 20 links; see text for details




Figure 24 shows sample calculations of global and local link densities (for two selected nodes, s_1 and s_2), computed for a graph consisting of 10 nodes and 20 links. The dotted lines in the middle and right-hand-side graphs denote, respectively the links between s_1 and s_2

and their nearest neighbors; and which are therefore excluded from the local density calculations. The light-grey nodes and links appearing in the graph on the right-hand-side of the figure represent components of G that are at least a distance two away from s_2 ; and are thus also excluded from the calculation of s_2 's local density.

Eigenvector centrality

Eigenvector centrality, $C_E(i)$, introduced by Bonacich [138], is essentially a generalization of degree centrality. In the context of social networks, it is based on the supposition that one's "importance" is not solely a function of the number of people one knows, but a function of how many people that one knows are themselves also important; i.e. $C_E(i) \propto \sum_{j \in \Gamma(i)} C_E(j)$ (where $\Gamma(i)$ is the set of i 's neighbors). Bonacich formalized this notion by defining centrality as the i^{th} component of principal eigenvector of G 's adjacency matrix, $\mathbf{A} = [a_{ij}]$. Recall that an eigenvector of a symmetric square matrix is any vector $\mathbf{e}$ which satisfies the equation:

$$e_i = \lambda^{-1} \sum_j a_{ij} e_j, \quad (30)$$

where λ is a constant (i.e., the largest eigenvalue), and e_i gives the centrality of node i .

One can show that an eigenvector is proportional to the row sums of a matrix, $\mathbf{M}$, equal to the sum of powers of $\mathbf{A}$, weighted by corresponding powers of the inverse of the eigenvalue [139]:

$$\mathbf{M} = \mathbf{A} + \lambda^{-1} \mathbf{A}^2 + \lambda^{-2} \mathbf{A}^3 + \dots = \sum_k \lambda^{1-k} \mathbf{A}^k. \quad (31)$$

Since the entries of the n^{th} power of $\mathbf{A}$ give the number of different walks of length n from node i to node j [140], eigenvalue centrality can thus be interpreted as effectively counting the number of walks of all lengths starting from a given node, weighted inversely by length. As such, it assumes that information flows by *all available means*, and is not confined to taking into account only the fastest route (i.e., flows

along geodesics); a fact which alleviates a conceptual drawback found in some other popular measures of centrality.

Eigenvalue centrality is also frequently used to design World Wide Web search engines (such as *Google*), for which an intelligent estimate of “relevance” of a hyperlinked page is of obvious importance [141].³⁵ The search variant assumes that a given page is “important” if other important pages link to it. Kleinberg [119] generalizes this ranking to include both the set of nodes that link to some node of interest and the set of nodes the node of interest itself links to. Nodes that link to highly ranked nodes are *hubs*: though they may not contain the desired information, they are likely to provide links to nodes that do.

Information Centrality

A metric related to eigenvalue centrality is *information centrality*, C_I , introduced by Stephensen and Zelen [142]. The idea is to combine all possible paths from one node to another, and to assign weights to the paths so that the “information” in the combined path is maximized. Geodesics are typically assigned weight one, while longer paths receive smaller weights that are functions of the information they carry (assumed to vary inversely with their length).

To define C_I , first consider the matrix $\mathbf{B}$ with the following elements:

$$B_{ij} = \begin{cases} 1 + \deg(i) & \text{if } i = j, \\ 1 - a_{ij} & \text{if } i \neq j. \end{cases} \quad (32)$$

Then, assuming $\mathbf{B}$ has an inverse, $\mathbf{B}^{-1}$,³⁶ the information centrality of node i , $C_I(i)$, is defined by [143]:

35. The more general “navigability” problem in complex networks is discussed on page 58.

36. In practice, only connected networks considered; i.e. isolated nodes and/or disconnected subgraphs are simply ignored (their information centralities would be equal to zero in any case, so there is no loss in generality).

$$C_I(i) = \frac{1}{\sum_{j \in V(G)} C_I(j)} \cdot \left\{ \frac{N}{N\mathbf{B}_{ii}^{-1} + \text{Trace}(\mathbf{B}^{-1}) - 2\sum_{j \in V(G)} \mathbf{B}_{ij}^{-1}} \right\}, \quad (33)$$

where $\text{Trace}(x)$ is the trace of matrix x (i.e., $\text{Trace}(\mathbf{x}) = \sum_i \mathbf{x}_{ii}$), and the first term, $1/\sum_{j \in V(G)} C_I(j)$, is a normalization factor that ensures that $0 \leq C_I(i) \leq 1$.

Loosely speaking, the metric assumes that if information travels among various routes between nodes, it is inversely proportional to the variance of the signal strength. The information centrality is the harmonic average of the of the signal strengths along all the possible paths (or “bandwidth”). $C_I(i)$ thus measures the amount of information contained in all paths that start and end on the i^{th} node; it may be interpreted as the fraction of the total information flow in the network that is controlled by the i^{th} node [142]. Because nodes with larger values of $C_I(i)$ tend to have a large number of short paths to many other nodes, they typically have a greater control over the flow of information within the network than nodes with smaller values.

For example, calculating the information centrality for each of the three graphs shown in figure 23, we find that $d_{Star} = \{0.234, 0.128, \dots, 0.128\}$, $d_{Circle} = \{0.143, \dots, 0.143\}$ and $d_{Line} = \{0.104, 0.168, \dots, 0.168, 0.104\}$.

Closeness centrality

The second of several related “centrality” measures (the first was the degree of an individual node; see above) is *closeness centrality*, introduced by Beauchamp [144] and Sabidussi [145]. Providing a sense of a node’s global importance, it is based on the idea is that a given node is central if it is best able to interact quickly with all other nodes in a graph; and not just with its immediate neighbors. Closeness centrality may thus also be interpreted as an index of expected *time of arrival*—at a given node—of whatever is flowing throughout the network.

Generalizing the notion of closeness as defined by immediate proximity between nodes i and j (i.e., in which i and j are said to be “close” if and only if distance $\text{Dist}(i, j) = 1$) to one that respects *all* distances

between a given node and all other nodes of a graph, the (normalized) “closeness centrality” of node i , $C_c(i)$, is defined by:

$$C_c(i) = \frac{N-1}{\sum_{j \in V(G)} \text{Dist}(i, j)}. \quad (34)$$

The metric simply measures the inverse of the sum of the distances between node i and all other nodes in G . Its maximum value is *one*, and is attained when i is adjacent to all other nodes. Its minimum value is equal to *zero*, and occurs whenever there is at least one node, j , that is disconnected from i (in which case, $\text{Dist}(i, j) = \infty$).

For example, calculating the closeness centrality for each of the three graphs shown in figure 23 (on page 83), we find that $d_{Star} = \{1, 0.545, \dots, 0.545\}$, $d_{Circle} = \{0.5, \dots, 0.5\}$ and $d_{Line} = \{0.286, 0.5, \dots, 0.5, 0.286\}$.

There are two subsets of nodes of a graph G , related to closeness centrality, that are useful in describing G : (1) the *centroid* of G (which is typically used for *trees*), and (2) the *center* of G (which applies to any graph).

To define the centroid, first count—for each node i (with $\text{deg}(i) > 1$)—the number of nodes that appear in each of the subtrees emanating from i . Let N_i be the maximum of these numbers. It can be shown that if the tree has N nodes then *either* there is exactly one node, j , for which $N_j = 1/2(N-1)$ —which is called the tree’s *centroid*—or there are two adjacent nodes i and j for which $N_i = N_j = N/2$, which is called the *bicentroid*. (It is easy to see that every tree either has a centroid or bicentroid, but not both.) In either case, the centroid of a tree is essentially equivalent to the tree’s center-of-gravity.

The “center” of G consists of the subset of nodes in G that share the same minimal value of the maximum distance to all other nodes.

Group closeness

In the same way as the degree of a single node may be used to define the group degree for the entire graph, deg_{group} (see equation 20), we can define the *group closeness centrality*, $C_{C,group}$, by summing over, and normalizing, the differences between the maximum possible value of

closeness centrality over all nodes ($= C_{C,max}$), and the value of C_C for each node:

$$C_{C,group} = \frac{\sum_{v \in V(G)} [C_{C,max} - C_C(v)]}{\text{Maximum} \left\{ \sum_{v \in V(G)} [C_{C,max} - C_C(v)] \right\}} = \frac{\sum_{v \in V(G)} [C_{C,max} - C_C(v)]}{(N-1)(N-2)/(2N-3)}. \quad (35)$$

$C_{C,max}$ attains its maximum value of *one* when there is at least one node from which the minimum distance to all other nodes is equal to 1, and the minimum distance from the other nodes to all $(N-2)$ remaining nodes is equal to 2. As is true for deg_{group} , this occurs for the *star graph*. $C_{C,max}$ attains its minimum value of *zero* when the minimum distances between any two nodes in the graph are all equal; for example, in a complete graph or in a circle graph.

Betweenness centrality

The *betweenness centrality* of node i (or simply, *betweenness*), $C_B(i)$, is a metric that measures how “important” a role a given node (or link) plays in the information flow throughout the entire graph. Intuitively, the metric is derived from the supposition that interactions between nonadjacent nodes, i and j , are a function not just of the local information content of i and j , but on the information residing at, and communicated by, other nodes that are located on some path between i and j . A node is deemed “central” if it lies “between” many other nodes.

Freeman’s [135] original definition of $C_B(i)$ is as follows:

$$C_B(i) = \frac{1}{(N-1)(N-2)} \sum_{\substack{u,v \in V(G), \mathfrak{U}(u,v) \neq 0 \\ u \neq v, u \neq i, v \neq i}} \frac{\mathfrak{U}(u,v;i)}{\mathfrak{U}(u,v)}, \quad (36)$$

where $\mathfrak{U}(u,v)$ is the number of geodesics (i.e., shortest paths) between nodes u and v , and $\mathfrak{U}(u,v;i)$ is the number of geodesics from u and v that pass through i . The sum is taken over all pairs of nodes such that $\mathfrak{U}(u,v) \neq 0$. An analogous measure of *link*-betweenness, $C_B(l)$, can be defined using the same equation, but with $\mathfrak{U}(u,v;l)$ interpreted as the number of geodesics between i and j that contain link l .

For example, calculating the betweenness centrality for each of the three graphs shown in figure 23, we find that $d_{Star} = \{1, 0, \dots, 0\}$, $d_{Circle} = \{0.2, \dots, 0.2\}$ and $d_{Line} = \{0, 0.6, .533, \dots, .533, 0.6, 0\}$.

Computation times

An important issue associated with many of the most useful network metrics, particularly for large networks, is computation time. For example, since the definition of betweenness requires us to compute the length of the geodesic between all pairs of nodes, calculating its value of is, in principle, at least as difficult as solving the *all-pairs shortest-paths* (APSP) problem. The fastest known algorithms for solving APSP on sparse graphs entail solution times that scale as $\sim NM$ on unweighted graphs (where N is the number of nodes in a graph, and M is the number of links), and scale as $\sim NM + N^2 \log(N)$ on weighted graphs (which is the type of graph that SOTCAC depends on) [56]. For dense graphs, the computation is even more costly and scales, in the worst case, as $\sim N^3$; although some progress has been made recently to lower this bound [135]. In practice, both closeness and betweenness centrality measures can be calculated fairly quickly for relatively small graphs (containing up to a 100 or so nodes).³⁷

While *betweenness* provides valuable insight into the information flow between nodes of a graph, its main drawback is that it assumes that communication *always flows along the shortest paths*. This is a reasonable approximation, but it is certainly not always true. Communication paths often take indirect, circuitous, and/or sometimes intentionally long routes. Internet traffic, for example, often becomes jammed and data packets are forced to make indirect excursions to reach their target.

A number of refinements and generalizations to the basic betweenness metric defined above have been introduced in recent years: (1) *flow betweenness* [146,147]; *random-walk betweenness* [143] (which estimates centrality by the fraction of times a random walker visits a given node, averaged over time and starting location); (3) *e-betweenness*

37. Appendix 3 lists several freely available and commercial graph analysis packages that include algorithms to compute network metrics.

[148]; and (4) *efficiency* [149,150] (see discussion of *information centrality* in next section). Because of their importance, we conclude this section with brief descriptions of *flow betweenness* and *e-betweenness*.

Flow betweenness

Flow betweenness was introduced by Freeman, Borgatti and White [147], and is derived from the concept of “maximum flow.” First treat each link of the graph as a physical “pipe” through which a unit amount (or some over value equal to a link’s assigned weight) of “liquid” is able to flow. Now, consider some arbitrary node i as the source of flow, and another node j as the target. The *Max-Flow Min-Cut Theorem* [56] allows us to calculate the maximum possible flow between i and j , m_{ij} . The flow betweenness centrality, $C_F(k)$, of node k is then defined using the fraction of the total flow m_{ij} passing through node k , $m_{ij}(k)$:

$$C_F(k) = \frac{\sum_{i,j \in V(G)} m_{ij}(k)}{\sum_{i,j \in V(G)} m_{ij}}. \quad (37)$$

e-betweenness


Motivated by an additional limitation of the basic betweenness centrality metric, as defined in equation 29, in dealing with imprecisely specified network topology, Carpenter, et al. [151] have recently introduced a refined metric called *e-betweenness*.

Carpenter, et al. note that, in trying to understand many real networks—in particular, dynamic *terrorist* networks—one can expect there to be many uncertainties and/or inaccuracies in the available intelligence data that is used for determining sets of shortest paths. Relatively small differences in a network’s structure—whether real, or apparent, as the structure may be inaccurately reflected in incomplete and/or erroneous intelligence reports—may result in large changes to the set of geodesics within the network; and therefore result in (potentially very) different estimates of centrality.

For example, the sample network shown in figure 25 (abstracted from [151]) illustrates how betweenness, as defined in equation 31, is vulnerable to even minor changes to a graph's weight values.

With the given numerical weights assigned to each of the links, using equation 31 to compute betweenness yields the estimate $C_B(A)=0$, despite the fact that the node labeled “A” in the figure is the intuitively “central” node. Carpenter, et al. [151] point out that if the links that are assigned weight $w = 1.1$ in the figure, are all instead assigned the weight $w = 1.3$, **A** *does* become the central node, and its value of betweenness becomes $C_B(A)=4$.

Figure 25. Sample *weighted* graph for which $C_B(A) = 0$; adapted from [151]



ϵ -betweenness is defined by first introducing the notion of a *simple path* between nodes i and j , $P_S(i,j)$: $P_S(i,j)$ is a regular path between i and j that never visits any node more than once. An *e-shortest path* between i and j , $P_e(i,j)$, is a simple path for which:

$$Length[P_S(i, j)] \leq (1 + \epsilon) \cdot Dist(i, j), \quad (38)$$

where $Length[P]$ is the length of path P . The ϵ -betweenness is then defined by:

$$C_B^\epsilon(i) = \frac{1}{(N-1)(N-2)} \sum_{\substack{u,v \in V(G), \mathcal{D}(u,v) \neq 0 \\ u \neq v, u \neq i, v \neq i}} \frac{\mathcal{N}_\epsilon(u,v;i)}{\mathcal{N}_\epsilon(u,v)}, \quad (39)$$

where $\mathfrak{N}_\varepsilon(u, v)$ is the number of ε -shortest paths between nodes u and v , and $\mathfrak{N}_\varepsilon(u, v; i)$ is the number of ε -shortest paths that pass through i . Using this expression to calculate ε -betweenness for the sample graph shown in figure 25, we find that, for $\varepsilon = 0.1$, $C_\varepsilon^b(\mathbf{A}) = 2$.

Efficiency centrality

Latora and Marchiori [150,152] have recently introduced a metric called *efficiency centrality*³⁸ that is designed to rank nodes according to the relative role they play—both locally and globally—in a network’s overall information flow.

The metric is based on two basic ideas: (1) that information propagates in parallel (i.e., that all nodes communicate concurrently), and (2) that the relative importance of a given node (or group of nodes) depends on how the network performs when that node (or group) is removed from the network. Latora and Marchiori show how information centrality can be used to identify the critical components of a network (and thereby aid in vulnerability analyses), and then apply it to the protection of communication flow on the internet and disrupting the dynamics of terrorist networks.

Before defining efficiency centrality, we must first introduce the node-to-node communication *efficiency*, e_{ij} , which is assumed to be inversely proportional to the shortest distance between nodes i and j , $e_{ij} = 1/\text{Dist}(i, j)$. Latora and Marchiori [152] define the global ($=E_{\text{global}}$) and local ($=E_{\text{local}}$) efficiency of a graph G as follows:

$$E_{\text{global}}(G) = \frac{1}{N(N-1)} \sum_{i, j \in V(G)} \varepsilon_{ij}, \quad E_{\text{local}}(G_i) = \frac{1}{N} \sum_{i \in V(G)} E(G_i), \quad (40)$$

where G_i is the subgraph of G consisting of neighbors of i . That is, G_i consists of $k_i = \text{deg}(i)$ nodes, and contains, at most, $k_i(k_i-1)/2$ links; thus:

38. Latora and Marchiori [152] actually call their measure *information centrality*. Because we have already defined another centrality measure by that name (introduced by Stephensen and Zelen [142] in 1989), and which differs from the measure being introduced here, we will hereafter refer to Latora’s and Marchiori’s metric as *efficiency centrality*; see also [153,154].

$$E(G_i) = \frac{1}{k_i(k_i - 1)} \sum_{u,v \in V(G)} \varepsilon_{ij}(G_i), \quad (41)$$

where e_{ij} appears as $e_{ij}(G_i)$ as a reminder that each term in the summation must be computed as though the entire graph consisted, temporarily, only of the subgraph G_i .

Node information centrality

Using the definition of efficiency that appears in equation 35, we can now define the *information centrality* of node i , $C_I(i)$ [152]:

$$C_E(i) = \frac{\Delta E_{global}(G)}{E_{global}(G)} = \frac{E_{global}(G) - E_{global}(G - G_i)}{E_{global}(G)}, \quad (42)$$

where $G - G_i$ is the graph that remains after node i and all k_i links attached to i are removed from G . Since the deletion of links can only result in communication paths between points that are at least as long, and possibly longer, as the paths in the original graph, the efficiency of the smaller graph, $E_{global}(G - G_i) \leq E_{global}(G)$, and $0 \leq C_E(i) \leq 1$.

Group information centrality

The definition of point information centrality is easily extended to an arbitrary group of nodes in G , $\mathcal{S} = \{n_1, \dots, n_g\} \in G$ [152]:

$$C_E(\mathcal{S}) = \frac{E_{global}(G) - E_{global}(G - G_{\mathcal{S}})}{E_{global}(G)}, \quad (43)$$

where $G - G_{\mathcal{S}}$ is the graph that remains after the set of nodes, $G_{\mathcal{S}}$, and all the links attached to each of the nodes in $G_{\mathcal{S}}$ are removed from G .

Graph efficiency centrality

The node and group information centralities can be generalized to provide a measure of how centralized the entire graph is, with respect to the node-to-node communication efficiency defined above. The *graph efficiency centrality*, $C_E(G)$, is defined analogously to how we previously defined *group degree* (equation 20) and *group closeness* (equa-

tion 30). The basic idea is to use local centrality to measure the extent to which the graph's global topology is organized around its most central point (or group of points).

$$C_E(G) = \frac{\sum_{v \in V(G)} [C_E(v_{\max}) - C_E(v)]}{\text{Maximum} \left\{ \sum_{v \in V(G)} [C_E(v_{\max}) - C_E(v)] \right\}} = \frac{\sum_{v \in V(G)} [C_{E,\max} - C_E(v)]}{(N+1)(N-2)/(N+2)}, \quad (44)$$

where $v_{\max}$ is the node with the highest information centrality, and the maximum possible value of $\sum_{v \in V(G)} [C_E(v_{\max}) - C_E(v)]$ is obtained for the N -node *star-graph*.

Comparison with other centrality measures

How does efficiency centrality compare to other measures of centrality, such as *degree*, *closeness*, and *betweenness*? For graphs that contain “obviously” central nodes, such as the central node in a star graph or the equicentral nodes of a regular graph, all four measures of centrality agree on which nodes are central. However, Latora and Marchiori [149] point out that this agreement breaks down for graphs between these extremes.

They cite the example of a graph, $G = G_1 \cup v \cup G_2$, that consists of two (possibly large) subgraphs G_1 (of order N_1) and G_2 (of order $N_2 < N_1$) that are not directly connected to each other, but a few nodes of each of which are linked to an “intermediate” vertex v . Thus, $\text{Dist}(G_1, G_2) > 1$ and $\deg(v) \ll N_1, N_2$. While the efficiency and betweenness centralities agree in their assessments of importance for this graph, in the sense that they both identify the node v as the most “central” by assigning it the highest measure, it is unlikely that v would be so identified either by degree or closeness; it is likely that either G_1 or G_2 contains nodes with a higher degree than v , and the node that is closest to all other nodes is likely to be in G_1 (with an even stronger likelihood of being in G_1 if $N_1 \gg N_2$).

Structural Holes

Burt [134] introduced several metrics that describe the extent to which an individual node (or group of nodes) fills what he calls “structural holes” in a network. Developed primarily as a way to iden-

tify and harness the latent entrepreneurial value in business organizations, these metrics effectively measure properties of a node's "ego-centered" network; i.e., the local subset of the network that—from an *individual node's* point of view—is the only part of the whole that is of immediate importance.

Since Burt's goal is to better understand the social structure of competition in business environments, his calculations focus heavily on rates of return on, and profits from, capital investment. Objectively speaking, this appears—*a priori*—to have little to say about the dynamics of terrorist networks. However, much of the *theory* behind Burt's otherwise business-centric use of structural holes, is based on considerably more abstract (and therefore more generally applicable) concepts. The basic question that *Structural Hole Theory* (SHT) addresses is [155], "*How does a node's local perception of the "ambient world" around itself, along with the local perceptions of its partners, relate to the value that node represents to the network (in terms of its ability to process, coordinate and control information)?*"

Burt partitions the "capital" in a network into three abstract categories: (1) *resource capital* (which, in a financial context, refers to lines of credit, reserves in a bank, investments due, and cash on hand; more generally, it refers to any consumable, dynamical entity, physical and/or virtual, that an agent needs to perform its actions); (2) *human capital* (which refers to the innate personality, knowledge and skills possessed by each agent); and (3) *social capital* (which refers to the set of relationships a given agent has with other agents, and through which agents may identify and exploit opportunities to harness their resource and social capital for personal and/or group gain).

The network, thought of as a competitive arena in which agents vie for higher returns on their investments, is defined by the coevolving social structure that both provides opportunities for, and constrains the possible actions that can be taken by, individual nodes (or agents). *Who should I establish a relationship with? Who do I trust enough to make a deal with? What information do I need to gain an advantage over my peers? Who do I need to connect with in order to establish an indirect link with those who have the information I need?* Agents who are best able to answer these, and other similar, questions regarding the nature and

flow of capital throughout the network, are best positioned to recognize and seize opportunities as they arise.

In a financial setting, “seizing an opportunity” is strictly entrepreneurial, and translates to mean that agents are generally motivated to harness their social networks to maximize the return on investment; in a terrorist organization, “seizing an opportunity” means being adept at marshalling resources, recruiting agents, and harnessing the latent social terrorist-network capital to achieve its ultimate goal of committing acts of terror. While the details of how agents go about recognizing and seizing opportunities are obviously different in different contexts, the social network dynamical rules according to which this evolves is essentially the same.

SHT provides metrics by which the “action potential” of any given terrorist’s ego-centered network may be objectively graded. The SOTCAC model, which is discussed in a later section, borrows many of the ideas outlined in this section.

Effective Network Size

A node’s *effective network size*, S , is based on the premise that links among a node’s neighbors attenuate the effective size of that node’s local network. The maximum effective size occurs when the nodes that a given node is linked to have no links among themselves. The value of S is reduced by the average number of links that neighboring nodes have among each other.


To define S , we first introduce the notion of *redundancy*, which measures the extent to which a node’s local ties (i.e. its “ego network”) are mutually interconnected. The two simplest ways in which this can happen are by *cohesion* and *structural equivalence* [134]:

1. *Redundancy by cohesion*, in which two nodes are redundant to the extent that they are linked by a strong relationship. Consider a node, A , that is linked to three neighboring nodes, each of which are connected to one another, and each of which thus, redundantly, provides the same network benefit to A .

2. *Redundancy by structural equivalence*, in which two nodes are redundant to the extent that they have the same contacts. Considering the same node *A* with its three ties, even in the case where none of *A*'s three neighbors have direct links to each other, each may nonetheless have ties to another (more distant) group of nodes, none of whom have direct ties back to *A*. In this case, the information *A*'s neighbors receive, and the nodes responsible for sending this information, are all redundant.

Structural holes are, essentially, the implicit boundaries separating groups of nonredundant nodes; they serve as buffers between nonredundancy, waiting to be “discovered” and exploited by the entrepreneur (or, in the context of terrorist net dynamics, by the organization’s mastermind or terrorist cell leader); figure 26 shows a simple example.

Figure 26. Illustration of two similar local topologies; one of which has no structural holes (a) and one that has three (b)



If a node’s effective network has many structural holes (and none are attached to the node itself), it is rich in “entrepreneurial” opportunity. Burt calls such a state of affairs, structural autonomy; which effectively measures the degree to which a given node (relative to other nodes) has unconstrained access to structural holes. The problem, for agents, is to have some metrics by which to recognize the structural holes around them.

Let l_{ij} be the (possibly asymmetric) matrix of values that represent the *strengths* of ties between nodes i and j . Heuristically, high values of l_{ij} indicate strong, cohesive links between the two nodes. Nodes i and j are “structurally equivalent” to the extent that their relationship with other nodes k are the same. Consider an arbitrary node i , with $\deg(i)$ neighbors, and ask, “What is the redundancy of one of i ’s neighbors, j ?”

The information that i receives via j is redundant (in the sense defined above) to the extent that i has a strong tie to another node, k , to which j also has a strong tie:

$$p_{ik}m_{jk} , \quad (45)$$

where p_{ik} is the fraction of i ’s network ties (“attention”) that are focused on its relationship with k :

$$p_{ik} = \frac{(l_{ik} + l_{ki})}{\sum_j (l_{ij} + l_{ji})}, \quad i \neq j , \quad (46)$$

and m_{jk} is the marginal strength of j ’s relationship with k (i.e. its interaction with k divided by the strength of the strongest relationship j has with other nodes):

$$m_{jk} = \frac{(l_{jk} + l_{kj})}{\text{Maximum}_{i \neq j} (l_{ij} + l_{ji})} . \quad (47)$$

The fraction of i ’s relationship with j that is redundant to i ’s relationship with other neighboring nodes is then given by the sum:

$$\sum_{k \neq i, j} p_{ik}m_{jk} . \quad (48)$$

Finally, i ’s effective network size, $1 \leq S(i) \leq \deg(i)$, is computed by summing over the nonredundant fraction of i ’s relationship with each of its neighbors (equal to one minus the redundancy):

$$S(i) = \sum_{j \in \Gamma(i)} \left\{ 1 - \sum_{k \neq i, j} p_{ik} m_{jk} \right\}. \quad (49)$$

If a node j is completely disconnected from all of i 's immediate neighbors, the j^{th} term (in brackets) is equal to one, which is interpreted to mean that j contributes a unit-valued nonredundant link in the network. As the strength of the relationship between j and other nodes increases, the bracketed term approaches p_{ij} , meaning that j is completely redundant with other nodes in i 's effective network. The value of $S(i)$ ranges from a minimum of *one*, to a maximum equal to $\text{deg}(i)$, achieved when each of i 's contacts are nonredundant.

Efficiency

The *efficiency* of a node, which is simply the effective network size normalized by the maximum possible value ($=S(i)/\text{deg}(i)$),³⁹ measures the extent to which a node has maximized the effective size of its egocentric network. Jumping ahead a bit toward our discussion of SOTCAC, one of the ways in which users can “tune” agents’ otherwise autonomous local rewiring decisions is to “weigh” their efficiency. That is, in choosing to take specific actions, terrorist agents may be more or less “motivated” to maximize the effective size of their ego-networks.

Community structure

Among the vast number of complex networks—ranging from uncorrelated Erdos-Renyi random graphs, to small-world networks, to scale-free networks (see the “zoology” illustrated in figure 12)—one property that has attracted particular attention among social network researchers, and is of obvious interest to intelligence analysts seeking to “discover” hidden cells and other structural patterns within terrorist networks, is *community structure*.

By “community structure” we mean a topology whose nodes and links are arrayed in such a way that nodes are both highly clustered within local groups and relatively loosely connected to the nodes living in other tight clusters.

39. Burt’s “efficiency” is not the *efficiency* metric, defined by equation 35.

Figure 27. An illustration of a network with “community structure”




Figure 27 shows a typical example of a small graph that contains five communities; the “communities” are indicated by the groups of nodes lying within the dashed circles. Intuitively, we expect this kind of modular structure to arise in social contexts, since people tend to divide naturally into groups based on common interests, geography, occupation, age, and so on. For the purposes of this paper, we are interested in two classes of problems: (1) how to *find* community structures in arbitrary networks, and (2) how to *characterize* them, mathematically.

Finding community structures


The problem of finding community structures in complex networks is closely related to the *graph partitioning problem* in graph theory⁴⁰ [90] and *cluster analysis* in social networks [42]. As for most “pattern

40. For example, in the context of computer science, where a node might represent a single component of a parallel processing computer, a typical graph partitioning problem consists of finding an optimal allocation of nodes that balances the processing load on nodes, and minimizes the total number of internode connections. While simple to state, the problem is believed to be NP-hard, meaning that solution times grow exponentially with the size of the problem [101].

detection” problems of this kind on graphs, one can visually discern communities only in relatively small graphs (that have, at most up to a few hundred nodes, and only by using a good graph visualization program); as the size of the network grows much beyond that, one must typically resort to using automated cluster-finding algorithms.

A powerful graph visualization tools for finding community structures in networks (along with other latent patterns) is the *force-directed* (or *spring-embedding*) algorithm described on page 34. The method consists of representing the links by linear springs (with a force constant proportional to an appropriate measure of communal “similarity”), and allow the system to relax to its equilibrium position.

Figure 28. Spring-embedding visualization of the patterns of friendships among children in a US school; taken from [157]




Example 1: social networks

Figure 28 shows a spring-embedding visualization of the social network of friendships among children in a US school [156], color coded according to the race of individuals each node represents (and divided, top/bottom according to age). Although the graph is very large and it is hard to distinguish individual links, the strong community structure is immediately obvious.

Example 2: terrorist networks

Figure 29 shows a schematic of the “community structure” evident in the global *Salafi* network.⁴¹

Figure 29. Schematic of the global *Salafi* network; after Sageman [25]



The figure shows that the network is dominated by four large clusters—labeled *Central Staff*, *Core Arab*, *Maghreb Arab*, and *Southeast Asian*—that are highly connected internally, but are only loosely associated with nodes in other clusters. Moreover, the internal architecture of each of these clusters is built around central hubs; Osama bin Laden (shown in red and labeled in the middle of the figure) is the obvious hub of the *Central Staff*;⁴² the two Arab clusters, in particular, approximate small-world networks. The *Southeast Asian* cluster consists of the Jemaah Islamiyah, which is more hierarchically organized than, say, Al Qaeda [25].

41. “Salafi” is a term that is used to describe fundamentalist islamic thought, and is here used to denote the global jihad that includes not just Al Qaeda, but the Egyptian Islamic jihad, Jemaah Islamiyah, and Abu Sayyaf group, among others; see [158]

42. Not shown in the figure are the central players in the *Core Arab*, *Maghreb Arab*, and *Southeast Asian* clusters: Sheikh Mohammed, Zein al-Abidin Mohammed Hussein, and Abu Bakar Baasyir, respectively [159].

Cluster finding algorithms

A number of algorithms for discovering communities in social networks have been introduced in recent years [42]. The pioneering work on community recognition is due to Laumann [160], who applied early forms of multidimensional scaling techniques to test for modularity in patterns of embedded social distance. “Social distance,” which was discussed earlier in the context of navigability of graphs,⁴³ measures the degree to which two nodes (or agents) perceive themselves as being similar, where “similarity” is defined according to the relevant social features that characterize a given multidimensional hierarchy (or social network). “Multidimensional scaling” consists of expressing social distance in terms of a physical metric, and is a direct precursor of the *force-directed* visualizations shown in the previous section. Laumann’s contribution was in using an objective measure of social distance (derived from the network) to discern association patterns in networks, rather than relying on subjective measures (such as self-ranked “similarity”) which were prevalent at the time.

A common modern technique to extract community structures from networks is called *hierarchical clustering* [41], which groups nodes into subsets so that all of the nodes within a given subset are maximally similar to one another. The method consists of first assigning each of the $N(N-1)/2$ possible links in an order N graph some social distance measure, S_{ij} , to represent the degree of structural equivalence between i and j , along an assignment of a threshold value S_T , such that i and j are assumed to be structurally equivalent if $S_{ij} \leq S_T$. The following four steps are then applied (Johnson [161]):⁴⁴

1. *Assign each node to its own cluster (i.e., start with N clusters). The distance between any two clusters is equal to the distance between the nodes within them.*
2. *Merge the two closest clusters into a single cluster.*
3. *Calculate the distances between the new cluster and each of the old clusters.*

43. See the discussion of Watts, et al.’s *social-distance*-based networks, page 61.

44. Many freely available social network analysis programs such as *AGNA* and *Pajek* (see **Appendix 3** for details) include a hierarchical clustering algorithm.

4. Repeat the second and third steps until all nodes are clustered into a single cluster of size N .

At the end of step four, the process yields a *hierarchical tree* rooted at the single cluster and branching off into as many levels as the number of iterations of the four-step process that were required to finally merge all the nodes into a single cluster. Each level in the hierarchy represents the clusters that were created at a given iteration step; thus a horizontal cut through the hierarchy just above step X , say, splits the network into an associated set of communities (*modulo* the social distance and inter-cluster distance functions that were used to generate the hierarchy).

The distance calculations that must be performed during the first and third steps can, of course, be done in a number of ways. For example, the inter-cluster distance, $D[C_i, C_j]$, can be set equal to either the *shortest* or *longest* distance from any node in one cluster, C_i , to any node in the other, C_j ; or $D[C_i, C_j]$ can be set to equal the *average* (or *median*) distance between the nodes in C_i and C_j ; *maximum flows* [162] and *weighed path sums* [46] between the nodes can also be used.

Girvan and Newman [163] have recently suggested combining a divisive form of the above algorithm with variants of *link betweenness*.⁴⁵ Rather than merging nodes into progressively larger clusters, Girvan's and Newman's approach is to systematically take away from a network edges that serve as links between large numbers of nodes. The network is thus broken down into smaller and smaller components, where the clusters at each stage of the process contain only those nodes that remain after the links with the *highest value* of betweenness on the previous iteration step have been removed. Among the many possible betweenness measures that can be used are

1. *Geodesic betweenness* (in which one counts the number of shortest paths between nodes in C_i and C_j), and
2. *Random-walk betweenness* (in which one also counts paths, but does not assume that the "signals" traveling along links travel only along the shortest-paths, but rather perform random walks until they reach their target).

45. See discussion on page 90.

Tyler, *et al.* [164] make several recommendations on how the computation time of Girvan’s and Newman’s algorithm can be improved. Radicchi, *et al.* [165] generalize Girvan’s and Newman’s algorithm by using loops of varying lengths to measure between-community edges. One drawback to the method (that is ameliorated by using the *modularity index*, discussed below), is that it does not provide any insight into how many communities a given network naturally contains; i.e., while it yields clusters—by construction—the cluster that emerge are based entirely on the form of the metric used to define the inter-cluster distances, and not on the clusters that the network innately contains, but are hidden.

An alternative algorithm, that takes as its starting point a slightly different form of the community-search problem—namely, “How can one identify an existing community to which an arbitrarily selected node belongs?”—is based on a *maximum flow* method, and is due to Flake, *et al.* [166] (who address it in the context of searching the World Wide Web).

Quantifying community structure

While algorithms for *finding* (or, more accurately, *recovering*) community structures in networks work well (for both real and simulated networks [56,167,168]) when some information about the communities exists prior to running an algorithm (even if that information consists only of a loosely defined social distance metric), the more typical problem is to find communities about which we know little, if anything, beforehand.

Lusseau and Newman [169] have recently introduced a *modularity index*, Q , that quantifies the strength of a given community structure;⁴⁶ it can therefore be used both to discriminate between alternative community finding algorithms and to objectively assess how well a given algorithm performs.

To define Q , suppose a network is divided into k communities. We label the community to which node i belongs, $C(i)$; and use the standard notation for adjacency matrix element: $A_{ij} = 1$ if node i is linked with j , otherwise $A_{ij} = 0$.

46. Q is based on a related index, called *assortative mixing*; see M. E. J. Newman, “Assortative Mixing in Networks,” *Phys. Rev. Lett.*, Volume 89, 2002.

The fraction of links that lie within the communities (that is, that connect nodes that within the same community), is then given by [170]:

$$\frac{\sum_{ij} A_{ij} \cdot \delta(C(i), C(j))}{\sum_{ij} A_{ij}} = \frac{1}{2m} \cdot \sum_{ij} A_{ij} \cdot \delta(C(i), C(j)), \quad (50)$$

where d is the *Kronecker delta function*, $d(i, j) = 1$ if $i = j$ and $d(i, j) = 0$ otherwise, and $m = \sum_{ij} A_{ij} / 2$ is the number of links in the network. The modularity index is obtained from this quantity by subtracting the fraction of links that lie within the communities that are expected to arise in a random network of the same size. Since the probability of having a link between i and j is $\deg(i) \cdot \deg(j) / 2m$, the modularity Q is defined by:

$$Q = \frac{1}{2m} \cdot \sum_{ij} \left[A_{ij} - \frac{\deg(i) \cdot \deg(j)}{2m} \right] \cdot \delta(C(i), C(j)), \quad (51)$$

which, with a bit of simple algebra, can be written in the simpler form:

$$Q = \sum_i \left[e_{ii} - a_i^2 \right], \quad (52)$$

where:

$$e_{ij} = \sum_{ab} A_{ab} \cdot \delta(C(a), i) \cdot \delta(C(b), j), \text{ and} \quad (53)$$

$$a_i = \sum_a \deg(a) \delta(C(a), i) / 2m. \quad (54)$$

If the fraction of links within a given network is the same as the number expected to exist in a completely random network, $Q=0$. Nonzero values of Q therefore represent deviations from randomness, and appear when community structure is present; Clauset, *et al.* [171] suggest using the value $Q = 0.3$ as a threshold above which one is assured of the network possessing significant community structure. By thus providing an objective measure of modularity, Q can be used to tune searches for communities. For example, rather than simply merging the two “closest” clusters in the search algorithm outlined on page 105, one can again start with each node as sole member of a community of one, but now iteratively link communities such that, at each step, the value of Q is maximized.

Zhou [172] has recently introduced an alternative modularity metric, called the *dissimilarity index*, $L(i,j)$, which measures the affinity strength that any two nearest-neighbor nodes, i and j , have for the same community:

$$\Lambda(i,j) = \frac{\sqrt{\sum_{k \neq i,j} [D_R(i,k) - D_R(j,k)]^2}}{N-2}, \quad (55)$$

where $D_R(i,j)$ is the average number of steps a random walk requires to get from node i to node j . If i and j , both belong to the same community, the distance $D_R(i,k)$ from i to any k (not equal to i or j) will, on average, be equal to $D_R(j,k)$, and thus $L(i,j) \sim 0$. If i and j belong to different communities, each is likely to have multiple paths within their respective communities that a random walk must sample in a search for a path to the other node, and $L(i,j)$ will thus tend to be large.


Vulnerability

This section has introduced a variety of both local and global network metrics; the objective being to identify some of the salient features of a network's *structure* and *function*. Being now armed with (at least a rudimentary) language with which such features may be described mathematically, it is natural to ask:

“In what ways can a network's structure be altered in order to degrade and/or disable its function?”

Answers to this question reveal ways both for a network to protect itself, and for network's adversary to destroy it. The question becomes particularly poignant as our attention turns to exploring ways to use SOTCAC's counterterrorist agents to disable a terrorist network's ability to function. For example, figure 30 illustrates the effects of removing two nodes (colored red) from an initially connected network; the result is one isolated node and an increase in distance between nodes A and B (from *two* steps to *five*). If the links are communication ties, this increase represents a substantial rise in expected delivery times of messages. The problem is to identify, as objectively as possible, the set of nodes whose removal does the maximum damage.

Figure 30. Illustration of the effects of node deletion on an initially connected network



We conclude our discussion in this section by briefly surveying some recent studies of the vulnerability of complex networks to some form of attack.

Any serious assessment of the vulnerability of a network must begin by positing a measure (or measures) of the network's *functionality*, m_F ; a network's *vulnerability to an attack A* is then characterized as the degree to which m_F is degraded by A. For some networks, deciding what measures are appropriate to use is straightforward. For example, if the problem is to study the vulnerability of a communication network, an obvious measure to use is the size of the *largest connected component*—as an indicator of the efficiency with which information flows within the network—that remains as nodes and/or links are removed from the system; another obvious measure is Latora's and Marchiori's *efficiency* (see page 94).

For other networks, such as terrorist networks, which perform multiple functions simultaneously, and whose "functionality" is typically a function of both context and time and is dispersed throughout its myriad levels, the set of appropriate measures to use are not at all obvious; and a combination of measures may prove more useful than any single measure.

The two main reasons why this section has introduced so many different network metrics is to (1) provide the reader with an appreciation of the enormous spectrum of (both conceptual and mathematical) measures that exists , and (2) to have a referenceable “palette” of measures to use to “tune” the behaviors of SOTCAC’s counterterrorist agents. As we have seen throughout this section, measures such as characteristic path length, average inverse geodesic length, the clustering coefficient, betweenness and/or information centrality, efficiency (along with other measures introduced in this section) are *all* useful; each characterizes an important aspect of a network’s behavior. Which measures are more, or less, appropriate to use to examine a specific issue depends partly on the scenario being considered, partly on the dynamic contexts that the agents of the simulation find themselves in, and partly on the subjective judgement of the modeler (or analyst). Indeed, a strong motivation for developing SOTCAC is with an eye toward providing the intelligence community with an analytical tool to help explore the efficacies of, and tradeoffs among, the variety of *a priori* equally viable “solutions” that exist to this difficult problem.

Network Reliability

Assessing a network’s vulnerability to attack is, conceptually, similar to determining a network’s *reliability*; which is a well-studied mathematical problem [173,174]. Thinking, again, of a graph as a communication network—in which nodes represent communication centers and links the communication channels between them, and in which both nodes and links are fallible and subject to random failures—the “reliability problem” consists of finding the probability that the network will continue to function properly in the event that one or more of its components fails; i.e., to find the network’s *reliability*.

An example of a typical problem is the *k-terminal reliability problem* [175], for which one must determine the probability that a given subset of k nodes is connected. Another common problem is the *all-node link reliability problem* [176], in which one starts with a graph in which all the nodes are perfectly reliable but whose links fail independently with some probability p , and must then determine the probability that the surviving links constitute a spanning connected subgraph of the original graph.

Unfortunately, for all such problems, the task of finding exact solutions is, in practice, generally difficult. For example, the *k-terminal* and *all-node link reliability* problems are both known to be NP-hard to solve [173].⁴⁷ Nonetheless, much work has been done in finding efficient solution algorithms for certain subclasses of graphs. For example, polynomial-time algorithms for reliability exist for acyclic graphs, complete graphs, and a few classes of threshold graphs [174,175].

The analogy between communication networks and terrorist networks is crude, at best. While meaningful estimates of the reliability of communication nets may be obtained using little more than a simple measure of information flow, the same cannot be said of terrorist networks, whose ability to “function”—and whose ability to *continue to function* when their components malfunction, fail, or are removed entirely from the network—depends on a vastly larger and more complicated set of features and distributed activities. Nonetheless, because the fundamental problem is essentially the same in both cases—namely, “*How does one characterize a net’s ability to maintain a threshold level of functioning when one or more of its components fails to function?*”—network reliability theory provides a rich source of ideas and algorithms. General references include those by Ball [174], Buzacott [177], Harms [178], Marlowe and Schoppmann [174], Shooman [179], and Bodlaender and Wolle [180].

Error tolerance versus vulnerability to attack

We earlier touched upon the problem of assessing the vulnerability of networks during our discussion of random graphs and scale-free networks (see pages 51-56). We commented that whereas *randomly* connected nets tend to degrade steadily and predictably—slowly losing their connectivity as nodes (or links) are either randomly or selectively removed, and eventually breaking apart into multiple disconnected subgraphs—*scale-free* networks are much better at tolerating

47. NP (= *Nondeterministic Polynomial*) problems are those for which the solutions, once found, are verifiable in polynomial time (using a nondeterministic Turing machine), but which may require an exhaustive search over all possible solutions (with solution times that grow exponentially with problem, or network, size) to actually solve. NP-hard problems are those that are at least as hard as any NP-problem, although they may even be harder to solve [101].

random errors than they are at tolerating attacks. In particular, while scale-free networks may not *initially* suffer any ill effects from random node removal—since it is statistically unlikely that any key (i.e., highly connected) nodes will be removed by a *random* attack—if the attack on the network is intelligently focused on its most important hubs (which are, typically, relatively few in number), the results may be catastrophic. Only a few of vital hubs need to be removed to cripple the entire system. Scale-free networks effectively trade vulnerability to attack for strong error tolerance.

For example, Albert, Jeong, and Barabasi [112] observe how the *diameter* of a graph (as a basic measure of *connectedness*) changes in response to random errors; where the errors are modeled by removing a fraction, f , of nodes from the network. Albert, et al., observe that while the diameter increases monotonically with f for random graphs—a behavior that results directly from the homogeneity of random graphs (since all nodes have roughly the same number of links, the removal of any node is likely to be as detrimental to the network’s connectivity as any other)—the diameter remains *constant* for scale-free networks, even if up to 5% of the nodes are removed (or fail). This result is attributable to the *inhomogeneity* of scale-free networks: since only relatively few nodes are highly connected, the probability that a node removed at random will significantly degrade connectivity is low.

The flip-side of error-tolerance is *vulnerability to attack*. To simulate targeted attacks, Albert, et al., systematically eliminate the most highly connected remaining node in both random and scale-free networks. As expected for random graphs, since each node is, on average, equivalent to other nodes, there is little appreciable difference between randomly removing nodes or targeting “selected” nodes. On the other hand, selective targeting has profound consequences in scale-free networks: Albert, et al. find that the diameter of scale-free networks is doubled when only 5% of the selectively targeted nodes are removed; a behavior that, as in the earlier example, stems directly from the inhomogeneity of scale-free networks.

As a final comparison between random and scale-free networks, Albert, et al., compare how the two types of networks are *fragmented*

by either random failure or targeted attack. As some fraction, f , of the nodes is removed, the degree of “fragmentation” is monitored via two measures: (1) $S(f)$ = the size of the largest connected cluster of nodes that remains (normalized by the total network size), and (2) $\langle s(f) \rangle$ = average size of all remaining clusters (except for the largest one). For *random* networks, $S(f)$ and $\langle s(f) \rangle$ exhibit a simple threshold behavior: both remain fairly constant for all $f < f_c$ at which point they abruptly decrease ($S(f_c) \sim 0$ and $\langle s(f_c) \rangle \sim 1$).

Scale-free networks, on the other hand, display no such thresholding: $S(f)$ *slowly decreases* and $\langle s(f) \rangle \sim 1$ for most f , indicating that the topology of scale-free nets remains largely unaffected, even when a sizeable fraction of nodes is removed.

Lai, Motter and Nishikawa [181]-[183] study the effects of attacking links (rather than nodes), observing that scale-free networks are generally more sensitive to attacks on short-range rather than on long-range links. Using “efficiency” (as defined by Latora and Marchiori [149]⁴⁸) as a measure of a network’s *function*, Lai, *et al.* find that short-range attacks are much more destructive, on average, than long-range ones; at least for scale-free networks whose scaling exponent, γ , lies between about 3 and 5.⁴⁹

Long-range attacks become more destructive only for small or large scaling exponents. In the case of small γ (< 2), long-range attacks are effective due to the appearance of densely connected subnetwork of nodes with large connectivity. In the case of large γ ($\rightarrow \infty$), long-range attacks become more effective due to the increase in the network’s homogeneity: as more and more nodes are likely to share the same connectivity, it becomes increasingly likely that links with the highest loads are those between distant nodes. (This is exactly the behavior seen in both the Watts-Strogatz [83] and Erdos-Renyi [79] random graph models; both of which generate homogeneous networks.)

48. Using this metric, a network is more “efficient” when it has small shortest paths; see equation 35.

49. Recall that the *scaling exponent*, γ , defines the decay of the degree distribution for scale-free networks: $P(k) \sim k^{-\gamma}$, where $P(k)$ is the probability that a node has degree k , and γ is the scaling exponent.

Attack Strategies

In the context of assessing a network's vulnerability to attack, an "attack strategy" is understood to mean a targeting procedure (or set of selection "rules") whereby the nodes and/or links to be removed from the network are selected in some mathematically well-defined manner. In other words, a strategy is an "answer" to the question, "*Which set of nodes and/or links must be removed from a network, and in what order, to maximally degrade the network's performance?*"

One obvious strategy, which is common in studies of the vulnerability of computer networks, is to target the most highly connected node; systematically removing the nodes in descending degree order, as they are removed one by one [88]. Alternatively, one can target nodes that rank most highly in *betweenness* (see equation 31) or *centrality* (see equation 23), or any of the individual metrics of local "value" or "importance."

The drawback to all of these basic strategies, however, is that they depend only on the initial degree distribution, and do not take into account the likelihood that the targeted net will adapt to node removals by rewiring itself.

Holme, *et al.* [136,137,184] consider strategies that likewise use degree and betweenness metrics, but in which the targeted node at time τ is the most highly ranked node that is identified in the net at time $\tau - 1$; that is, new targets are selected only after the net has had a chance to react to prior attacks.

Cascade attacks

Motter and Lai [181] study cascade-based attacks on complex networks. Focusing on the physical "loads" on nodes (which represent, say, the information processing requirements for communication flows), the authors show that when networks are able to adapt to local overloads, attacks often lead to a cascade of failures that may cripple the entire system. There are three basic mechanisms responsible for this: (1) *load redistribution*, in which the load of a disrupted (or failed) node is redistributed to other nodes; (2) *high-load node failure*, in

which redistributed flows are automatically channeled through lower-load nodes to compensate for high-load failures, and thereby exceed the network's remaining capacity; and (3) *heterogeneity*, which effectively guarantees that there will be few, if any, remaining high-load nodes to redistribute the remaining load. Operating together, these three mechanisms make it possible for even a single critical node disruption to trigger a system-wide cascade of failure. Watts [185] relates cascade failures to percolation problems and the spread of epidemics.

As a concrete example of cascade effects in networks, consider the electric power grid blackout that occurred in North America on August 14, 2003. On that date, an estimated 50 million people living in the midwest and northeast portions of the United States and Ontario, Canada, collectively lost roughly 61,800 megawatts of electric load. The unexpected severity of the event highlights the importance of understanding the relationship between local behavior and global properties.

In their analysis of the structural vulnerability of the power grid, Albert and Nakarado [186] find that—consistent with the more theoretical observations made above—while the grid is robust to most perturbations, failures (or disruptions) that affect critical transmission substations greatly reduce its ability to function. They conclude that “...the transmission hubs ensuring the connectivity of the power grid are also its largest liability in case of power breakdowns.”


Of course, what we know of cascading failures in electrical power grids (as one real-world example) cannot immediately be transferred to the behavior of terrorist networks. Terrorist networks are inherently “social” systems, and, as such, are considerably more adept at adapting to, and redistributing, any local (or global) load requirements. One expects the impact of high-load node failure, for example, to be much less severe for terrorist networks than for “hard-wired” (or programmed) physical systems. Nonetheless—from the point of view of *counter*-terrorist organizations tasked with disrupting the activities of the terrorist network—an understanding of the dynamics of cascade failures in complex networks provides valuable insights into how to attack terrorist organizations:

Using *ordered set theory* to break Al-Qaeda cells

Farley [187] begins his analysis of how to best “break” terrorist cells, by following Woo’s [188,189] advice and first asking the simpler question (than the more general one posed at the bottom of the last page): “How many nodes must be removed from the terrorist net before it becomes disconnected?” This obviously simplifies its more abstract predecessor by (1) effectively equating “disabling” a TN with “disconnecting”⁵⁰ it, and (2) positing that what “can be done” to a TN is to “remove its nodes.” This form of the question is thus both mathematically sound, and sufficiently well-defined to base a “break the network” algorithm on.

How do we decide which nodes must be removed? Farley argues that the information resident within the graph *itself*—i.e., only in its nodal and link structure—is *insufficient* for constructing an algorithm that meaningfully degrades the graph’s ability to function. What is needed, according to Farley, is a way to break the *chain of command* that exists within the graph, and not just the connectivity among independent, equally valued, nodes.

Figure 31. Hypothetical terrorist cell, shown *before* an attack on its connectivity (a), after an attack removes nodes *d* and *g* (b); and after an attack removes nodes *b* and *c* (c)



To illustrate why this is important, consider a hypothetical TN that consists of 15 members and has the hierarchical structure shown in figure 31-a. If terrorists labeled *d* and *g* were captured and removed

50. A disconnected graph consists of two or more subgraphs, all of whose nodes are isolated from all nodes in the other subgraphs [30].

from the net, the TN would be disconnected; it would be broken into four components (see figure 31-b). However, the *chain of command*—stemming from leader *a*, down to “foot soldiers” *m*, *n*, and *o*—would remain intact. Compare this strategy to the results of the strategy shown in figure 31-c, in which the terrorists labeled *b* and *c* are removed from the TN. Using this alternative strategy, the leader *a* and foot-soldier *o* are isolated; and, while a single, large component remains intact, its *de facto* leaders are only low-level operatives, two-steps removed from the TN’s leader.

Farley [187] applies *ordered set theory*⁵¹ to graphs to define an objective measure of effectiveness of removing *k* nodes from a TN. While graphs alone, via their links, are able to show only that nodes are somehow related—but not what the relationships are among the *roles* the nodes play within the graph—*ordered graphs* explicitly order the nodes, from top (highest ranking) to bottom (lowest ranking), according to their relative importance within the graph.

Once again, consider the notional TN shown in figure 31-a. In order theory, the *maximal node* represents the topmost node (or *nodes*; in this case, we have the single leader *a*); and the *minimal nodes* represent the lowest lying nodes (in this case, “foot soldiers” *i*, *j*, ..., *o*). A *maximal chain* is any path that starts from a maximal node, ends on a minimal node, and is ordered from highest to lowest rank. The terrorists labeled *a*, *b*, *e*, *g*, and *l*, for example—and in that order $a > b > e > g > l$ —form one of the 13 possible maximal chains (note that each also represents a possible chain of command in the TN):

- (1) $a > b > d > f > i$; (2) $a > b > d > f > j$; (3) $a > b > d > g > k$; (4) $a > b > d > f > i$;
- (5) $a > b > e > g > k$; (6) $a > b > e > g > l$ (7) $a > b > e > h > m$; (8) $a > b > e > h > n$;
- (9) $a > c > e > g > k$; (10) $a > c > e > g > l$; (11) $a > c > e > h > m$;
- (12) $a > c > e > h > n$; (13) $a > c > o$

If we wish to break the TN’s chain of command, we must remove at least one node from each of the TN’s possible chains of command. A set of nodes that intersects each maximal chain is called a *cutset* of the

51. An *ordered set* is defined as a pair (X, R) , where X is an unordered set, and R is a (*partial*) *order relation* of X ; R must be *reflexive* ($a \leq a$ for all $a \in X$), *anti-symmetric* ($a \leq b$ and $b \leq a$ imply that $a = b$), and *transitive* ($a \leq b$ and $b \leq c$ imply that $a \leq c$) [90]. In the discussion above, we take $R \equiv \leq$.

ordered graph. For example, the nodes b and c form a cutset, because each of the 13 maximal chains contain either b or c .

Using these definitions, Farley [187] calculates the probability that an N -member TN will be disrupted by a random removal of k members, $G(N, k)$:

$$\Gamma(N, k) = \frac{Cut(TN, k)}{\binom{N}{k}} = \frac{Cut(TN, k)}{N!/k!(N-k)!}, \quad (56)$$

where $Cut(TN, k)$ is the number of cutsets in the ordered set TN with k members, and the binomial coefficient, $\binom{N}{k}$, counts the number of ways of selecting k subsets out of N members. Farley's formula generally yields a significantly lower probability for disrupting TNs than does its more conventional cousin, in which "disruption" is equated with "disconnectedness" [190].

While Farley's formula is straightforward and easy to calculate for small networks, it is unfortunately of little practical utility in dealing with real terrorist networks. Aside from the realities that intelligence is rarely complete (or reliable) enough to adequately define the target, and that real TNs generally do not take the "clean" hierarchical form assumed in equation 51, cutset calculations are NP-hard [174, 191]. The virtue behind Farley's method, and the reason for presenting it here, is that it emphasizes the need to ascertain the *relative value* that each terrorist represents (as a dynamic component) within the organization, in the context of *information flow* within the network's self-organized *command and control* structure.

SOTCAC: *conceptual design*

“Actions and decisions...[are]... critically important. Actions must be taken over and over again and in many different ways. Decisions must be rendered to monitor and determine the precise nature of the actions needed that will be compatible with the goal. To make these timely decisions implies that we must be able to form mental concepts of observed reality, as we perceive it, and be able to change these concepts as reality itself appears to change. The concepts can then be used as decision-models for improving our capacity for independent action. Such a demand for decisions that literally impact our survival causes one to wonder: How do we generate or create the mental concepts to support this decision-making activity?”—John Boyd, Destruction and Creation

SOTCAC (Self-Organized Terrorist-Counterterrorist Addaptive Coevolutions) is a conceptual model that uses autonomous, intelligent agents to represent the components of *coevolving* terrorist and counterterrorist networks, and includes interactions among notional terrorist and counterterrorist intelligence agents, terrorist cells, training, logistical, and miscellaneous support networks, weapons and financial resource networks, and physical terrorist targets. SOTCAC serves as the core “logical engine” within which the *terrorist network* (TN) $\leftrightarrow$ *counterterrorist network* (CTN) coevolution takes place, adjudicates the communications between, and interactions among, all terrorist and counterterrorist agents, and provides a visualization of the emerging graphical structures. It is fervently hoped that SOTCAC, as it is developed, will prove to be useful not just for studying the dynamics of terrorist networks (for which it is being explicitly designed), but for helping usher in an entirely new class of general-purpose multi-agent-based dynamic graph models that can be used to explore the fundamental properties of *complex adaptive evolving networks*.

It is anticipated that a genetic algorithm will be used (see chapter seven in [11]) to search for agent-types and social networks that are “optimally” suited for performing a given mission. For example, in the case of TNs, the mission may be to destroy a set of assigned targets, and to do so as quickly and secretively as possible; the CTN’s mission is to prevent the TN from succeeding in its mission, and/or to eliminate the potential of the TN’s ability to do conduct future missions.

SOTCAC addresses, and is designed to allow analysts to interactively explore, three broad classes of problems:

1. **Dynamics.** The self-organized emergence, growth, and development of TNs, as driven by the general rules of social network formation, the personalities, skills, and needs of agents, and local and global mission constraints (such as balancing the need to find and marshal the resources required for a mission while maintaining stealth).
1. **Relationships.** The relationship between a TN's structure and internal dynamics (as characterized by, say, its topology and its local and global information processing capability), and its ability to perform its mission; or, succinctly, the relationship between a TN's *form* and *function*.
2. **Coevolution.** The coevolutionary process that fuels the mutual interdependency of actions by TNs and CTNs, as agents (and cells) belonging to each structure affect, and adapt to the actions of, agents belonging to the other. On the one hand, components of a TN always act in ways that minimize infiltration by counterterrorist agents (and other assets; see discussion in section **Counterterrorist Network**); on the other hand, the CTN's mission is to gather intelligence about the TN and its activities, and to otherwise infiltrate, disrupt and/or destroy the TN. From a CTN-analyst's point of view, the "problem" is to identify those key features of the TN (such as its topology and information processing capability) that can be used to pinpoint exploitable weakness and/or critical vulnerabilities to disruption or attack.

Modeling ontology

Figure 32 shows a schematic of EINSTEIN's modeling ontology. An "ontology" (as the term is used in artificial intelligence research) means a *specification of a conceptualization*, and is the set of fundamental concept definitions that describe the classes, structures and relationships that characterize a complex system of agents. (This ontology is introduced at this stage of the discussion as a reference to which SOTCAC's own conceptual design is later compared.)




Figure 32. Schematic of the decision-making process in EINSTEIN

Figure 32 shows that each agent “lives” on a battlefield (which thus constitutes its environment, and can sense and react to five classes of context-specific forms of information:

- *Internal features*, which represent both static and dynamic aspects of an agent’s unique personality and experience. Some of these features are visible to the outside world (i.e., by other agents), some are not. Internal features include energy level, fear of combat, health, morale, and personality “type” (inexperienced, timid, aggressive, obedient, leader, follower,...).
- *Environmental features*, which represent all local forms of information, such as terrain, intensity of combat, the presence and type of nearby paths and waypoints, and estimates of the tactical utility of specific battlefield positions.
- *Battlefield entities*, which refers to any member of an object class to which an agent can assign a motivation for either moving toward or away from that member. This class includes, as a subset, the class of (friendly and enemy) agents, battle specific field positions, own and enemy flags, waypoints, and patrol areas.
- *Object Features*, which include both absolute and relative measures. For example, if the entity is a specific enemy agent, absolute factors include firepower, health, position (near flag, in open, near boundary, etc.), threat level, and vulnerability. In the case of relative measures, examples include the distance between the two agents, relative health states, relative firepower, and relative vulnerabilities.
- *Communicated Information*, which refers to information either communicated to, or received by, an agent. While, in EINSTEIN, communication is modeled fairly crudely (and merely extends the effective range over which agents are able to probe their local environment), communication is a vital component in SOTCAC, where it is used as the foundation of network formation, dynamics and evolution.

SOTCAC's ontology

Figure 33 shows a schematic of the decision-making process in SOTCAC. It is deliberately drawn using the same component blocks used in the schematic of EINSTEIN's ontology (shown in figure 32), to highlight some of the differences and similarities between the two programs.

The fact that the same central diagram appears in both figures illustrates that, at their core, the ontologies for EINSTEIN and SOTCAC are exactly the same, and consists of: (1) *agents*, that are defined by both an internal personality and a set of rules by which they interact with an external space; (2) an *environment*, which consists of various objects and features that agents can sense and react to; and (3) agent *actions*, that are local optimizations over the set of possible moves, and which are predicated partly on an agent's "understanding" of its local space, and partly on an agent's unique motivations, tasks and goals. The system evolves by iteratively applying the same sequence of logical steps—i.e., *Sense::Process::Act*—for each agent. The basic template used for designing both programs is thus essentially the same.

The two programs also obviously differ in several important respects:

1. *Focus of central conflict—land combat versus terrorism.* In EINSTEIN, blue "fights" red in a conventional battle, while in SOTCAC, counterterrorist agents hunt for information about, and attempt to nullify the organizational and coordinating capability of, terrorist agents.
2. *Notional environment—a physical battlefield versus an abstract information space.* EINSTEIN's agents move about on a notional representation of a real battlefield, and are anchored to sites located on a fixed grid of possible locations; while SOTCAC's agents are also free to roam over a notionally physical space (a feature that is used both for mimicking "chance" meetings between possible recruits and recruiters, and for modeling the physical movement of active terrorist agents to reconnoiter fixed target locations), most of the dynamics takes in within, and is partly a function of, a mathematical space of graphs (which is used to model the growing and evolving social network ties among agents).




Figure 33. Schematic of the decision-making process in SOTCAC; compare to figure 32

3. *The details of the primitive properties and behavioral characteristics of their respective classes of agents and their mutual interactions* (each of which is defined in a manner consistent with its associated focus and environment). For example, in EINStein, agents are defined by such features as sensor range, movement range, and health, all of which are appropriate for EINStein’s “combat” domain. In contrast, SOTCAC’s agents are defined by features that, loosely speaking, are all attributes of social networking forces. That is to say, features such as *allegiance* (to a cause, or group of agents), *grievance*, *influence*, and *trust* (among others), collectively describe some of the main forces responsible for creating, growing, sustaining, and evolving social networks.

Design overview

SOTCAC uses agents to model the coevolution of terrorist and counterterrorist networks. All activity in SOTCAC takes place within two, tightly coupled—and dynamically reciprocal—domains:⁵²


- A *physical* domain, which represents a notionally physical space in which agents move and interact in a direct, face-to-face fashion (and in which the physical “targets” of the TN are also situated); and
- An *information* domain, which represents the abstract space that contains the TN’s dynamic social network at time t (as well as the CTN’s “best guess” as to what the topology of this social network is at time t).

Figure 34 shows a schematic of these two domains.⁵³

52. These two domains are selected more for expediency and general modeling constraints and requirements than for completeness. For example, the information domain represents a simplified distillation of a much larger set of social network domains that real terrorist networks inhabit; including *cultural*, *familial*, *financial*, *legal*, *political*, and *religious* spaces [196]. The purpose of *any* model, multiagent-based or not, is to identify, and derive insights from, the subset of *critical features that drive the real system*; it is not to provide an exhaustive taxonomy of parts. See R. Smith, “Counter Terrorism Simulation: A New Breed of Federation,” *Simulation Interoperability Workshop*, Spring 2002.

53. Many details of SOTCAC’s design are missing from this “first pass” schematic, but will be introduced during the ensuing discussion.

Figure 34. Schematic of SOTCAC's *physical* and *information* space representations of TN and CTN coevolutions; see text for details



The box at the top of the figure shows a configuration of T-agents (labeled Ta_i , and colored *red*), CT-agents (labeled CTa_i , and colored *blue*), support agents (and which represent sources of financial, logistical and financial resources sought after by T-agents; these are labeled Sa_i , and are colored *yellow*, with *red border*), and terrorist targets (labeled T_i , colored *grey*, with a *black border*). The left and right boxes along the bottom contain, respectively, the social network ties of TN and CTN's best-guess of what those social ties look like.

Note the CTN's knowledge/inference base is neither *complete* (since several components of the TN's actual structure are missing) nor completely *accurate*; the black squares and grey dotted lines between presumed terrorists represent incorrect information regarding the TN that CTN believes to be true. A key role for the analyst is to uncover "optimal" mixes of physical/virtual CTN-assets (along with the "rules" according to which they move and collect data, and that define how the CTN processes INTEL) for which the *difference between what the CTN believes to be the TN and the TN's actual structure is minimized*.

TN and CT-agent activity in the *physical* and *information* domains is "reciprocal" in the sense that the set of primitive behaviors of one side is mirrored (to the extent possible), by the set of primitive behaviors of the other. For every possibility of something being dynamically *created* (an agent, link, information, wire-tap, etc.), there is a reciprocal possibility of the same thing being *deleted*.

For example, just as CT-agents—while they maneuver in the physical domain—base some of their actions on the T-agent activity they perceive within their sensor range, so too do CT agents, as they maneuver in the same physical space, tailor their own actions, in part, according to whether they sense the presence of CT-agents. On a more abstract level, for every *Link* rule that specifies the exact conditions under which, say, two previously unlinked agents will forge a connection, SOTCAC possesses a complementary *Unlink* rule that specifies the conditions needed to disconnect linked agents.

Of course, reciprocal actions rarely occur simultaneously during any given coevolution. It is only by recognizing and seizing the more opportune possibilities, as they arise dynamically (and partly due to chance), that one side is able to "defeat" the other.

The most important interactions of the model take place in the information domain. For it is in graph space that the terrorist organization---viewed as a *complex adaptive system*---first emerges (and subsequently evolves) as a self-organized, dynamic network of aggrieved, disenfranchised, and/or repressed individuals, recruited to join the organization; of their cultural, religious, familial and/or general social ties; and of the command, control, communication and coordination links that naturally form, in a decentralized manner, within the organization as it inexorably moves toward accomplishing its mission.

Implicit in the figure is the fact (elaborated upon later in this section) that each agent has an associated set of, typically changing, sensory and behavioral properties that determine, among other things, how far it is able to *see*, what properties of other agents it may *know*, with whom it may *communicate*, the degree to which it *trusts* information that is communicated to it, and what *actions* it may take. Initially, before any social structures emerge, both terrorist and counterterrorist agents move about randomly. However, as agents begin making contact—meaning that they come within a threshold “contact” distance of one another—the process of registering, and exchanging, various kinds of information starts to take place as well.

For example, one terrorist agent, whose temporary role is that of *recruiter* (see below), may choose to link with a possible recruit, and to maintain an informal tie with that agent until the recruit has acquired a threshold skill level from a trainer, at which time the recruiter may “create a link” between the agent and a cell leader. A counterterrorist agent may “see” two suspected agents together and infer that the two agents are, in fact, operatives of the TN and that they have a strong connection. The CTa’s new INTEL is used to update the CTN’s *belief matrix* (using a method that is explained below), which defines the structure of the TN that the CTN believes it has at that moment. A counterterrorist that comes within an *eavesdropping distance* of two known Ta’s, has an associated probability, P_e , of eavesdropping on what one agent is communicating to the other.

SOTCAC’s design borrows several important elements from two existing multiagent-based models, both developed at CNA: (1) the EINSTEIN land combat model [16], and (2) an agent-based variant of an

older human-based wargame, SCUDHunt [17] (which is designed to explore information flow in command and control structures).

From EINSTEIN, SOTCAC inherits the basic conceptual template for administrating a heterogeneous set of agent “personalities” and personality-weight-based decision rules; according to which all actions are locally motivated, goal-driven optimizations over an agent-specific “action space.” While, in EINSTEIN, the action space is (at least notionally) *physical*—in the sense that agents’ “action” consists of choosing a battlefield position to move to, and engaging in combat with enemy agents at other positions—SOTCAC’s agents’ action space is both *physical* and *abstract*. SOTCAC’s agents are not only tasked with performing abstract information processing chores (that are considerably more involved than what EINSTEIN’s agents are capable of), but must decide both where to move in a physical space, and how to *alter the topology* of their local (social network) structure.

In SCUDHunt, agent personalities consist of parameters that define how an agent obtains, interprets, and uses game-generated information, and includes the interpretation of sensor reports, trust (of other agents), strike-plan logic, and sensor-placement logic. The state of the game is defined by a *belief matrix*, $B(x,y)$, which is a measure of an agent’s belief that a SCUD is at location (x,y) . As a game progresses, agents update the values of the components of their belief matrix. The way in which partial beliefs, derived from snapshots views of an agent’s immediate environment, are added to a SCUDHunt-agent’s current belief, proceeds in exactly the same way as EINSTEIN-agents combine the value of two or more environmental features to modify a component of their current personality weight vector.

SOTCAC inherits SCUDHunt’s logical structure that describes how INTEL data (that is extracted from a “ground-truth” battlefield view) can be used, in a mathematically well-defined manner, to augment, and refine, one’s *beliefs* (or concurrently updated “best guesses”) about what the “ground-truth” picture really is. In SOTCAC, “ground truth”—at each time, t —consists of the TN’s actual structure and set of activities taking place at t ; i.e., TN’s agent composition, link structure, the sets of parcels of information being communicated among agents, and its cells’ precise structure. No single agent, belonging to either the TN or CTN, has complete access to ground truth.

The CTN's data-fusion activities (which it must coordinate among N physical CT-agents and M virtual CT-assets) is not rendered directly by SOTCAC; instead, it is monitored indirectly by snapshots of its evolving *belief matrix*. The "belief matrix" represents the CTN's "best guess" as to what the TN looks like at time t . The set of action options that the CTN selects from at time $t+1$ —for example, *the node or link to "eavesdrop" on, the link to pump false information into, the link to jam, or node to infiltrate* (among other options; see discussion below)—is a function of what the CTN "believes" it knows about the TN at time t .


What *is*, and *is not*, being simulated

Figure 35 shows a schematic timeline of the coevolution of *terrorist* (TN) and *counterterrorist* (CTN) networks, as modeled in SOTCAC. Apart from summarizing some of the basic components and dynamics of the model (all of which are covered in detail in the following sections), the figure highlights, in grey, two important aspects of terror network analysis that SOTCAC *does not* attempt to model. Namely:

1. The early, nascent, stage, of the terrorist organization, in which various social, cultural, economic, political and religious factors conspire to spawn the first seeds of terrorism, but at which time the terrorist *network*, as a self-organized, self-sustaining entity able to conduct missions, does not yet exist, and
2. A much "later" stage (in the TN's evolution), at which point a mature, resource-laden TN finally emerges, and harbors sufficient manpower, skill and experience to conduct and successfully accomplish its mission.

SOTCAC focuses on the set of *intermediate steps* of this timeline, providing the analyst tunable parameters with which to explore various "dynamic forces" that guide a TN's self-organizing, emergent form; its ability to react, and adapt to, attack; and the coevolutionary processes that concomitantly shape both TNs and CTNs, as each tries to destroy the other.

Figure 35. Schematic timeline of terrorist network (TN) and counterterrorist network (CTN) coevolution in SOTCAC



While the patterns of behavior that appear at the periphery of figure 35—*pre-formative dynamics* on the left-hand-side, and the dynamics and mechanisms of the terrorist *attacks* themselves, on the right—are obviously equally as important to understand, to simulate them in a meaningful way requires a different set of social, dynamic, and military components, and neither is modeled directly by SOTCAC. The dynamical behaviors at the two end of this spectrum are, however, the subject of other recent studies.

For example, at the pre-formative end of the spectrum, MacKerrow [31] is using an agent-based program called the *Threat Anticipation Program* (TAP)⁵⁴ as the “engine” with which to explore the complex dynamics of militant terrorist group formation.

MacKerrow’s model includes such factors as social grievance, perceived oppression and corruption, various economic, cultural, social and religious influences, and searches for indicators of “terrorist instabilities” that occur when agents with high levels of social grievance have access to terrorist organizations with similar grievances. Casebeer and Thomas [32,33] apply a similar set of systems-level tools to model the general formation of violent non-state actors.

SOTCAC models this part of the TN formation timeline only to the extent that it includes an “ambient pool” of possible recruits to join the TN, and considers a few general motivations such as *grievance* and *risk aversion* (as factors in deciding whether to incorporate these contacts as trainees; see discussion below).

To help understand the “attack stage” of the TN timeline in figure 35, Pate-Cornell and Guikema [34] use systems analysis (as well as elements of risk assessment, decision theory and mathematical game theory) to model terrorist threat probabilities. Their work can also be viewed as a terrorist-threat-specific prioritization of possible defensive countermeasures. Their analysis carefully considers all of the relevant factors at play during (and immediately prior to) the “attack stage” appearing on the right-hand-side of the timeline in figure 35.

54. TAP is being developed at the *Defense Threat Reduction Agency* (DTRA): <http://www.dtra.mil/>.

Two recent studies that address the same “middle ground” of the timeline of TN evolution as does SOTCAC are those by Carley [35] and Raczynski [36]. Carley is developing what she calls a *meta-matrix* approach to conducting dynamic network analysis. The approach consists of combining knowledge management, operations research and social networks to represent a multidimensional view of the dynamical relations among agents, resources and tasks. The idea is to explore ways in which changes to one component of the meta-matrix cascade into changes in other components. SOTCAC’s design complements Carley’s method by examining an analogous cascading of effects between coevolving multidimensional representations of terrorist and counterterrorist networks.

Raczynski uses agents to model the dynamic interaction between terrorist and antiterrorist organizational structures that is, superficially, similar to SOTCAC. However, Raczynski’s agents are rudimentary, and interactions depend less on personality-weight-prescribed local optimizations than on probabilities (or tunable “rates” at which certain scripted events occur). Due credit must be given to Raczynski’s model as being one of the first agent-based simulations of terrorism to recognize the fundamental importance of, and to focus its attention on, the *evolution of the terrorist network structure*, and the consequences that evolution has on the terrorist net’s ability to conduct attacks.

Terrorist network

“[Terrorism]...The calculated use of violence or threat of violence to attain goals—political, religious or ideological in nature—by instilling fear or using intimidation or coercion. Terrorism involves a criminal act, often symbolic in nature, intended to influence an audience beyond the immediate victims.”—US DoD Directive 2000.12H

T-agents

The basic dynamical component of SOTCAC, as in EINSTEIN, is the *agent*, which embodies the properties, characteristics and behaviors of notional members of a TN (and the TN’s support network).⁵⁵ As do EINSTEIN’s combat agents, SOTCAC’s *terrorist agents* (or *T-agents*, for short) also represent a heterogeneous mix of personalities, motivations, and goals. However, unlike EINSTEIN’s agents’ relatively limited feature set—which was tasked only with describing local

movement and simple combat engagements—T-agents “live” in both physical and social domains, and process considerably more information; thus, T-agents generally require a more complex mix of features to tailor their unique behavior.

Additional features (not used by EINSTEIN but important to SOT-CAC) include *allegiance, experience, knowledge, leadership, rank* and *skill*. Each of these will be introduced and discussed in the sections that follow.

Consider the kinds of basic questions that an agent needs to answer:


- *What are my properties? Which properties are fixed, and define my core personality? Which properties change as I evolve during a scenario and accrue experience?*
- *What do I know...about myself?...about my environment?...about my local network...about the TN?*
- *What do I own?*
- *What do I need to do?*
- *Where should I move to in the physical domain?*
- *What is my mission? What is my task? What are the requirements for fulfilling my task and/or mission?*
- *What do I need...to find? What can I barter with?*
- *How do I get what I need? Where ought I to focus my search in physical space?*
- *Who am I connected to (i.e., who can I “talk” to right now?) Who do I know?*
- *Who else exists that I can communicate with?*
- *Who do I establish (or break) a connection with?*
- *What reasons do I have for establishing (or breaking) links? What factors must I take into account? How do I assess those features (using my personality)?*

55. The notional counterterrorist organization also consists, in part, of agents (called *CT-agents*) that are able to move around the same physical space as T-agents. The properties and behaviors of CT-agents are discussed in the section **Counterterrorist network**).

T-agent types

SOTCAC has seven, partially overlapping, classes of *terrorist* agents (see figure 36):

Figure 36. Schematic illustration of the basic components of a terrorist network, as represented in SOTCAC; see text for details



Recruit

Recruits, $\{r_1, r_2, \dots, r_{N_r}\}$, are members of an ambient “pool” of random members of the population that come into contact with, or are selectively targeted by, recruiters, $R_j \in \text{TN}$, to join the terrorist organization. Not having any *a priori* social network links, recruits initially ‘live’ only in the physical domain, and move about randomly (their movement being constrained only by a user-defined, and fixed, *movement range*, see below). Recruits represent a heterogeneous mix of raw personality traits (such as *allegiance*, *risk aversion*, *social grievance*, and a miscellany of *knowledge* and *skills*).

Whether or not a given recruit, r , becomes a provisional member of the TN (i.e., whether it is “recruited”) depends on two conditions: (1) r must first come in “contact” with a recruiter, R , (which is achieved by wandering into R ’s *recruit range*), and (2) the value of uniformly drawn random number, x , must exceed the *recruitment probability threshold*, $P_{\text{Recruit}}^*(r, R)$:

$$P_{\text{Recruit}}^*(r, R) = [1 - \gamma_{RA}(r)] \cdot \gamma_A(r) \cdot \frac{\deg(R)}{\text{Maximum}\{\deg(R_i)\}}, \quad (57)$$

where $0 \leq \gamma_{RA}(r) \leq 1$ is r ’s *risk aversion* index (see below), $0 \leq \gamma_A(r) \leq 1$ is r ’s *allegiance index*, and the last term, $\deg(R)/\text{Maximum}\{\deg(R_i)\}$, is an R -dependent *preferential link attachment* factor, that is included to ensure scale-free growth ([88; also, see discussion on page 57]). Thus, recruits are (a) more (or less) likely to join the TN if they possess a greater (or lesser) degree of innate *allegiance*, and if they are generally less (or more) *risk averse*, and (b) favor joining the TN via recruiters that have already successfully “recruited” other recruits.⁵⁶

Once recruited, new recruits have a direct, but tenuous, connection to the TN through their recruiter and the trainer to whom they are immediately sent for instruction. Recruits have no *skills*, *experience*, *rank* or *influence*. Before they become active agents, or *mission operatives*, recruits must first spend a *TrainingTime* period in the vicinity of their assigned trainers, and successfully evade capture by CT-agents.

A small fraction, $f_{\text{ TerroristPool }}$ of this pool of possible recruits consists of already trained professional terrorists that may be known to recruiters or cell leaders (or even to active operatives, as “trusted” members of their virtual ego-maps), and may be enticed to join a mission. Members of this subpool of recruits generally possess unique skills and experience; if they decide to join a cell they represent a “shortcut” to a cell leader who wishes to marshall his required manpower and resources as quickly as possible.

56. This same scale-free ‘*rich get richer*’ growth and formation rule also applies to TN cell formation; see discussion below.

Infiltrator

A certain (user-defined) fraction of the ambient pool of possible recruits consists of CT-agents that attempt to *infiltrate* the TN (by posing as recruits). If a CT-agent successfully infiltrates the TN in this manner, it can be subsequently used (by the CTN), within limits, to eavesdrop and otherwise disrupt local activities of the terrorist organization. A double agent constantly faces the specter of being discovered by the TN. as each of its actions (directed toward providing the CTN information about the TN that it is motivated to find; see discussion in the **Counterterrorist network** section) has an associated probability of being “discovered” by a T-agent.

Recruiter

Recruiters, $\{R_1, R_2, \dots, R_{N_R}\}$, are members of the TN whose function is to parse members of the population (that they come into contact with either through chance or by identifying certain desirable traits) to find individuals that may want to join the TN. Recruiters possess a heterogeneous mix of abilities and skills; and their degree of influence depends, in part, on their social network of leaders (seeking their assistance in meeting their own mission requirements) and trainers.

Trainer

Trainers, $\{T_1, T_2, \dots, T_{N_T}\}$, provide recruits with a default set of knowledge and skills; without which, recruits cannot be “promoted” to the status of “operative” and become active members of existing cells. Not all trainers are able to teach exactly the same set of skills, and each is part of his own social network of cell leaders and operatives.

Mission Operative

Operatives, $\{O_1, O_2, \dots, O_{N_O}\}$, under the “command” of a cell leader, represent the basic unit of “manpower” of the TN; i.e., the *foot soldier*. Members of this class represent the active terrorist agents that—upon being assigned a mission role—are tasked with, first, acquiring the requisite knowledge and skills for the mission; second, marshalling whatever resources are necessary to perform the mission; and, finally, with conducting the terrorist attack.

Operatives are characterized by a heterogeneous mix of abilities, skills and personality (such as allegiance, experience, influence, rank, and risk aversion).

Leader

Leaders, $\{L_1, L_2, \dots, L_{N_L}\}$, are responsible for creating, organizing and sustaining (the growth and adaptability of) cells, which are composed of operatives under their command. Leaders orchestrate and coordinate all cell activities necessary for the mission that is assigned to them, and may act as liaisons among operatives, support agents, trainers and the mastermind (i.e., the highest ranking member of the leader class). The probability that a given leader successfully accomplishes his set of mission-specific tasks, depends on his experience and influence; it also depends, implicitly, on the social networks he has already formed with other agents.

Note that cells, as modeled in SOTCAC, only notionally represent their real-world counterparts; referring only to the set of all agents that have at least one link (at any time) to a leader agent. In particular, differences and interactions among, say, command and control cells, logistics cells, intelligence and/or reconnaissance cells, and sleeper cells, are not explicitly modeled in any way.

SOTCAC, by default, includes one *senior* leader, who assigns missions to cells, and to whom cell leaders occasionally report back to. Senior leaders are *virtual*, and do not occupy any position in the physical space; they represent, notionally, the “spiritual” leaders of a terrorist organization, such as Osama bin Laden (Al-Qaeda), Abdullah Ocalan (PKK), or Yasser Arafat (PLO).

Support Agents

Support agents, $\{S_1, S_2, \dots, S_{N_S}\}$, are either active TN operatives, or are entities outside the organization that have links to an active TN operative. Support agents do not directly participate in conducting terror attacks; rather, they provide arms and weapons to operatives and/or leaders, as well as various levels of financial, logistical, and miscellaneous resource support. Support agents provide different levels of support, depending on their influence and experience.

They are motivated to “find” members of the TN who they believe have a need for the resources they possess. They do so both in the physical domain, by favoring positive movement weight values for operatives in need of resources, and in the information domain, by demonstrating an increased propensity for creating “need” and “query” links in their local social nets.

There are four basic types of support agents:

- *Financier*
- *Weapons specialist*
- *Logistics specialist*
- *General support*

T-agent characteristics

Agents are equipped with *sensors* (with which they observe and extract information from their immediate environment), an *ego-space* (which is their local map of social network ties; it also represents information about the TN’s overall structure that may be compromised in the event of a targeted attack by the CTN; see below), a *memory* (), *tasks*, and *actions* that it may select from at each time step.

Primitive

- *Ability*: measures an agent’s innate ability to perform its assigned tasks. Agents of lesser ability are more rapidly overloaded by activity (and thereby cease functioning for a time) than agents of higher ability.
- *Allegiance*: measures the degree to which an agent is loyal to the TN. The smaller the value of allegiance, the greater the probability that an agent will refuse to follow an order issued by a higher ranking agent (which has the additional effect of decreasing that agent’s own rank and value in the organization). In recruits, lower values of allegiance effectively lower the probability that a recruiter will select them to join the TN.
- *Independence*: measures the degree to which agents are able to operate on their own, without guidance from, or coordination with, agents of superior rank. Independence provides a local

glimpse of the degree to which the TN's command and control structure, as a whole, is organized as a hierarchy. Examples at both extremes exist in the real world. For example, while Al Qaeda is known to be strongly decentralized, local cells of the Jemaah Islamiyah tend not to operate on their own without specific orders from leaders.⁵⁷

Leadership is typically used to discriminate among recruits and operatives; cell leaders naturally possess a high value of leadership (see below), and are thus already prone to be independent.

- *Leadership (charisma)*, $0 \leq \mathcal{L} \leq 1$: measures the degree to which a leader is able to recruit new agents into his cell, as well as his ability to maintain unity and cohesion in his cell. Leaders with higher values of leadership require less overt coordination (via intermittent links with cell-subordinates that are subject to discovery) than do leaders with less charisma.
- *P-S inclination index*: the physical-social inclination index, $-1 \leq I_{P/S} \leq +1$, measures the degree to which the agent focuses its actions on the *physical* (P) or *social* (S) space.
- *Risk aversion*, $0 \leq \mathfrak{R} \leq 1$: measures the degree to which agents tolerate the risk of incurring the cost of failing to achieve the desired beneficial outcome of a possible action. *Maximally* risk averse agents (i.e., those with $\mathfrak{R} = 1$) take only those actions whose cost is zero (or minimal), without regard for expected return; *minimally* risk averse agents (with $\mathfrak{R} = 0$) take only those actions that maximize benefit, regardless of the cost involved. The propensities of agents with other values of risk aversion are interpolated between these two extremes.
- *Movement range*, R_{Move} : defines the maximum distance an agent can “move” from its current location (at time t) to its new position (at time $t+I$).

57. Jemaah Islamiyah is a militant Islamic group active in several Southeast Asian countries that's seeking to establish a Muslim fundamentalist state in the region. Because of this particular group is strongly hierarchical, the arrests of several of the Jemaah Islamiyah's leaders (in 2002 and 2003) may presage the death of the entire organization [25].

- *Sensor range*, R_{Sensor} : defines the range of an agent's *vision* in the physical space. Agents sit at the center of an R_{Sensor} -by- R_{Sensor} box, and “see” everything that occupies any of the $(2*R_{Sensor}+1)^2$ sites within this box. Note that this does not imply that an agent correctly identifies what is present at a given site, only that an agent has the *potential* of registering another agent within its sensor range.

Composite

- *Anti-capture index*: measures the degree to which a T-agent is able to resist or negate an “attack” (i.e., *capture*) by a counterterrorist agent. A T-agent's ability to resist attacks increases (i.e., the index value increases to a maximum value of one) as the T-agent successfully evades capture, acquires experience and skill.
- *CTa-detection probability*, $P_{Ta \leftarrow CTa}$: probability with which a T-agent is able to correctly “recognize” an otherwise unidentified agent, within its CTa-detection range, as a CT-agent. This probability increases as the T-agent acquires experience and skill.
- *CTa-detection range*, $R_{Ta \leftarrow CTa}$: range (in the physical domain), at which a T-agent is able to “see” a CT-agent. Whether, or not, the T-agent correctly identifies this agent as a CT-agent (and thereby adapts its personality weight vector values to “avoid moving toward” that agent), depend on the Ta 's CTa-detection probability. The CTa-detection range increases as the T-agent acquires experience and skill.
- *Value*: measures the total, dynamic, value that agent represents to the terrorist organization. It is a weighed sum of ability, allegiance, experience, leadership, rank, resource marshalling and skill.

Dynamic

- *Degree*: a topological index of the number of active links an agent has with other agents at time t (i.e., a measure of current communications activity and workload).

- *Position*, $\bar{x}(t) = (x(t), y(t))$: defines an agent's (x, y) position in the physical space at time t .
- *Rank*: measures the relative authority the agent has, and can exert over other agents of lesser rank, within the organization.
- *Resources*: which is an interim catalog of acquired resources required by an agent's assigned mission; specific classes of resources include *skills*, *weapons*, and *money*.
- *Skill*: which measures the cumulative information acquired by an agent via direct training or indirect learning through its social network ties
- *Visibility*: measures the degree to which a T-agent is "visible" to a CT-agent; and naturally scales accordingly with an agent's ability, skill and experience.
- *Work load*: which measures an agent's "tasking burden," and is a function of the number of current and past social ties, and number of unfulfilled mission requirements.

Social network maps


LoTo-map

With every T-agent, s , there is an associated *LoTo-map* (short for *local topology map*), $L_t(s)$, that represents an instantaneous snapshot of s 's local neighborhood of direct and indirect contacts at time t ; i.e., $L_t(s)$ encodes a s -centric view of the local topology of the TN. Since $L_t(s)$ contains no information regarding how s *interprets* (or *organizes*) the elements in its neighborhood (this processed form of information is instead stored in s 's *ego-map*; see below), the *LoTo-map* does not necessarily reflect what s *knows* about its neighborhood.

Figure 37 shows a schematic of $L_t(s)$, centered on s and containing all agents out to a distance $D = 3$ removed from s : (1) the set $\{a_i(s)\}$ represents all of s 's nearest neighbors (i.e., agents with whom s is currently linked with); (2) the set $\{b_{ij}(s)\}$ represents all of s 's *next-nearest* neighbors (i.e., all agents X that are a distance $\text{Dist}(s, X) = 2$ away from s), indexed according to which nearest-neighbor a given next-nearest neighbor is connected to (for example, $b_{31}(s)$ is the

“first” nearest-neighbor of s ’s “third” nearest-neighbor, $a_3(s)$); and (3) the set $\{g_{ijk}(s)\}$ represents all of s ’s *second*-nearest neighbors (i.e., all agents X that are a distance $\text{Dist}(s, X) = 3$ away from s), indexed according to which next-nearest-neighbor a given second-nearest neighbor is connected to. Note that certain subsets of the indexed sets $\{b_{ij}(s)\}$ and $\{g_{ijk}(s)\}$ may overlap; although, from s ’s point of view, the actual node that is being referenced is always unambiguous.⁵⁸

Figure 37. Schematic illustration of σ ’s *local topology map* at time $t (=L_t(s))$; see text for details



For example (as illustrated in figure 37), $b_{14}(s)$ and $b_{21}(s)$ —indexed according to two different nearest-neighbors of s —represent the same agent that is a distance $D=2$ away from s .

Note that two components of $L_t(s)$ comprise part of s ’s *knowledge state* (which may be “discovered” by CT-agents, and hence the CTN, in the event that s is captured): (1) the set of direct contacts (i.e., s ’s immediate neighbors at time t); and (2) the set of degrees of s ’s neighbors, $\{\deg(a_i(s))\}$ (i.e., it is assumed that the maximum information s has,

58. SOTCAC always keeps a “ground truth” list of the nodes and links of the TN; in the form of a master *adjacency matrix*, $M_{ij}(t)$, indexed by arbitrary, but persistent, labels for all nodes (the number of nodes, of course, is a function of time). Thus, while indexed arrays—such as the sets $\{a_i(s)\}$, $\{b_{ij}(s)\}$, and $\{g_{ijk}(s)\}$ —are convenient because they simplify formal discussions of how agents process information and adjudicate moves, it is important to keep in mind that they *do not* represent the actual form of information as it is processed by SOTCAC on the source code level. If SOTCAC’s master node-label list is $\{n_1, \dots, n_N\}$ (at time t), for example, SOTCAC “knows” that s ’s neighbor ($=a_3(s)$), say, is really just node n_j .

and may reveal, about any of its neighbors is the number of contacts that a given neighbor has; but not the identity of those contacts).

Active social space

The LoTo-map defines the *active social space* that is available to s at time t , in which s may “search” for information and required resources. For example, s may send any of the following queries to any of its immediate neighbors: (1) “*Do you have X?*” (where “ X ” is a particular resource); (2) “*What do you possess?*” (which requests an explicit inventory of possessions); (3) “*Who do you know?*”; (4) “*Send me X*”; and (5) “*Have b send me X through you.*”

For example, if s queries α_3 with “*Who do you know?*” (which takes one time step in the simulation), α_3 responds (on the next time step; *if it chooses to*) by sending s the list $\{b_{31}(s), b_{32}(s), b_{33}(s)\}$, and s , in turn, may elect (on the following time step) to communicate directly with any of α ’s neighbors; i.e. to make its prior b ’s (whose identities were unknown by s), s ’s new α ’s. On the other hand, depending on s ’s personality, s may instead elect to have a resource X (possessed by some b) passed from b to s , indirectly via α (who acts as “broker”). The tradeoff is one between *efficiency* (or *speed* of transmission) versus *vulnerability*: if a resource is transferred directly from b to s , the transfer is fast (requiring a minimum of one time step and a maximum that is determined by the type of resource being transferred to complete) but creates a temporary new link between b and s that may be visible to surrounding CT-agents; if the resource is, instead, transferred indirectly, first from b to a , followed by a transfer from a to s , no new links are introduced, thereby reducing the probability of being discovered by a physical CT-agent, but the transfer takes at least twice the time (possibly more, if during the transfer, a ’s “attention” is temporarily diverted to other matters).

Ego-map

An agent s ’s *ego-map*, $E_t(s)$, is a dynamic, local map of s ’s existing and prior social contacts at time t ; and is both an extension, and refinement, of $L_t(s)$. $E_t(s)$ is also, implicitly, a pointer to where other known agents are located (see below), a reference for potential future contacts, and a general aid for finding information and/or other mission-

critical resources. $E_t(s)$ represents what s “knows” about the TN from its local (i.e., ego-centered) vantage point, categorized according to time (past, present and potential future). To appreciate the difference between s ’s *Loto* and *ego* maps, note that while $L_t(s)$ contains only the raw s -centric view of TN’s local topology (or, information about “who is connected to whom”), $E_t(s)$ contains s ’s *accrued local interpretations* of its evolving local neighborhood (or, information regarding “who s knows about either from prior direct contacts or indirectly via other channels). For example, s ’s ego-map contains a list of agents with whom s may choose to establish a direct link (and can do so immediately, at any time, without going through intermediate contacts to gain the identity of the other agents); s ’s LoTo-map does not.

Figure 38. Schematic illustration of the four basic categories of contacts in a T-agent σ ’s ego-map, $E_t(s)$, at time t ; see text for details




Figure 38 shows a schematic of the four basic sets of contacts to which any of the T-agents in $E_t(s)$ may belong at time t :

- *Trusted Prior (TP) contacts*,
- *Current Contacts (CCs)*,
- *Terminated Contacts (TCs)*, and
- *Potential Contacts (PCs)*.

The color-coding of the sets and links in the figure do not intentionally convey any meaning (beyond that of graphically distinguishing among the different contact sets), save for the solid black links between s and the set of CCs: these are the *only* links—at time t —that are vulnerable to detection (and possible compromise) by CT-agents. As long as s itself is invisible to the CTN, all of the links in $E_t(s)$ *other than those between s and its CCs* are also invisible to the CTN and its CT-agents.

Trusted Prior (TP) Contacts

TPs represent the network of family, friends, acquaintances, and/or other past associates that an agent knows and trusts (and has possibly had numerous contacts with). It is a *virtual* network because its members and structure are both *a priori* unknown to the CTN. CT-agents cannot “see” the links between any T-agents and their networks of TPs. Such links may only become “visible” to the CTN in two ways: (1) if a T-agent chooses to activate virtual links to one or more TPs, and a CT-agent is in a position (and has the requisite ability, skills and experience) to detect the activation; and (2) the agent s is either captured or converted into a double-agent (in which case the information content of s ’s ego-map, including s ’s virtual network of TPs, is subsumed by the CTN and incorporated into its *belief matrix*; see **Counterterrorist network: beliefs**).

Ericson [85] discusses the importance of the role that trusted priors generally play in secret societies. More recently, in the context of the 9/11 terrorist attacks, Krebs [193] shows that the 19 hijackers all came from a network that had high closeness and degree centrality; Voss and Joslyn [194] reveal that the “Hamburg Cell”—which was the coordinating center of the entire operation—was also largely based on trusted prior relationships. An invisible network of trusted prior relationships is what renders terrorist networks both hard to detect and resilient [195].

An important research issue to explore is the degree to which the efficacy of the TN depends on the size and topology of the “seed” TP contact network. The two extreme cases occur when (1) no recruit that subsequently joins the TN has any TPs (so that the emerging TN is relegated to using only those assets that its components are able to “dis-

cover” on their own, “out in the open”; i.e., there are no “invisible” components for the TN to exploit), and (2) all recruits effectively “know” all other recruits before joining (so that all members of the TN have virtual access to all other members). Just as a starting *spatial disposition* of combat-agents is critical to consider in designing a scenario in EINSTEIN, the seed TP network is a critical factor in SOTCAC for exploring the coevolution of TN and CTN systems.

Current Contacts (CCs)

CCs constitute the set of agents with whom a given agent, s , is currently linked (i.e., the set defined as $a_i(s)$ in s ’s *loto-map*; see above). While a particular link between s and some other agent CC_i , $l(\sigma, CC_i)$, may assume a variety of forms—for example, it could represent a face-to-face meeting, an email or telephone exchange, or a correspondence via conventional mail (see **Communication** below)—its relevance to $E_t(s)$ is embodied entirely in the fact that $l(\sigma, CC_i)$ exists at all: links between s and a CC are always vulnerable to discovery, intrusion and/or destruction by the CTN.

Having open communication links with CCs offers both (short- and long-term) benefits and drawbacks to s . Benefits include the trade for, or acquisition of, the information, skills, and other miscellaneous resources that s needs to accomplish its assigned mission; as well as simply augmenting s ’s evolving social contacts and establishing a growing network of trusted contacts that s may later choose to cooperate with. Drawbacks include exposing s (as well as s ’s linked CC partners) to detection or abduction by a CT-agent, and adding to s ’s overall workload.⁵⁹

Terminated Contacts (TCs)

TCs are T-agents that s has had prior contact with (during the current run) but to which it is not linked with at the current time t . Depending on how those prior contacts were established (and their

59. Each link represents one unit of “work” that is assigned to an agent; the effective rate at which an agent is able to “assimilate” communicated information (and/or channeled resources) diminishes with increasing workload, and depends on the agent’s innate ability and experience.

type and duration), the CTN may or may *not* be aware of their existence. However, once an agent becomes a member of s 's set of TCs, it effectively becomes part of s 's extended virtual network that TPs are also part of: i.e., s is “aware” of a TC's existence, and may choose to reactivate a link to it if s perceives a need to do so; but, as long as a link between s and one of s 's TCs remains virtual, it is hidden from CT-agent sensors. As for TPs, s 's virtual net of TCs may only become “visible” to the CTN if a CT-agent is in a position to detect the activation of one or more of those links, or if s is either captured or converted into a double-agent.

Potential Contacts (PCs)

PCs constitute the set of agents that s is neither currently linked to nor has been linked to in the past, and about which s knows—via intermediary agents belonging to s 's TP, CC, or TC sets—only the fact that they “exist.” For example, in figure 38, s knows that PC_4 exists, via PC_4 's link to CC_1 , one of s 's current contacts.

Agents belonging to s 's PC class represent potential sources of information, resources and other contacts (currently unknown to s) that s may choose to establish communications with. For example, s may query CC_1 about what CC_1 “knows” about PC_4 (for example, “*What skills does PC_4 possess?*” or “*Who does PC_4 know?*”), and thus establish the parameters necessary for deciding whether or not to link (and barter, negotiate, or cooperate) with PC_4 .

T-agent personality

As in both EINSTEIN and SCUDHunt, SOTCAC's agents all possess a unique, dynamic, vector-valued *personality*, the components of which define the relative value an agent assigns individual tasks and/or motivations; and thereby regulates how agents behave. However, unlike EINSTEIN's agents, all of whose decisions, and therefore personality weights, are confined solely to one battlefield, SOTCAC's agents live in two spaces and thus require two personality vectors: (1) $\vec{W}_{Phys}$, that determines how they act in the *physical domain*, and (2) $\vec{W}_{Infor}$, that defines how they act in the *information domain*.

Figure 39. Schematic illustration of the different distance functions used in SOTCAC's *information* and *physical* domains




Figure 39 illustrates, schematically, these two coupled spaces, and highlights the different metrics that are associated with each.⁶⁰ In the physical domain—shown on the bottom—distances are measured using the conventional Euclidean distance function:

$$\text{Distance}_{\text{physical}}(A, B) = \sqrt{(x_A - x_B)^2 + (y_A - y_B)^2}, \quad (58)$$

where A and B are any two points in the space, and x_A , x_B , y_A , and y_B are their x and y coordinates, respectively. The farther two agents are from each other—on the two dimensional grid—the less likely they are either to “see” one another or interact in any way.

60. The information domain in figure 39 depicts only a “collapsed” view of a much larger multidimensional space that encompasses the possible kinds of agent ↔ agent links (such as face-to-face meetings, phone, email, and internet chat rooms); not too mention the even larger set of social network domains that real terrorist networks inhabit, including *cultural*, *familial*, *financial*, *legal*, *political*, and *religious* spaces [196].

In SOTCAC's information domain (or social network space)—appearing on the top half of figure 39—distances are measured using the graph metric function:

$$\text{Distance}_{\text{Information}}(A,B) = \text{Minimum}_{\{\text{Paths } \mathcal{P}\}} \{ \# \text{ links between } A \text{ \& } B \text{ on } \mathcal{P} \} , \quad (59)$$

where the *Minimum*{...} is taken over all paths ($= \mathcal{P}$) between *A* and *B*. For example, figure 39 shows that while the physical distance between the agents labeled *A* and *B* spans essentially the entire space, and is therefore quite large ($\text{Distance}_{\text{Physical}}(A,B) \sim 28$), the distance between the same two points in the information domain measures five units, because the shortest path between them contains five links. The figure also shows that despite also being relatively far apart physically, the points *C* and *D* are direct neighbors in information space.

Just as SOTCAC's global dynamics is driven by the interplay between two spaces (namely, the coevolution between what the TN *does* and what the CTN *perceives* the TN is doing), the TN's local actions result from a tight coupling between the physical and information spaces. While some actions might require close physical proximity, other actions may preclude close contact, and be conducted covertly over long-distance communication lines (for example, when cell leaders want to coordinate the activities of their cell members, but do not wish to compromise their identity or location).

Physical domain

T-agents use the physical domain primarily for three purposes: (1) *training*, which requires that unskilled recruits be present at a “training camp” for a threshold period of time, (2) *face-to-face* meetings with other T-agents or supporting agents (that maximizes the reliability of information flow, but also renders the participating agents vulnerable to CT-agents, and (3) *reconnaissance* of assigned targets (which can motivate T-agents to either approach or stay-away-from certain areas of the physical space, depending on the dynamic context and other components of the T-agent's personality).

T-agents generally try to avoid being “seen” by other agents they have either recognized as belonging to the CTN or suspect of being CT-

agents. T-agents are especially fearful of being physically surrounded by (what they *believe* to be) the quorum number of CT-agents needed to capture them ($= N_{Capture}^{TN} (Ta \mapsto Cta)$). Thus, the “do not get caught” rule acts as a trigger event:

If the number of CT-agents you are surrounded by exceeds the quorum number, give 100% weight to moving away from CT-agents at maximum speed, and temporarily set all other physical movement weight vector components to zero.

Information domain

Minimize I-visibility

Analogous to a T-agent’s desire to “not be seen” by CT-agents in the physical domain is an innate tendency to minimize its visibility in the information domain. In practice, this means that an agent must weigh its need to communicate with other T-agents (in order to acquire resources, coordinate activities, etc.) with the requirement that all communications must be kept to a minimum. The dynamic effect may be regulated using an agent’s *risk aversion* parameter.

Maximize information gain

During any transaction with other agents, an agent is motivated to maximize expected information (and/or material resource) gain.

Maximize compartmentalization

Defines the degree to which a T-agent wishes to constrain *all of his actions* (physical and social) to members of his own cell. At one extreme (weight value = *one*), a T-agent never associates (or forms any links) with another T-agent that does not belong to his own cell; i.e., only cell members are visible (in both the physical and social domains). At the other extreme (weight value = *zero*), a T-agent does not discriminate between cell type, treating all T-agents belonging to different cells equally.

Social network metric propensities

T-agents may use many of the social network metrics introduced earlier in this paper (see section **Complex networks: metrics**, page 83) to tune their behavior in the information space. Leaving the details of how this is done to a later section, the basic idea is to endow T-agents

with a “social sense” that depends on what communicative properties they would like to see in their local neighborhood.

For example, some agents, who may lack the requisite skills for forging links or simply be inexperienced, may seek to form ties only with other a few other agents who are physically close by and whom they already know well. Other, more experienced, agents, with entrepreneurial skills, may seek to establish a local link structure that creates structural holes (see discussion beginning on page 102) that they can exploit. Still others may have personalities that render them strongly motivated to play the “central” role in cell-to-cell bargaining among operatives and support agents. In all of these cases, a given CT-agent’s information-space personality is defined by an appropriate set of values of basic social metrics such as *degree*, *betweenness*, and *centrality*.

T-agent actions

T-agents (and, as we shall see later, especially CT-agents) always behave according to what they *perceive* to be true about their local environment; not necessarily what their local environment’s true state is. For example, a T-agent may fail to “see” a nearby CT-agent and/or erroneously tag another agent as belonging to the CTN.

Movement

T-agents move about the physical space in exactly the same way as agents in EINSTEIN move about the battlefield (albeit driven by a different set of primitive motivations): weighing expected utility versus risk, agents adapt their personality to their state and context. For example, new recruits may want to stay close to trainers. Mission operatives may also need to temporarily cluster for face-to-face meetings. On the other hand—environmental conditions permitting—all cell members may generally elect to be as far apart from one another as possible so as to elude detection.

Acquiring manpower: recruiting

Recruiters and cell leaders are always motivated to find new recruits. However, the degree to which a leader focuses attention on recruiting, relative to his other tasks, is inversely proportional to how close the number of his subordinates are to his *CellMaxNum* value. A

leader's primary task to ensure that the requisite manpower (for the mission assigned to its cell) is always available.

Acquiring skills: training

Both *recruits* and *mission operatives* must receive training before their rank and/or skill index are upgraded: recruits need basic training prior to being “activated” to T-agent status (which they receive from a physical collocated *trainer* during a training interval $Dt_{Training}$); mission operatives may need to acquire a set of specialized skills that are required for their cell's assigned mission.

Basic training, notionally, represents the acquisition of such skills as familiarization with basic weapons, grenades, laying and fusing landmines, cover and concealment, use of false documents, and training in foreign culture survival. Specialized skills include training in mortars and rockets, grenade launchers, urban battle tactics, sniper fire, bomb making, creating improvised explosives, and camouflaged/stealth operations (see [197]).

Acquiring resources

Most of the activity of mission operatives, once they have been assigned a role in the mission for which their cell leader is responsible, is to acquire the required resources to conduct the mission; resources include *weapons*, *money* and *skills*. This is done by using their *LoTo*-maps (*local topology map*; see page 150) to locate resources (and/or members of the TN who know where a given resource is located), and, if necessary, bargaining with other agents to acquire what they need.

Communications

Communications are an integral component of SOTCAC (and the basis of the TN's social-networking functions), and are defined as the combined set of rules and actions by which two agents establish, maintain, and mutually transfer information and/or material resources via a “communication link” that connects them. SOTCAC includes a variety of communication *types*, characterized by *content*, *directionality*, *strength*, *duration* and *visibility*.

Briefly, communications are predicated on three basic motivations (from a T-agent's point of view): (1) *Coordination* (i.e., "Who do I need to be linked with in order to coordinate and/or complete my mission tasks?"); (2) *Acquisition* (i.e., "Who do I need to link with in order to get X?", where "X" represents a *skill*, *resource* or *weapon*), and (3) *Stealth* (Who must I break off communications with to reduce the risk of discovery or intrusion?).

Of course, each of these motivations entails additional considerations (as well as finer contextual distinctions) that must be taken into account. For example, acquisition may be either *direct* or *indirect*: if *direct* – meaning a direct transfer or exchange of resources between two agents – the exchange may be quick but is vulnerable to detection by CT-agents; if *indirect*, the exchange may be effectively cloaked from outside discovery (albeit not completely, since the intermediate steps of the exchange may still be detected), but will require additional time to complete (a channeling of resources or information through one agent-to-agent link takes one step to complete). How a given acquisition proceeds – between two specific agents – depends, in part, on the two agents' personalities (for example, how risk averse are they?), and, in part, on the general dynamic context. For additional details, see discussion in section **Communications**.

Defection

Agents have a probability of defecting to a CT agent (which is inversely proportional to allegiance, experience, rank and value). Since the social cost of forsaking one's commitment to the TN is, in the real world, very high, the probability of detecting is generally low. They must be within a "defection range" of an active CT agent. In the event that a T-agent erroneously attempts to defect to another T-agent, he is removed from the TN (along with his social net, knowledge, skills and whatever resources he possessed at the time of defection).

If a CT-agent is close enough to the defecting T-agent at the time that agent *erroneously* defects to another T-agent (i.e., the defecting agent is acting under the mistaken belief that the agent he has defected to is a CT-agent), the *observing* CTN's belief-map is updated with the INTEL regarding the TN's removal of the node (an act that is automatic upon one of its agents observing another defect).

Cells

Promotion

If an “attack” on the TN by the CTN results in the destruction of a cell leader (along with that cell leader’s ego-map), one of that leader’s prior subordinates (nominally, the one with highest rank) may be elevated in rank to assume the leader role in his cell. A promotion also leaves a “hole” in the cell; namely, the node previously occupied by the new leader. The new leader’s first task, therefore is to restore the value of the hole. For example, if the new leader was previously a mission operative, he must now find a new mission operative to replace his role in the cell’s mission, and provide him with whatever additional skills and resources are required to complete the mission.

Cells are groups of T-agents, all of whom are subordinate to a single cell *leader* (which is a type of T-agent; see above). SOTCAC’s cells loosely represent their real-world counterparts, but make no operational distinction between ostensibly different kinds of cells (such as between command and control cells, reconnaissance, INTEL or logistics cells). Likewise, “sleeper” cells are not currently modeled. Thus all cells in SOTCAC are notional *tactical operations cells* that perform a variety of overlapping functions.⁶¹

While the members of a cell may have latent (or *virtual*) links to members of the terrorist organization that are not part of the same cell (for example, prior existing friendship and/or kinship bonds)—in fact, some T-agents may actively exploit such contacts while pursuing desired skills and resources—cell members are, by default, motivated to confine their social networking to other cell members.

The *strength* of a cell, which determines its ability to channel its manpower to acquire resources, is defined, in part, by the leadership and influence exerted by the cell leader, and, in part, by its *adhesion* and *cohesion* indexes (see below).

61. The absence of an explicit model of the dynamics of specialized terrorist cells does not necessarily incur a loss in generality. Al-Qaeda, for example, often uses combined calls to maximize the utility of existing manpower [197].

Size

A cell's maximum size, *CellSizeMax*, is a user-tunable parameter value, and defines the maximal number of subordinates that may be assigned to a cell leader. Of course, the size of real terrorist cells varies greatly, and ranges from a single terrorist (for example, a suicide bomber), to small cells consisting of less than five members (which are useful for abductions and hijackings), to larger cells that consist of up to 20 members or more (which appear to be preferred by Al-Qaeda) [198]).

The actual cell size that forms during a run in SOTCAC depends on the manpower requirements of the cell's assigned mission; the cell size must be less than *CellSizeMax*. The cell leader remains motivated to "search" for, or find (through recruiters or other T-agents on his ego-map), agents to join his cell as long as he perceives that his mission manpower requirement is not satisfied. Once the manpower requirement is met, the cell leader is no longer motivated to achieve that mission goal. However, if one or more agents are subsequently captured by the CTN, the cell leader resumes his search for new members.

Adhesion

White and Harary [199]—in a paper discussing the general structural cohesion within social groups—consider the relationship between connectivity and density. In particular, they introduce a concept called *adhesion*, which is a function of edge-connectivity, and is measured by the minimum number of edges that must be removed in a connected group to result in its disconnection. White and Harary show that a graph (possessing any degree of edge-connectivity) may be disconnected by a removal of a single node. This means that "*the unilateral power of actors can be high even when there are many relations connecting people.*"

Colloquially speaking, there are two related social "forces" that tend to keep groups together:⁶²

- *Cohesive forces*, that depend on the number and strength of many-to-many links within a group (see **Cohesion** below), and

- *Adhesive forces*, that measure the number and relative strengths of many-to-one links that exist within a group. Adhesion depends most strongly on the cell leader's *leadership* (or *charisma*).

White and Harary [199] state that...

"...what holds the group together where this is the major factor in group solidarity is the strength of adhesion of members to the leader, not the cohesiveness of group members in terms of social ties amongst themselves...The key to structural cohesion thus rests in how the pattern of relations makes unilateral action impossible."

Generally speaking, a cell is adhesive to the extent that the social links among its members are pairwise resistant to being disconnected.

Coordination Strength

Recall that a cell's leaders's chief task is to ensure that the manpower, finances, skills, weapons, and other resources (such as logistical support) needed to conduct the cell's assigned mission are all marshalled as covertly and quickly as possible.

Since these two requirements—*stealth* (to prevent accidental "discovery") and *speed* (to be the first cell to conduct its assigned mission and to boost the cell leader's experience, rank and influence)—are generally antithetical, the cell leader must weigh the benefit versus cost of each approach. SOTCAC provides a tunable parameter called *coordination strength*, $0 \leq \gamma_{\text{Coor}} \leq 1$, that defines the degree to which a cell leader is willing to sacrifice stealth for efficiency.

$g_{\text{Coor}}=0$ means that the cell leader *does nothing* to directly coordinate the marshalling of resources among his subordinates. $g_{\text{Coor}}=1$ means that the cell leader plays the central role in coordinating the search for resources for all cell members.

62. As pointed out by Fellman and Wright [18], the difference between these two forces is far from academic; particularly in the context of marshalling available counterintelligence resources, when it is important to appreciate the operative qualitative difference: it would be equally ill-advised to remove the strongest member(s) of a strongly adhesive cell as to target the leader of a strongly cohesive one.

Cohesion

The degree to which a cell is able to maintain its integrity that is a function of the sense of unity and collective identity that gradually builds-up among its members. This depends on the existence of many, strong many-to-many ties within the group; and are established, for example, by face-to-face meetings.

Recall Farley's [187] use of ordered set theory to break apart terrorist cells (see discussion starting on page 123). The key observation is that a terrorist network's vulnerability to being disconnected—i.e., its *cohesive strength*—may be monitored by measuring the net's *cutset* resistance. Recall that an (i,j) -cutset defines a subset of nodes such that every path that connects nodes i and j passes through at least one node of the cutset.

A cell's adhesive and cohesive strengths are—somewhat counterintuitively at first—global weaknesses, if assessed from the point of view of an overall vulnerability to some counterterrorist attack. On the one hand, in terms of a cell's ability simply to function as a group, being strongly adhesive and/or cohesive implies that the cell contains a highly redundant—and, hence, robust—network of communication pathways. On the other hand, that same robust redundancy increases the likelihood that parts (if not the entire structure) of the cell may be inadvertently revealed to counterterrorist agents.

Social network links

Just as physical movement lies at the core of EINSTEIN's dynamics, *communication* and *social network relationships* play central roles in SOTCAC (though movement also plays an important auxiliary role).

A link in SOTCAC is any kind of relationship between two agents. It can therefore take on a variety of forms and have many different properties, some of which depend also on context. For example, some links, such as kinship and/or friendship bonds, exist solely within the ego-maps of agents, and are not directly accessible to (nor can they be infiltrated by) CT-agents.

Other links, such as email exchanges, telephone calls or face-to-face meetings, take place in the open, and may be “observed” either by a physical CT-agent (that is close enough to “eavesdrop” on a conversation, or surreptitiously using one of several virtual INTEL assets under the CTN’s control).

Before deciding whether to establish a link—of whatever type—with another agent, each agent must first weigh the relative *benefits* and *costs* of doing so; each agent does so according to its own unique, and dynamic, personality. Benefits include factors such as *expected information gain*, *monetary transfer*, and *weapons acquisition*; costs typically include estimated *probability of detection*, *intrusion*, and/or *discovery* and *compromise of identity*.

Invisible links

There is an “invisible” web of latent bonds in the TN, that represents the set of agents that are all mutually linked to one another by the fact that each member “knows” one or more of the others from some past affiliation, but chooses—at the current time *t*—*not to communicate with them*. By constituting secret pathways for information to flow, these invisible webs provide a vital source of hidden strength to a TN.⁶³ The web remains virtual until an agent “activates” one or more latent bonds to achieve his local goals.

One of the tunable parameters that SOTCAC provides the analyst with is the density of virtual links that lies at the core of the TN; with it, the analyst may explore questions such as, “*To what extent does a TN’s mission success depend on a prescribed degree of virtual kinship connections?*”

Communication

Communication links are links between T-agents that involve an exchange of *information*.

63. In the context of the 9/11 terror attacks, Krebs [193] emphasizes the fundamental role the terrorists’ hidden web of trusted prior contacts played in their mission planning; see discussion in **Appendix 2: mapping Al-Qaeda**.

Each form of communication is characterized by a set of features that determine how the link is used by CT and CT-agents:

- *Type*
- *Content*
- *Directionality*
- *Strength*
- *Duration*
- *Visibility*
- *Vulnerability*

Type

- *Face-to-face.* Face-to-face meetings entail a high degree of risk (and therefore may not be agreed to by an agent that is especially risk averse), and assumes that a threshold level of trust exists between the agents that agree to meet. Face-to-face meetings provide a high degree of certainty of correctly communicating “intended” information, with a relatively low risk of compromise (by the CTN’s virtual INTEL sensors).

The only way in which face-to-face meetings may be “discovered” by the counterterrorist organization is by direct visual contact: one or more CT-agents, possessing a threshold level of experience and ability, must be within a *FaceToFaceDetection-Range* of at least one of the T-agents participating in the face-to-face meeting, and must correctly recognize at least one of the T-agents as belonging to the TN.

- *Indirect.* Indirect communications use an intermediary agent—who acts as courier—to transfer information (or resources) from one agent (typically the leader of the organization or cell leader) to another (typically an operative or support agent). Such so-called *cutout communications* [197] generally involve at least one break in an otherwise direct line of communications, and maximize security and protection from discovery, but at the expense of an increased likelihood of error. Moreover, in the event that a T-agent is “captured” by the CTN and his local topology map is compromised, the identity of senior operatives

(and cell leaders) with whom the T-agent had only indirect lines of communications with, will remain unknown.

- *Email/Chat room.* Essentially another form of indirect contact, email and chat rooms provide the terrorist an electronic medium in which to establish communication links, with the added security benefit of encryption.
- *Phone.* Linking via a conventional (or satellite) telephone is a moderately secure way of communicating. The level of security afforded a terrorist by a particular channel in the real world varies greatly, of course (compare, for example, throwaway global phone cards with cellular systems and long-term residential numbers). SOTCAC makes no effort to explicitly model the level of security (or reliability); the analyst is free to define whatever notional communication channels are desired using primitives such as *strength*, *visibility* and *vulnerability*.
- *Regular mail.* Mail can be used to send messages that minimize the probability that the identity of either sender or recipient will be compromised. In the real-world, regular mail may be used to deliver text, computer disks, memory chips, and SIM cards, all of which may also be encrypted. As for all of the above forms of “communication” it is up to the analyst to decide how best to use SOTCAC’s available primitive features to define an appropriate notional link.

Content

- *Query*
 - *Targeted* to a specific agent (that the sending agent “knows”):
 - What CT-agents do you see?
 - Who do you know? (i.e. send me your entire ego-space)
 - Do you have $\mathbf{X}=\{\text{mission requirement class, amount}\}$?
 - Send $\mathbf{I}=\{\text{communication ID}\}$ to $\mathbf{X}=\{\text{agent ID}\}$
 - *Untargeted* message (communicated to all agents in the sending agent’s ego-space):

- I have $\mathbf{X}=\{\text{mission requirement class, amount}\}$
- I need $\mathbf{X}=\{\text{mission requirement class, amount}\}$

- **Information**

- I know $\mathbf{X}=\{\text{agent ID}\}$
- I know $\mathbf{X}=\{\text{ego-map}\}$
- I am $\mathbf{X}=\{\text{agent ID}\}$: identify self, but do not provide additional information
- I have $\mathbf{X}=\{\text{source ID, amount}\}$: = without identifying self, provide information regarding *resource ID*.
- I have $\mathbf{X}=\{\text{agent ID, resource ID, amount}\}$: = identify self, and provide additional information regarding *resource ID*.

Vulnerability

Each communication link is vulnerable to various modes of “attack” by the CTN:

How vulnerable a given link is depends on its type, information content, the context of the exchange, and (user-defined probabilities of) the particular kind of CTN-directed attack.

For example, an email exchange between two terrorist agents has associated probabilities of “intercept” ($=P_{INT}$) and “location ID” ($=P_{LOC-ID}$): P_{INT} is the probability that the CTN intercepts the email and extracts the information being exchanged; P_{LOC-ID} is the probability that the CTN establishes the physical location of the two agents (and thereby also, implicitly, the probability that the CTN issues an order to send the nearest CT-agents to move toward the terrorist agents in the hope of gaining further intelligence). A face-to-face meeting between two terrorist agents likewise has an associated probability of intercept, but depends on the CTN having at least one *CTa* positioned within an “eavesdropping” distance of one of the terrorist agents.

Adaptive topology

The TN's dynamic topology, at a given instant in time, consists of the set of (1) *active agents*, (2) the set of all *covert* (i.e., invisible, “expired,” or virtual) links to other agents, and (3) all *extant* (and overtly visible) links. The composition of these sets of objects—*agents* and *links*—obviously changes over time as the TN coevolves with the CTN; or, more precisely, as the TN coevolves with the CTN's *beliefs* regarding the TN. We have already discussed some of the rules by which agents may be added to the topology (for example, by being recruited by a T-agent) and deleted (say, by being “captured” by a CT-agent) from the TN;⁶⁴ in this section, we discuss the agent-directed dynamics of topological change and adaptation.

SOTCAC's adaptive topology is loosely based on three sets of heuristic guidelines:

1. The general form of link-*creation* and link-*deletion* rules introduced in the SDCA model (discussed on pages 76-81);
2. The penalty-function-based movement adjudication rules and *meta*-rules used by EINSTEIN; and
3. The “lessons learned” from social network theory regarding the general dynamics governing human communication and information sharing.

SDCA

Recall, from our earlier discussion of the SDCA model, that the essential ingredient of that model is its CA-inspired link-*creation* and link-*deletion* rules. Although they are applied to the topology of a system rather than to the set of values residing at the nodes of a fixed topology (as done by conventional CA rules), the SDCA link rules are manifestly CA-like in that they are *local* within the changing topology; i.e., a given node has access only to information in its local neighborhood and is allowed to change only its local topology.

64. A few more agent-*addition* and agent-*deletion* rules are introduced in the **Counterterrorist network** section below.

The basic mechanism behind SOTCAC’s adaptive topology (which is explicitly defined in a later section; see **SOTCAC’s link rules**) is also manifestly CA-like and borrows heavily from the SDCA model. Links between agents—which are abstract representations of different kinds of “flows” (for example, information, money, and other material resources)—may be created and deleted only *locally*; and agents are only allowed to alter the topology of their local neighborhood, using only local information.

The chief difference between SDCA’s and SOTCAC’s link rules is that while the SDCA model treats links as independent objects that *turn themselves* on or off, in SOTCAC, links are treated as passive conduits that must be turned on or off by *agents*. Moreover, in SOTCAC, links may be directional and require that two agents (i.e., the “sender” and the “receiver”) simultaneously agree to an exchange (subject to their own, typically differing, personalities and motivations) before a link between them is established.

EINSTEIN’s rules

*Brief review of EINSTEIN’s action selection logic*⁶⁵


In EINSTEIN, the global state of the combat model at (discrete) time t , $\Sigma(t)$, is a formal “snapshot” of the system at t that records the identity and locations of all objects, agents, and their internal states. Of course, individual agents typically have access only to some subset of the information contained in $\Sigma(t)$.

Define the *local state*, as perceived by an agent $\mathcal{Q}$, $\Sigma_{\mathcal{Q}}(t) \subseteq \Sigma(t)$, to be the set of all features of $\mathcal{Q}$ ’s local environment that are filtered by $\mathcal{Q}$; i.e., it is the set of features that are either sensed directly by, or communicated to, $\mathcal{Q}$. The two fundamental axioms of EINSTEIN’s action selection logic are then:

1. *All agent actions derive from time varying assessments of the relative value among features $f \in \Sigma_{\mathcal{Q}}(t)$.*
2. *The local state is defined by the matrix of $\mathcal{Q}$ ’s penalty function values: $(Z_{\mathcal{Q}})_{ij}$, evaluated for all sites within a movement range, r_M , of $\mathcal{Q}$ (where $i = x_{\mathcal{Q}} - r_M, \dots, x_{\mathcal{Q}} + r_M$, $j = y_{\mathcal{Q}} - r_M, \dots, y_{\mathcal{Q}} + r_M$, and $\mathcal{Q}$ is at the site $\vec{x}_{\mathcal{Q}} = (x_{\mathcal{Q}}, y_{\mathcal{Q}})$; see figure 40).*

65. The material in this section is extracted and distilled from [11] and [16].

Figure 40. Schematic illustration of EINSTEIN's action selection



Assessments are functions of personality, and consist, in part, of making distinctions between what available features are, and are *not*, relevant to $\mathcal{Q}$ for selecting an appropriate set of actions in a given context. While one set of features might be important to consider in one context, a different set of features might be important for another. No agent can credibly mimic being “intelligent” unless it is able to tailor its actions to specific needs, and adapt to changing contexts. Agents must therefore have some way of identifying which features are most important in a given context, and to deduce which actions are appropriate for the given features.

In EINSTEIN, the user specifies the features that are visible to each agent, and agents select their action, via the penalty function, $Z_{\mathcal{Q}}$, by mapping a given context (i.e., a given vector of feature values visible to it) to motivations for moving toward (or away from) nearby sites. As the value of $Z_{\mathcal{Q}}$, at a given site, $\bar{x}$, increases, $\mathcal{Q}$'s desire to move to $\bar{x}$ decreases, relative to other sites. As the value of $Z_{\mathcal{Q}}$ decreases, $\mathcal{Q}$'s desire to move increases. The site at which $Z_{\mathcal{Q}}$ attains its minimum value ($\bar{x}'_{\mathcal{Q}}$) is the site at which $\mathcal{Q}$ expects to best satisfy its (personality-weight specified) objectives.⁶⁶

General classes of motivations in EINSTEIN include:

66. The fact that $\mathcal{Q}$ seeks to *minimize*, rather than maximize, the value of its penalty function is an artifact of the author's training as a physicist. In physics, one typically solves for the *minimal energy states* of a system.

- *Moving toward, or away from, other agents.*
- *Moving toward, or away from, specific sites or areas on the battlefield.*
- *Minimizing, or maximizing, various local battlefield characteristics (i.e., static indices of the cost of moving over terrain), vulnerability (to possible enemy fire), and visibility.* Static indices represent default, and unchanging, measures that are calculated once prior to the start of a run and in the absence of agents.
- *Minimizing, or maximizing, various local, dynamic combat characteristics (such as the projected vulnerability and visibility if moving to site (x,y) as a function of the actual disposition of local forces, relative local firepower concentration, combat intensity, territorial possession,...).*

Fundamentally, Z_q 's *form* implicitly embodies the filters by which q "sees" the world, and the matrix of Z_q 's *values*, evaluated over the space of all possible moves, explicitly determines how q "reacts" to its world.

From q 's point of view, all behavior ultimately reduces to a calculus of *feature values*: q 's contexts are defined by features (i.e., by q 's perception of the local state), q identifies the features that are important in its current context (which vary according to an agent's "personality"), and q selects an action that essentially represents q 's "best guess" (as defined by Z_q) as to which move leads to a local state in which the values of the perceived features come closest to what q "wants" them to be.

If q were the only agent occupying the battlefield, and the environment was unchanging, q would quickly find the one site (or sites) that best satisfies its needs and stop there. What makes the model interesting, of course, is the presence of multiple agents, all mutually interacting within a changing environment. Each agent's landscape of penalty function values is thus continuously deformed by the actions of other agents. Just as one agent moves closer to "solving" its local problem, other agents move farther away from solving theirs, and all agents face the specter of needing to tune their solutions to constantly shifting problems.

Although the number of terms in $Z_q(x,y)$ is, in practice, quite large, $Z_q(x,y)$ always has the same general form:⁶⁷

$$Z_{\mathcal{Q}}(x,y) = \sum_A w_{\mathcal{Q}}(A) \cdot \mu(A;x,y), \quad (60)$$

where:

- $-1 \leq w_{\mathcal{Q}}(A) \leq +1$ is a numerical weight value that represents $\mathcal{Q}$'s motivation for maximizing (if $w > 0$) or minimizing (if $w < 0$) its expected gain from performing the action A . The label " $\mathcal{Q}$ " appears as a subscript on the motivation w to remind the reader that motivations, which are an integral part of an agent's personality, may be uniquely assigned to individual agents.
- $\mu(A;x,y)$ represents a measure of how well $\mathcal{Q}$ expects it will perform action A in the event that it chooses to move to site (x,y) . The lower and upper bounds of $\mu(A;x,y)$ both depend on A . In simple cases, such as when $A = \text{"move toward squad-mate,"}$ m is equal to the distance between a squad mate and a candidate site to which $\mathcal{Q}$ may move. For other actions, such as $A = \text{"maximize coverage of assigned patrol area,"}$ m is a more complicated function of two or more features.

Because much of EINSTEIN's behavior, both local and global, depends on how the user defines the penalty function, it is important to understand the subtle conceptual difference between *weights*, w , and *measures*, μ , to which the weights are assigned. Weights represent an agent's motivation to perform a given action, and generally depend on one or more dynamic features of the environment according to functions defined by the user. Only their relative values are meaningful. Informally, we may say that w specifies how strongly $\mathcal{Q}$ *either wants, or does not want, to do something* (relative to the set of actions it can perform in a given context).

Measures are also user-defined functions of environmental features (though the features do not have to be exactly the same set as used to define weights), but define how well an agent *expects to perform the action associated with its corresponding weight*. Informally, μ measures

67. Note that the underlying action-selection methodology that is encoded in the expression for $Z_{\mathcal{Q}}(x,y)$ is formally equivalent to the *von Neumann-Morgenstern utility function* (with risk aversion) used in economic decision theory; see [200].

how well $\mathcal{Q}$ expects to do, assuming that $\mathcal{Q}$ has chosen its course of action (consistent with its weights).

Positive weight values are interpreted to mean that an agent is motivated to perform the associated action. Negative values are interpreted to mean that an agent is motivated to *not* perform the associated action (or, more precisely, to perform whatever set of actions are necessary so that the measure associated with performing action A , $\mu(A;x,y)$, is minimized). If the value of a weight is equal to zero, then the agent effectively *ignores* the action (or actions) that are associated with that weight (and is thus also “blind” to the features that the weight is a function of, for the range on which the weight is zero).

EINStein’s rules adapted to SOTCAC

In the spatial domain, EINStein’s movement rules translate almost directly into a form appropriate for SOTCAC’s T-agents and CT-agents. Although SOTCAC’s agents are obviously motivated by features that are either irrelevant or entirely absent in EINStein—for example, although an agent’s “health” is a critical parameter in EINStein, neither T-agents nor CT-agents “care” about this particular feature; on the other hand, SOTCAC’s agents respond to perceived “peer pressure” and “social rank,” neither of which appears as a feature in EINStein—the way in which SOTCAC’s agents use local environmental features to determine their moves is identical to how EINStein’s agents process their moves.

As a simple example, suppose we have a single action,

$$A = \text{minimize distance between } \mathcal{Q} \text{ and agents friendly to } \mathcal{Q}$$

Then, $w(A) > 0$ means that $\mathcal{Q}$ wants to “get closer to” all friendly agents; $w(A) < 0$ means that $\mathcal{Q}$ wants to “get farther away from” all friendly agents. In this case, the measure $\mu(A;x,y) = \text{distance between } \mathcal{Q} \text{ and agents friendly to } \mathcal{Q}$. A point worth emphasizing here, as it becomes an important focus of discussion in the next section, is that $w(A)$ is generally *not a fixed value* (as it was in earlier versions of EINStein); instead, $w(A)$ takes on a range of values (which is now a higher-level

“signature” of an agent’s overall personality), and is a function of one or more environmental features, as sensed locally by $\mathcal{A}$.

In general, while $w(A)$ always represents the motivation to perform a single action (and is usually associated with a single feature), the value of $w(A)$ usually depends on several features.

The battlefield site to which $\mathcal{A}$ moves, $(\hat{x}, \hat{y})$, is given by

$$(\hat{x}, \hat{y}) = \text{Position } (x, y) \text{ such that } Z_{\mathcal{A}} \text{ is minimum: } \underset{D[(x_{\mathcal{A}}, y_{\mathcal{A}}), (x, y)] \leq r_m}{\text{Minimum}} \{Z_{\mathcal{A}}(x, y)\}, \quad (61)$$

where the search for the minimum value of $Z_{\mathcal{A}}(x, y)$ is conducted over all positions (x, y) whose distance from $(x_{\mathcal{A}}, y_{\mathcal{A}})$, $D[(x_{\mathcal{A}}, y_{\mathcal{A}}), (x, y)]$, is less than or equal to $\mathcal{A}$ ’s movement range, r_m .

Social network interaction rules

A variety of mechanisms of social bond formation and interaction have been proposed in the literature; for example [201]-[203]:

- *Homophily*: agents tend to communicate more (or less) with other agents who are more (or less) “like” them.

Homophily refers to using *similarity* (between prospective communication partners) as an underlying mechanism that determines the likelihood of two people establishing, maintaining, altering, destroying, and reconstituting social networks [42]. Homophily has been studied by Byrne [204] and Turner [205].

Byrne was an early proponent of the similarity-attraction hypothesis, that asserts that homophily reduces the psychological stress that otherwise would build up as a result of social differences or inconsistencies. Turner introduced the idea of “self categorization,” which asserts that communication links arise as individuals attempt to define (and refine) their social identity by examining various categories (such as age, gender, and race) within which they can compare their own standing to that of others; in this context, homophily is a mechanism by which an agent legitimizes its own social identity by choosing to associate with others it perceives as falling within the same social categories.⁶⁸

Letting $0 \leq L_{ij} \leq 1$ represent the strength of the communication link between agents i and j , and $f_k(x)$ the value of the k^{th} feature of agent x , the homophily link mechanism can be expressed, formally, as [203]:

$$L_{ij} \propto \Psi \left[\sum_k w_k \cdot |f_k(j) - f_k(i)| \right], \quad (62)$$

where w_k is the weight for the k^{th} feature, $|x|$ is the absolute value of x , and Ψ is some appropriate function. In words, the strength of a bond between i and j is a function of the weighed differences between their attributes.

- *Proximity*: agents tend to communicate with other agents that are “close” to them.

In SOTCAC, the proximity coupling mechanism is followed trivially in social space, since all links are forged and deleted only locally. On the other hand, in the physical space, T-agents typically wish to minimize the risk of discovery, and thus tend behave antithetically to this premise: i.e., T-agents are motivated to *not* directly communicate with other nearby T-agents.

For example, if the physical distance between i and j is D_{ij} , then one obvious form is given by (up to a maximum distance $\hat{D}$):

$$L_{ij} \propto \left(D_{ij} / \hat{D} \right)^n, \quad (63)$$

where $n > 0$. Thus the strength of a communication link (in SOTCAC) *decreases* with decreasing physical distance between agents.

68. An additional argument in favor of using homophily as a TN link-creation mechanism in SOTCAC is that it automatically provides one of the two ingredients necessary for generating scale-free random graph models: namely, *preferential link attachment* [206] (the other ingredient being *growth*, which is also trivially satisfied by explicitly adding new recruits into the TN).

The general role that proximity mechanisms play in social network formation is discussed by Ibarra, *et al.* [207], Monge, *et al.* [208], and Rice and Aydin [209].

- *Social exchange*: agents tend to communicate with other agents that have previously communicated with them.
- *Resource dependency*: agents are likely to communicate with other agents that possess resources that they need and/or that need resources that they can provide.

For example, if $r_{a,k}(i)$ is the k^{th} resource that agent i has *available*, and $r_{n,k}(j)$ is the k^{th} resource that agent j *needs*, we can write, formally:

$$L_{ij} \propto \Psi' \left[\sum_k w'_k \cdot |r_{a,k}(i) - r_{n,k}(j)| \right] \cdot C_B(i), \quad (64)$$

where w' is the weight for the k^{th} resource, Ψ' is some appropriate function, and $C_B(i)$ is the *betweenness centrality* of agent i .⁶⁹ In words, the strength of a bond between i and j is a function of the cumulative, weighed potential for an exchange of resources between them; $C_B(i)$ is added to account for the possibility that there exist multiple agents ($j_1, j_2, \dots, j_n$), all of whom provide i the same resource: in this case, agent i tends to forge only those links that minimizes its dependence on other agents, or, equivalently, maximizes its local “importance” within the network, and hence its *centrality*. (Specific resources in SOTCAC include raw information, money, skills, and weapons.)

The role that resource dependency plays in network formation and social systems is discussed by Bienenstock and Bonacich [210] and Cook and Yamagishi [211].

- *Reciprocity*: a given agent X tends to interact with another agent Y (by either forming or severing bond with Y) according to how X perceives being treated by Y in the past.

69. Recall that the *betweenness centrality* of a node is a metric that measures how “important” a role that node plays in the information flow throughout the entire graph; it is defined in equation 31.

That is, agents tend to interact with other agents in a *tit-for-tat* fashion; they reciprocate positive exchanges with positive exchanges, and negative exchanges (for example, a doublecross) with negative exchanges.⁷⁰

- *Self-interest*: agents tend to simultaneously minimize the total “cost” of communication and maximize individual reward.
- *Evolutionary*: agents tend to locally and/or globally optimize their (own ego-centered) assessment of network “fitness”.
- *Valuation*: agents tend to maximize the collective “value” (or resources) gained from communicating with other agents (*value* mechanism; see below).

As manifest assertions of “bounded rationality,” all three of the preceding mechanisms—*self-interest*, *evolutionary* and *valuation*—underlie most, of not all, multiagent-based simulations. Self-interest, as used here, is essentially a strong form of the *rationality principle*, the application of which to the study of human systems was pioneered by Homans [215]: Homans likened an individual’s action selection to a local cost-benefit analyses in which various options are weighed, and the one that yields the greatest (or optimal) payoff is selected.

Simon [216] later modified this principle by introducing what he called *bounded* rationality: “*Boundedly rational agents experience limits in formulating and solving complex problems and in processing (receiving, storing, retrieving, transmitting) information.*” Simon’s weaker form of ratio-

70. The often counterintuitive ubiquity of *tit-for-tat* strategies in social network settings has been well studied by Axelrod in the context of the *prisoner’s dilemma* [212]–[214]. The prisoner’s dilemma is a game between two people who must choose whether to cooperate with one another. The payoff is such that the best choice for each player, individually, is noncooperation; however, if both players choose to not cooperate the reward is less than if they choose to both cooperate. Thus, while noncooperation is the optimal individual choice, the best mutual decision is to cooperate. After analyzing the efficacy of alternative strategies for repeated rounds of the prisoner’s dilemma game (in both real and simulated settings), Axelrod found that the best strategy is to always cooperate on the first move and then do whatever the other player did on the preceding move.

nality thus posits that agents seek to merely satisfy, rather than exactly solve, their needs.

EINSTEIN's weight-based agent-action selection logic (see equation 55) is already patently bounded rational; as are all of SOTCAC's agent's decisions in the physical space. However, one heretofore missing ingredient necessary to describe agent action selection in the social domain is the relationship between an agent's (social networking) properties and the character of communication links that the agent tends to form.

The theory of *social capital* [217] suggest that an agent i 's *features* ($=f_k(i)$), such as its ability to find information, skill, ability to adapt to increased workload, *etc.*), are influenced by the "social capital" that i is able to accrue from the network in which they are embedded. Social capital represents any attribute, derived from the communication flows and relationships in the network, that is of (either indirect or direct) benefit to a given agent. For example, i 's "value" to a network ($v = f_k(i)$) is, formally, some function of the sum of the strength of links that i has with other agents:

$$v = \Psi'' \left[\sum_k L_{ik} \right], \quad (65)$$

where Ψ'' is some appropriate function, and L_{ik} measures the strength of the link between agents i and k . Of course, this is only the simplest such formal relationship. In general, each of an agent's many features depends on multiple network characteristics and measures of accessible social capital. The exact form of the functional relationship between features and social capital determines the scope and details of a particular model.

For example, Monge and Contractor [203] suggest three theoretical mechanisms of generalizing the social-capital-based self-interest relation embodied in equation 60: *diversity*, *embeddedness* and *holes*.

Diversity. An agent's features may be influenced not just by the strength of ties an agent has with other agents, but also by the diversity of features that an agent's linked partners possess, as a group. For-

mally, i 's k^{th} feature, $f_k(i)$, is a function of the product of i 's link strengths, L_{ik} , and the *variances* of selected features ($=s_s(j)$) of agents to which i is linked [203]:

$$f_k(i) = \Psi'' \left[\sum_k L_{ik} \sigma_s(j) \right]. \quad (66)$$

Embeddedness. An agent's behaviors can be influenced (positively or negatively, depending a given agent's personality) by the degree to which an agent's local ties are reciprocal; that is, according to how strongly an agent feels it is embedded within a local network of mutual links. In the case where network links are directional (that is, L_{ik} is not necessarily equal to L_{ki} ; for example, if information or resources flow from i to k , but not *vice versa*), we can express this formally, by writing:

$$f_k(i) = \Psi''' \left[\sum_k L_{ik} \cdot L_{ki} \right]. \quad (67)$$

Monge and Contractor [203] suggest that the form of Ψ''' may assume both linear-positive forms (increasing embeddedness yielding enhanced features) and nonlinear-negative (increasing embeddedness yielding diminishing returns and/or, after some threshold level is reached, resulting in decreasing influence).

Structural holes. Recall that *structural holes*⁷¹ are the implicit boundaries that separate groups of nonredundant (i.e., disconnected) nodes; they are a network's local buffers between nonredundancy, waiting to be "discovered" and exploited by observant agents. In the current context, structural holes thus represent a potentially rich source of social capital for agents. Agents can strengthen their value to a network by exploiting existing holes and/or actively forging links that spawn new ones (with them at the center). Formally:

71. The "structural hole" concept was introduced by Burt [134]; see discussion in section **Complex networks: metrics**, page 102.

$$f_k(i) = \Phi \left[\sum_{s|L_{is}, L_{js} > 0} \sum_{j|L_{ij} > 0} L_{ij} \cdot L_{is} \cdot L_{js} \right], \quad (68)$$

where Φ is some appropriate function (that typically assumes larger values for smaller values of its argument: smaller values of the triadic relationship expressed by the summand are expected to yield a stronger brokerage potential), and the double summation is over all nodes j that are linked to i and nodes s that are linked to both i and j . L_{ik} measures the strength of the link between agents i and k .

From the standpoint of *self-interest*—as a core motivation underlying agent action—it is assumed that agents always seek to maximize access to exploitable social capital. For example, in addition to being influenced by the strength of existing holes, agents can also catalyze the formation of new holes by simultaneously forging links with unconnected agents and finding ways to minimize the interaction among those (otherwise loosely connected) groups.

- *Altruism*: agents tend to maximize the collective value of the communication to the group to which they belong.
- *Group cohesion*: agents tend to interact with other agents that belong to the same social group.

Group cohesion manifests itself as a dynamic element not just as a propensity for agents to link with group-mates (such as, say, with other members of a terrorist cell in SOTCAC), but also as the basis for strengthening (or weakening) links. For example, equating group cohesion with *group density* ($=r_G$), defined as the average link strength in G ($=\langle L \rangle_G = |G|^{-1} \sum_{i,j \in G} L_{ij}$), we can write, formally [202]:

$$L_{ij} \propto (\rho_G)_{ij} - \langle \rho_G \rangle, \quad (69)$$

where $(\rho_G)_{ij}$ is the density of the group to which i and j both belong, and $\langle \rho_G \rangle$ is the average density of the groups in the network.

The study of group cohesion, and its role in social network dynamics, arguably lies at the core of social network analysis, and dates back to the work of Back [218], Homans [219] and Seashore [220] in the 1950s. More recent studies are by Evans and Dion [221] and Moody and White [156]; Friedkin [222] provides a short review.

SOTCAC's link rules

SOTCAC's *link-creation* and *link-deletion* rules are, essentially, the set of social network theoretic communication heuristics applied to EINSTEIN's personality-weight-prescribed movement rules. The main difference between SOTCAC's rules and EINSTEIN's rules, is that while EINSTEIN's agents are confined exclusively to the physical domain, SOTCAC's agents maneuver in both the physical and information domains. However, although SOTCAC's link rules apply directly only to the abstract graph space that contains the TN's evolving topology, the physical space remains an indirect participant by providing a backdrop of physical features that T-agents can use for adapting their default link rule characteristics.⁷²

Figure 41 shows a schematic illustration of a portion of the TN, as seen from a T-agent s 's point of view. The red nodes (that appear within the region shaded in red) represent the T-agents that are within a distance $D = 2$ of s .

In general (at the user's discretion), s 's local r -neighborhood, $S_r(s)$, consists of all agents that are within a distance $D = r$ of s :


$$S_r(\sigma) = \{i \mid \text{Dist}(i, \sigma) \leq r\}. \quad (70)$$

Just as EINSTEIN's agents can "see" (and react to environmental features located) only as far as their (user defined) sensor range permits, SOTCAC's agents "know" (and react to agents and other information located in *graph space*) within their local r -neighborhood.⁷³ It is within $S_r(s)$ that all of the adaptive "rewiring" of the TN's topology takes place. Agents can choose to *sever existing links*, *create new links* between themselves and other agents with whom they are not currently linked, or to "*do nothing*" at the current time.

72. For example, the "*discovery risk aversion*" rule (discussed below) is activated, partly, as a function of the relative physical positioning of nearby T-agents and CT-agents.

73. Conceptually speaking, there is little difference between EINSTEIN's *physical* space and SOTCAC's *graph* space. While the physical space is notionally a battlefield (and is obviously not a social network), it is represented mathematically, within EINSTEIN, as a fixed N-by-N lattice of sites; i.e., it is a *graph*. Thus, SOTCAC's general local r -neighborhood is a direct analog of EINSTEIN's local ($\text{range} = r$) sensor field.

Figure 41. Schematic illustration of an agent s 's *local neighborhood* in social network (i.e., *graph*) space



Motivations

T-agents base all of their actions (in graph space) on three general classes of motivations: (1) *mission-centric*, which derive from the requirements that must be met by each cell (and, hence, each cell member) before a terrorist mission can be launched (for example, an agent may be motivated to find and acquire a specific quantity of a certain resource); (2) *topological*, which are those that derive from measures of local structure (for example, an agent may desire to maintain a minimum level of connectivity with his cell leader); and (3) *functional*, which includes social network metrics for communication flow (for example, an agent may be motivated to trade one kind of link for another to maximize the efficiency of exchange).

Examples of specific motivations for creating (and/or maintaining) links include:

- *Acquiring mission-required resources.* Recall that all active T-agents are members of a TN cell (under the command of a cell leader), and that each cell is assigned a physical "target" that requires a specific set of mission requirements be met before a strike against it can be launched. These requirements take the form of acquiring sufficient manpower, financing, weapons

and skills. While it is a leader's responsibility to maintain a threshold number of mission operatives (which requires that links be established with recruiters, recruits, and/or trainers), mission operatives must find and acquire necessary resources.

- *Discovery risk aversion.* Agents are generally averse to being “discovered” by nearby CT-agents. Prospective link partners (in graph space) are therefore afforded lesser (or greater) weight depending on an agent's estimate of the likelihood that a link may be “seen” by one or more counterterrorist agents.
- *Cell-cell mix proclivity.* Regulates the degree to which an agent does (or does not) want to associate with agents from other cells. Depending on the type of link (for example, information versus exchange and/or transfer of material resources), agents may be more (or less) inclined to maintain strict compartmentalization.
- *Coordination.* Cell leaders must periodically communicate with cell members, ostensibly to coordinate activity, but also to maintain cell cohesion. The user controls the coordination frequency; which can also be a function of the rank and skill of operatives (less skilled and lower ranking operatives requiring more frequent “pinging”).
- *Maximize structural autonomy.* As an example of a functional motivation, agents may wish to maximize their “structural autonomy” by actively seeking to establish ties with other unconnected agents; i.e., to create links indirect ties between otherwise unlinked agents, with themselves as mediators. Recall that structural “holes” (Burt [134]; see page 102) represent a rich source of social capital that agents can exploit both directly (by using their central position to maximize access to information not accessible by others, and thus streamlining their acquisition of resources), and indirectly (by forging links that create holes that may be exploited in the future). Agents can maximize their entrepreneurial networking opportunities by fashioning their local neighborhoods to provide multiple structural holes around their neighbors, but none around themselves.
- *Maximize familiarity.* Agents tend to forge links with agents with whom they are already familiar. This set consists of not just those agents that belong to the same cell (some of whom a


given agent may actually *mistrust*), but other agents that a given agent has had prior contact with (see *trusted priors*, page 154).

- *Other topologically endogenous motivations.* Many more topologically endogenous motivations are possible: *degree, diversity, betweenness, closeness, centrality, efficiency*, etc.

Motivations for deleting existing links include:

- *Lifetime of particular link-type has been reached (ex: duration for financial transfer) and/or required resources have been acquired*
- *Perceived risk of discovery exceeds either receiving or transmitting agents' risk tolerance.*
- *An agent's maximal processing workload has been exceeded.*
- *Linked agent is no longer perceived as sufficiently trust-worthy.*

Figure 42. Schematic of link-creation and link-deletion in SOTCAC; see text for details



Constraints

The analog to EINSTEIN's *meta*-rules, which act as local constraints on agent behaviors, are topological constraints that can likewise be used to regulate SOTCAC's social network rewiring:

- *Maximum number of extant links.* Specifies the maximum number of currently “open” communication channels that a given agent is allowed to have. (Can be used a rough measure of “information overload” and correlated with an agent's skill and experience.)
- *Maximum resource flow.* Limits the total amount, $r(R)$, of resource R that an agent can transfer at one time.

- *Minimum distance from agent X.* Specifies the minimum distance that a given agents wants to be from another agent of type X; for example, “X” may be cell X, in which case, the given agents desires to be a minimum distance from all agents belonging to cell X. (The relative importance of this particular constraint is obviously commensurate with the size of an agent’s local r -neighborhood.)
- *Disconnect Intolerance.* Defines the degree to which an agent does not want to be disconnected from other T-agents or cell mates; i.e., the minimum number of extant links that an agent seeks to possess at all times. If less than a threshold number of links are active, an agent will unilaterally create one or more links with (randomly selected) members of his own cell.

Counterterrorist network

“Counterterrorism... Action or strategy intended to counteract or suppress terrorism.”—American Heritage Dictionary, 4th Edition

The *counterterrorist network* (CTN) is the dynamical, coevolutionary complement of the TN. Its central mission is to attack, disrupt and/or destroy as many of the TN’s self-organized activities as possible, in order to prevent the TN from accomplishing its terrorist goals; and to minimize (or eliminate entirely) the TN’s ability to do conduct future missions. For accomplishing these tasks, the CTN is equipped with two kinds of agents: (1) *field agents*, that interact directly with one another and suspected T-agents, and members of the terrorist support organizations; and (2) *virtual agents*, that act as covert INTEL-gathering resources, and do not interact with T-agents in the physical domain.

Functions

The CTN performs three basic functions:

1. *Collects INTEL about the TN’s activity,*
2. *Generates a “real time” map of the TN’s composition and structure,*
and

3. *Performs strike missions (along with other actions) against selected “high value” components of the TN.*

Just as the terrorist network “lives” in two coupled domains—*physical* space and *information* space—so too does the CTN. The physical domain constitutes the dynamical arena in which terrorist and counterterrorist agents move and interact. The CT-agents are also equipped with a special set of INTEL-processing rules, that specify the conditions under which certain kinds of information may be collected about neighboring T-agents.

The information space harbors the results of CTN’s *inference engine*, which collates, fuses and updates all INTEL regarding TN activity. The inference engine generates a visualization of what the TN “looks like” at time t , *from the CTN’s point of view*; the visualization represents an imperfect, “best-guess” snapshot of the TN’s agents, structure, and activity.

Thus, the CTN’s dynamic role in the $TN \leftrightarrow CTN$ coevolution consists of these basic steps: direct the actions of, and collect INTEL regarding the location and actions of nearby terrorist agents from, CT-agents distributed in the physical domain; specify the data-collection requirements for, and direct the actions of, the set of covert assets (or virtual agents/sensors); collect and interpret all information reported by physical and virtual agents; fuse all new INTEL with existing knowledge to infer the current state of the TN; make decisions about where to send physical agents, what part of the TN the virtual sensors ought to focus their attention on, and make strike recommendations based on what the CTN’s inference engine decides are the most “important” components of the TN.

Implicit in these steps, is the fact that essentially all of the CTN’s dynamical functions—such as “deciding” where to send its agents, “interpreting” information, and “recommending” strike plans—are *subjective*. Thus, it is important to provide the analyst with some degree of *objective* control over how the CTN makes decisions; i.e., the analyst must be able to define, and tune, the CTN’s personality for making decisions. The mechanism for this is described below, and follows closely how agent personalities are defined in EINSTEIN [16], and the agent-based variant of SCUDHunt [17].

INTEL Assets

Physical Agents

The CTN's physical agents, $\{CTa_1, CTa_2, \dots, CTa_N\}$, occupy the same physical space as (and interact with) TN agents. Although they share the same basic class of properties that govern the behavior of their TN counterparts (see *CT-agent characteristics*, below), CT-agents base their actions on a different set of motivations.

For example, while T-agents are typically motivated to stay clear of CT-agents (as well as other T-agents that they recognize as members of their own cell), CT-agents are generally impelled to find and/or track T-agents; and of enlisting the aid of other nearby CT-agents when the context warrants.

CT-agent characteristics

Typical CT-agent characteristics include:

- *Movement range*, $CTMoveRange$: as for T-agents, $CTMoveRange$ defines the maximum distance that a CT-agent can move from its current location (at time t) to its new position (at time $t+1$).
- *Sensor range*, $CTSensorRange$: defines the range of a CT-agent's *vision* in the physical space. Agents sit at the center of an $CTSensorRange$ -by- $CTSensorRange$ box, and "see" everything that occupies any of the $(2*CTSensorRange + 1)^2$ sites within this box. (Just as for T-agents, this does not imply that a CT-agent correctly *identifies* what is present at a given site, only that the agent has the *potential* of registering another agent within its sensor range.)
- *Ta-detection probability*, $P_{CTa \leftarrow Ta}$: probability with which a CT-agent is able to correctly "recognize" an otherwise unidentified agent, within its Ta-detection range, as a T-agent.
- *Ta-detection range*, $R_{CTa \leftarrow Ta}$: range (in the physical domain), at which a CT-agent is able to "see" a T-agent. (Whether, or not, the T-agent correctly identifies this agent as a CT-agent depends on the CTa 's Ta-detection probability.)

- *Ta-follow range*, $R_{Ta-follow}$: minimum range that a CT-agent—who is “following” a suspected (or previously “tagged”) T-agent—wants to maintain between himself and the T-agent.

Virtual Resources

The CTN’s virtual resources represent all of the counterterrorist organizations’ non-HUMINT data collection and countermeasures assets: reconnaissance satellites and/or reconnaissance-equipped unmanned air vehicles (UAVs), communications intelligence resources (COMINT), signals intelligence resources (SIGINT), and general electronic intelligence resources (ELINT; which includes jamming and electronic deception). All of these resources are notional, in that their presence is *implicit*; only their dynamical effect on the dynamics is modeled explicitly.

For example, while the CTN’s virtual SIGINT capability ostensibly includes the ability to eavesdrop on the TN’s radio conversations (as well as an attendant code-breaking capability or other cryptological expertise), there are no physical agents that are explicitly tasked with this function. Only the SIGINT’s data-gathering and processing capability is modeled, such as the probability of intercept and probability of deciphering the message.

Similarly, although the CTN’s inference engine is an overarching information agent that “lives” within the CTN’s *belief matrix* (described below), it is entirely virtual: the $i - j$ grid of belief values (that represent the CTN’s confidence in the existence of a link between terrorist-agent/suspect i and terrorist-agent/suspect j), is the abstract, information-space, analogue of the physical $x - y$ lattice of possible T-agent positions. The actions that the CTN decides to take at any time t , using either its physical and/or virtual CT-agents, depend on how the virtual inference-engine-agent processes its matrix of beliefs; and this processing takes place entirely in an abstract space.

Other virtual assets are handled in the same way.


CTN Actions

At any given time t , the CTN takes one or more of the following actions (not all of which are necessarily viable at a given time, and in a given context):

- *Assimilate CT-agent-filtered data.* A basic function performed by all CT-agents is to provide the CTN a “filtered” view of their local environment. How the CTN processes this CT-agent-centric view, and how it fuses it with its own beliefs, is discussed in the sections **CTN beliefs** and **CTN inference personality**; here we outline two alternative (user selectable) methods by which an individual CT-agent processes raw data.

Figures 43 and 44 show, schematically, how a CT-agent—labeled CTa , and positioned at the center of the gray circle (that represents the extent of its $CTSensorRange$)—interprets the presence of eight T-agents, $Ta_1, \dots, Ta_8$; some of whom (i.e., those that are joined by a red link) are engaged in face-to-face meetings. CTa interprets this situation using one of two methods:


Figure 43. Schematic illustration of how the CTN “infers” latent TN-structure using raw data filtered by a CT-agent, using *method 1*; see text for details



- *Method #1.* The first method is probabilistic; it uses the Ta -detection probability, P_{CTa-Ta} , and a user-defined probability of detecting an existing face-to-face link (= $CTProbFaceToFaceLink$) to determine how CTa interprets “ground truth.”

In a three-pass approach, the *CTa* (1) “probes” each agent within its *CTSensorRange* by rolling a virtual “die” to obtain a random number, $0 \leq x \leq 1$. If $x \leq P_{CTa \leftarrow Ta}$, *CTa* correctly identifies the given agent as a T-agent (thus tagging it, and incorporating it into its ego-map), otherwise the given agent remains unidentified; (2) for each *existing face-to-face link* (between those T-agents that have been correctly identified in step 1),⁷⁴ *CTa* throws another virtual die to obtain a random number, $0 \leq y \leq 1$. If $y \leq CTProbFaceToFaceLink$, *CTa* correctly identifies the given link, otherwise the link remains invisible; and, (3) for all pairs of T-agents (that have been correctly identified as such in step 1, but who *are not* actually linked), *CTa* throws a third virtual die to obtain a random number, $0 \leq z \leq 1$. If $z \leq \epsilon_{link}$, where $0 \leq \epsilon_{link} \leq 1$ is the probability of misidentifying a nonexistent link as open, *CTa* correctly identifies the given link, otherwise the link remains invisible.

Figure 44. Schematic illustration of how the CTN “infers” latent TN-structure using raw data filtered by a CT-agent, using *method 2*; see text for details



— *Method #2*. The second method is decision-based; that is, it depends less on probabilistic adjudication and more on a context-driven intelligence (see figure 44).⁷⁵

74. This “ground truth” being known only to SOTCAC, and not either the *CTa* or the CTN.

First assuming that *CTa* has correctly identified each of the eight T-agents as belonging to the TN (there may be other T-agents within *CTa*'s *CTSensorRange* that—just as for method 1—because of *CTa*'s *Ta-detection probability*, have not been detected), the CTN is apprised of these agents' (x,y) positions. Second, whenever the distance between any two T-agents— Ta_i and Ta_j —is less than a threshold (= *CTFaceToFaceMtgRange*, which may be different for different CT-agents), *CTa* infers that they are having a “face-to-face” meeting and reports the existence of a link $l_{ij}=1$ between Ta_i and Ta_j .

The two “face-to-face” renditions appearing on the right-hand-side of figure 44—the first showing what *CTa* reports using a *CTFaceToFaceMtgRange* value of *one* (as marked in red on the left-side of the figure), the second showing what *CTa* reports using a *CTFaceToFaceMtgRange* value of *two*—illustrates that different values of *CTFaceToFaceMtgRange* can significantly alter *CTa*'s inferences based on what it “sees.”

The main difference between these two CT-agent filtering methods is that while method #1 allows the user freedom to effectively control how close to (or far from) “reality” a CT-agent's inference is, on average (but without regard to the logic by which the inference is reached), method #2 allows the user to experiment with alternative inference-logic algorithms. Moreover, while method #1 filters raw data using


75. The method described here is marginally “intelligent,” at best; its importance, at this early juncture, derives less from its ability to *drive* outcomes of runs, and more from serving as a conceptual placeholder for future decision-based algorithms that will allow analysts to explore tradeoffs among different ways of interpreting raw data. Method #2 is anticipated to play an analogous role, in SOTCAC, to EINSTEIN's evolving *targeting logic* (by which combat-agents decide which enemy agents to fight). While early forms of this logic amounted to little more than directing agents to shoot blindly and randomly at their targets, in time, the algorithm evolved to a sophisticated multi-parameter-dependent intelligent targeting logic by which users can finely tailor an agent's “combat personality.” While this targeting logic has become a bona-fide feature of EINSTEIN only in its most recent incarnations, it nonetheless played a vital role in earlier versions of the simulation during which time it served as a conceptual placeholder for the need to “intelligently” evolve an adaptive targeting logic.

ground truth data, method #2 entails no guarantee that *any* value of *CTFaceToFaceMtgRange* generates ground truth. Since the two methods obviously take fundamentally different approaches to filtering raw data, the user must decide issues of suitability in the context of the specific set of questions being addressed by the model.

- *Issue an order to CTa_i to move toward (or away from) either a specified site in the physical space, a specific T-agent or set of T-agents.* (The viability of this action depends on whether the CTN has a requisite number of CT-agents with an *IssueMoveOrderRadius* of its desired target.) This action may also be taken by individual CT-agents, and takes on a particular relevance in the context of capturing a terrorist agent (see below). A CT-agent may also be directed to “follow” a particular T-agent (at a discreet distance; see *Ta-follow range* above).
- *Capture a terrorist agent (Ta_i).* Depending on Ta_i ’s personality (allegiance, experience, etc.), the CTN may extract a greater or lesser amount of information about the TN (as represented by Ta_i ’s ego-map of TN) possessed by Ta_i .

At the user’s discretion, a threshold number of nearby supporting CT-agents may be required (to simultaneously surround a T-agent “tagged” for imminent capture) before a T-agent is captured. Nominally, the *CTAgentCaptureNumber* = 1; if a greater number is required, the first CT-agent (= CT_i) to move within the *CTMinCaptureRange* of the tagged T-agent communicates the sighting to all CT-agents within a *CTCaptureCommRange* of CT_i position, who, assuming they are not in the process of capturing another T-agent somewhere in their own vicinity, create a new T-agent-specific movement weight component and assign that component the maximum possible positive value (= +1); i.e., upon receiving a call for assistance from CT_i , all agents within range move as quickly as possible to surround the tagged T-agent.

Figure 45. Schematic of a typical scenario in which three T-agents (Ta_1, Ta_2, Ta_3) are captured by CTa_1 ; see text for details



In general, if a CT-agent “sees” a total of N_T T-agents within his $CTMaxCaptureRange$, those T-agents may all be captured—as a group—if the total number of CT-agents outnumber the T-agents by an amount $\Delta_{CT\ Capture} (= \text{Number of CT-agents} - N_T)$.⁷⁶ Figure 45 shows a schematic of a typical scenario in which there are five T-agents ($Ta_1, \dots, Ta_5$) within a counterterrorist agent’s ($=CTa_1$) sensor range ($=CTSensorRange$); three of these T-agents (Ta_1, Ta_2, Ta_3) are also within CTa_1 ’s $CTMaxCaptureRange$.

Assuming that $\Delta_{CT\ Capture} = +1$, CTa_1 communicates its need for support to the four CT-agents ($CTa_1, \dots, CTa_4$) within its $CTCaptureCommRange$; each of these counterterrorist agents adjusts its movement vector to allow for maneuver toward CTa_1 . (Notice that while CTa_1 “sees” Ta_4 and Ta_5 , because these two T-agents lie outside CTa_1 ’s $CTMaxCaptureRange$, they cannot be cap-

76. This is an analogue of EINSTEIN’s combat *meta*-rule, which defines the local conditions under which agents choose to engage in combat: *if (Number of nearby Friends) - (Number of nearby Enemies) $\geq \Delta_{Combat}$ then fight, otherwise temporarily disengage* [16].

tured.) The T-agents Ta_1, Ta_2 , and Ta_3 are captured as soon as CTa_2, CTa_3, CTa_4 , and CTa_5 move to within range (i.e., within a distance $CTMaxCaptureRange$ of CTa_1).

- *Insert a counterterrorist agent (CTa_i) into the “slot” occupied by terrorist agent (Ta_j);* in effect, convert an existing Ta_j agent into a CTa_i (i.e., make him a double-agent). The probability for this happening is obviously a function of Ta_i 's allegiance to the TN, as well as on its experience, rank, and value. However, if the “conversion” is successful, the CTN's *belief-map* has access to an “insider's view” of all the information TN normally provides Ta_j up until such time as the TN discovers the insertion and destroys the agent-node.

Recall that there is also a chance of *infiltrating* the TN (as it forms) with a double agent who poses as a possible recruit. If such an agent is successfully absorbed into the TN, its dynamics are from that point onward identical to those regulating the behavior of *inserted* agents.

Once inserted, a double agent faces the specter of discovery: there is a probability, *ProbInsertionDiscovery*, that he will be uncovered as a CTN spy. If discovered, the agent is effectively “killed” and removed from the TN.

- *Eavesdrop on (or wire-tap) a link (that the CTN believes exists) between Ta_i and Ta_j ($=CTL_{ij}$),* and thereby extract the information being passed along that link. (The ramifications of doing so, including the scenario in which the CTN's belief that CTL_{ij} exists is incorrect, are discussed below.)
- *Intrude on a link CTL_{ij} , and either jam the signal* that is being passed (decreasing the probability that the information being sent will be received), *or embed a false signal* (which maintains the integrity of the communication, so that the message is still received, but the components of that message—such as agent ID, location, and orders—are deliberately falsified).
- *Probe a node or link:* essentially a weaker form of “capturing” an agent, the CTN may elect to “probe” a node (for a duration

dt_{probe} ; provided certain user-defined trigger-conditions are met) to obtain information about incoming/outgoing links from Ta_i (i.e., a subset of Ta_i 's ego-map).

- *Destroy a TN node or link*, the probability of the CTN succeeding at which is partly a function of the vulnerability of the targeted node or link, and partly a function of the strength of the cell (or cells) to which they belong. The destruction of a node also entails the destruction of the overall value that node provided to the TN; i.e., its role and rank within the terrorist organization, its knowledge, skills and resources, and all social network links that it served as a hub to.

While the act of destroying a node also destroys the nodes ego-map (which is thus *not subsumed* into the CTN's belief map; in effect, the information previously possessed by the T-agent is destroyed along with the T-agent), in the event that the CTN has targeted a nonexistent target (i.e. one that it has mistakenly believed to exist), the appropriate components of the CTN's belief map are updated accordingly (see discussion below).

The actions that the CTN chooses to take at time t depend on what the CTN *believes* to be the “state of the TN” at that time.


CTN Beliefs

SOTCAC represents the counterterrorist network's activity in information space by visually rendering what the CTN's *believes* the TN composition, structure, and activities are at time t (see figure 46).

These beliefs—not all of which are necessarily correct, as they may be based on outdated, incomplete, and/or false INTEL (see below)—are functions of time, INTEL-source, and the CTN's data-fusion personality; they are embodied in four dynamic objects:

1. *OperativeID*-belief vector, $b_O(t)$: this represents the number, and identification tag (ID), of agents in the physical domain that the CTN believes are members of the TN possesses, at time t .
2. *Composition*-belief matrix, $b_C(t)$: this represents the CTN's “best guess” about what role each of agents it believes are members of the TN plays in the terrorist organization, at time t .

Figure 46. A schematic view of what the counterterrorist network *believes* the TN's structure is at a given moment in time; this view represents the CTN's activity in the information domain; see text for details



3. *Structure*-belief matrix, $b_S(t)$: this represents the CTN's beliefs, at time t , regarding the TN's network topology.
4. *Activity*-belief vector, $b_A(t)$: this represents what the CTN believes the TN is doing.

Collectively, these four quantities—which, henceforth will be generically referred to as the CTN's *belief map*—describe the “state” of the TN, as understood by CTN, consistent with its INTEL-gathering, information-interpretative and decision-making personality (see **CTN inference personality** below).

***OperativeID*-belief vector**

The OperativeID-belief vector represents the number, and identification tag (ID), of agents in the physical domain that the CTN believes are members of the TN, at time t . Until the CTN positively identifies an agent X as a terrorist (modulo the threshold degree of certainty

required by the CTN’s inference engine; see below), X is tagged a “suspect” and may be reconnoitered to update the belief map.

The number of positively IDed agents, $N(t)$, is not constant, but fluctuates as the CTN continuously adjudicates the inflow of CT-agent sightings of TN activity in SOTCAC’s physical domain.

Composition-belief matrix

The CTN’s *composition*-belief matrix at time t , $b_C(t)$, summarizes what the CTN *believes* is the TN’s composition at time t . The components of $b_C(t)$ — $b_C(i,j;t)$ —take on values between -1 and $+1$, and represent the CTN’s confidence that agent Ta_i plays role X_j within the terrorist organization, where $X_j \in \{recruit, recruiter, trainer, operative, leader or support agent\}$. If $b_S(i,j) = -1$, that means that the CTN believes that Ta_i *definitely does not* play the role of X_j ; $b_S(i,j) = +1$ means that the CTN believes that Ta_i *definitely* plays the role of X_j ; $b_S(i,j) = 0$ means that the CTN either currently has no reliable INTEL about agent Ta_i and/or is neutral in its assessment as to what role Ta_i plays. Values of b_C , between these two extremes, represent various confidence levels.

Figure 47. Graphical view of the CTN’s *composition*-belief matrix; different shades of grey at a site (i,X) represent varying degrees of confidence that the CTN has in its belief that agent Ta_i is either a member of the TN and/or is of type X




Figure 47 shows a schematic view of $b_C(t)$, in which shades of grey are used to represent varying degrees of the CTN’s confidence that

CTN's confidence that a given agent plays one of the six generic roles of terrorist agents. As noted above, the number of agents that the CTN believes the TN to possess, at time t ($= N(t)$) is not constant, but varies depending on the dynamics of the CTN's INTEL-collection and data-fusion processes.

Structure-belief matrix


The CTN's *structure-belief* matrix at time t , $b_S(t)$, summarizes what the CTN *believes* is the TN's social network structure at time t . $b_S(t)$ is partly a function of the information the CTN's agents gather while interacting with T-agents in SOTCAC's physical domain, partly a function of other INTEL that the CTN's COMINT (and other virtual) assets collect, and partly a function of the CTN's INTEL-fusion personality.

The components of $b_S(t)$ — $b_S(i,j;t)$ —take on values between -1 and $+1$, and are interpreted to mean the CTN's confidence that a link exists between terrorist agent Ta_i and terrorist agent Ta_j . If $b_S(i,j) = -1$, that means that there is *definitely no link* between Ta_i and Ta_j ; $b_S(i,j) = +1$ means that the CTN is *certain* that a link exists between Ta_i and Ta_j ; $b_S(i,j) = 0$ means that the CTN has no belief (and therefore no reliable INTEL) regarding the existence of a link between Ta_i and Ta_j . Other values of $b_S(i,j)$, between these two extremes, represent various confidence levels.

Figure 48 shows a schematic view of $b_S(t)$, in which shades of grey are used to represent varying degrees of the CTN's confidence that a link (i,j) exists between T-agents i and j ; because links are symmetric, only the top half of the matrix is shown. The *rows* and *columns* label each of the T-agents, $\{Ta_1, Ta_2, \dots, Ta_{N(t)}\}$, that the CTN *believes* to exist at time t .

The belief matrix is the critical component of SOTCAC's representation of the CTN, as the decisions regarding *CTa* placement and movement (within the physical domain), wire-tapping, *Ta* capture, and TN node/link disruption, intrusion and/or destruction, are all functions of $b_{CTN}(t)$.

Figure 48. Graphical view of the CTN's *structure*-belief matrix; different shades of grey at a site (i,j) represent varying degrees of confidence that the CTN has in its belief that a link really exists between T-agents Ta_i and Ta_j ; see text for details



Activity-belief vector

The CTN's activity-belief vector represents what the CTN believes the TN is *doing*; i.e., what mission it has assigned to itself (and to what cells), what the mission requirements are, and how close the TN to completing acquiring the manpower and resources necessary to initiate terrorist strikes.

CTN inference personality

How the CTN's inference engine actually parses the raw intelligence data (regarding the identity and properties of specific TN nodes and links), and the way in which the CTN updates its belief map, is a func-

tion of the CTN's *inference personality* (IP). The IP—which is defined, and fixed throughout a run, by the analyst—consists of various parameters that define how the IE obtains, interprets and uses SOTCAC-generated information:⁷⁷

- *Interpretation of INTEL data*
- *Valuation of INTEL data (as gathered by specific physical CT-agents and virtual collection assets)*
- *Fusion of INTEL reports (belief-map refinement)*
- *Tactical action plan logic*

Interpretation of INTEL data

The first part of the CTN's IP consists of parameters that define how the CTN interprets sensor reports. The basic idea is that each of the TN's nodes and links, for which the CTN either has no prior belief value (such as at the start of a run) or some nonzero value $b_S(i,j;t)$ that has been accrued over several turns, is updated using information supplied by either CT-agents (acquired in the physical domain) or the CTN's virtual assets (via targeted captures, insertions, intrusions, or probes). For example, before the CTN's beliefs about a specific node Ta_i are updated ($=b_C(t)$), the CTN must (1) interpret the data that each of its assets has gathered regarding Ta_i (which depends, in part, on how well CT-agents are able to collect the data, and, in part, on how the CTN "values" the collected data), and (2) fuse the combined intelligence data.

CT-agents provide reports of the following types: (1) *nothing significant to report* (about suspected T-node 'X'), (2) *unidentified* (as to type of) agent 'X' present, (3) *T-agent 'X' of type 'T' present*, (4) *unidentified link 'X' present* (sender or receiver identified as T-agent), and (5) *link of type 'X' exists* between T-agents Ta_i and Ta_j .

The CTN's *ReportInterpretation* vector, $\vec{I}_R$, defines the CTN's belief that object 'X' exists for which an agent (or sensor) has reported R (where R is any of the five possible CT-agent reports defined above); i.e., $\vec{I}_R$ represents the CTN's interpretation of the data collected by its CT-agents. The components of $\vec{I}_R$ range between the values -1 and $+1$.

77. The methodology described here is based on a prototype inference algorithm developed by the author for the agent-based SCUDHunt wargame [17].

For example, the component $\bar{I}_R = (-0.5, +0.5, +0.9, -, -)$ means that the CTN interprets a report from CT-agent A as follows: (i) “*nothing significant to report*” is treated as a negative 0.5 partial-belief that node ‘X’ is a T-agent; (ii) “*unidentified (as to type of) agent ‘X’ present*” is treated as a positive 0.5 partial-belief that ‘X’ is a T-agent; (iii) “*T-agent ‘X’ of type ‘T’ present*” is treated as a positive 0.9 partial-belief that ‘X’ is a T-agent of type ‘T’; and (iv) the last two entries (= “-”) indicate that CT-agent A is unable to gather link-related intelligence.

The utility of using the *ReportInterpretation* vector to filter the raw data before it is assimilated by the CTN, is that it provides the analyst the flexibility to explore the dynamical (i.e., coevolutionary) consequences of relying on different interpretations of the same data. By effectively decoupling the CT-agent INTEL-collection dynamics from the CTN INTEL-data fusion, SOTCAC allows the analyst to focus attention, *separately*, either on how well CT-agents are collecting relevant data (which is a function of CT-agents’ innate properties) or on how well the CTN is assimilating and fusing the data that is collected (which is a function of the CTN’s INTEL-processing personality).

Valuation of INTEL data

The second component of the CTN’s IP is the set of partial-belief assessments of INTEL-data that is reported by *specific* CT-agents (CTa_i). While the *ReportInterpretation* vector defines the CTN’s generic interpretation of INTEL reports, the *AgentValuation* vector, $\vec{V}_A$, refines that interpretation, by tailoring it to specific agents; indexed by CTa_i , the value of its components range from *zero* to *one*, and represent how strongly the CTN “values” CT-agent CTa_i . $\vec{V}_A(CTa_i)=0$ means that the CTN *mistrusts all data* reported by agent CTa_i (and therefore neither adds to nor subtracts from its marginal belief regarding any component of the TN structure and function that is due to information supplied by CTa_i); $\vec{V}_A(CTa_i)=1$ means that the CTN *trusts, implicitly, all* CTa_i reports (and therefore does not alter its generic interpretation of the specific content of CTa_i ’s reports).

Fusion of INTEL reports

Before deciding on a course of action to take at time t , the CTN must first process and fuse N sensor reports (meaning their *interpretations* and *valuations*) to determine its “best guess” as to which agents are

T-agents (along with their “type” and the topology of observed communication links); i.e., the CTN must update its *belief map*. One mathematically consistent way to do this, is to first weigh individual bits of information (of the form, “*How much do I trust what?*” and “*What have I learned from what?*”) and then combine all the prior partial beliefs using the *Durkin-summation function* that is commonly used in fuzzy logic applications.

Durkin summation

Durkin summation is a heuristic technique of combining certainty factors (i.e. beliefs) according to evidential reasoning in fuzzy logic [223]. Certainty factors are arbitrary (but self-consistent) measures of an expert's “belief” regarding some hypothesis. The maximum value of a certainty factor is $+1$ (meaning *definitely true*) and the minimum value is -1 (meaning *definitely false*). *Positive* values < 1 represent relative degrees of belief, and *negative* values > -1 represent relative degrees of disbelief. For example, suppose the CTN has a set, B , of N partial-beliefs (originating from N CT-agents and/or virtual assets) regarding the hypothesis that agent ‘X; belongs to the TN: $B=\{B_1, B_2, \dots, B_N\}$. Now use the Durkin fuzzy-summation function, $\oplus$ (where $|x|$ is the absolute value of x , and $\text{Minimum}(x,y)$ is the minimum of x and y):

$$B_i \oplus B_j = \begin{cases} B_i + B_j [1 - B_i], & \text{if } B_i, B_j > 0, \\ B_i + B_j [1 + B_i], & \text{if } B_i, B_j < 0, \\ \text{else } [B_i + B_j] / (1 - \text{Minimum}\{|B_i|, |B_j|\}), & \end{cases} \quad (71)$$


to compute the fused update to the CTN's belief at time t :

$$\beta(x \in \text{TN}; t) = \beta(x \in \text{TN}; t-1) \oplus \{B_1^+ \oplus B_2^+ \oplus \dots\} \oplus \{B_1^- \oplus B_2^- \oplus \dots\}, \quad (72)$$

where B has been partitioned into a union of two sets: one containing only positive valued components ($= \{B_1^+, B_2^+, \dots\}$), and the other containing only negative valued components ($= \{B_1^-, B_2^-, \dots\}$). While this function may, at first, appear somewhat strange, it possesses several intuitively desirable properties. For example, apart from its simplicity, we immediately note that all partial beliefs are bounded between -1 and $+1$: as long as $|x| \leq 1$ and $|y| \leq 1$, $|x \oplus y| \leq 1$. Adding a zero belief to a

nonzero certainty leaves the existing certainty unchanged: $x \oplus 0 = x$. Equal certainties that differ in sign combine, as expected, to yield zero effective certainty: $x \oplus (-x) = 0$. Figure 49 provides a few illustrative plots of the behavior of $x \oplus y$.

Figure 49. Behavior of Durkin-summation function $\oplus$; see equation 71



All three subexpressions of the Durkin-sum have very natural interpretations. For example, the Durkin-sum of two positive certainties (defined by the top-most expression), can be expressed in words as, “Reduce the influence of one certainty by the remaining uncertainty of the other, and add the result to the certainty of the first.” Likewise, the Durkin-sum of two negative certainties (given in the middle row), is equivalent to taking the Durkin-sum of two positive certainties and negating the result. Finally, combining a positive and negative certainty is equal to multiplying the difference by the reciprocal of the smallest remaining uncertainty.⁷⁸

78. Equation 66 represents a purely heuristic method of combining weights. Since single-feature weight components are nothing more than, essentially arbitrary, functions, they cannot justifiably be interpreted as, say, probabilities, for which a more rigorous inference engine based on Bayesian reasoning could be used. However, despite the lack of mathematical correctness, the performance of expert systems based on certainty factors has, on occasion, *outperformed* Bayesian reasoning (at least in systems designed to mimic human diagnostic judgment [223]).

Tactical action plan logic

The last element of the CTN's IP consists of personality weights that determine the set of actions that the CTN takes at time t (using its CT-agents and other assets), based on its updated belief map, calculated using INTEL data gathered at time $t-1$. An important design criterion for SOTCAC is for the program to be both *flexible enough* to encompass a variety of decision “types” (so that the user has some parametric variability to experiment with) and *simple enough* to not overwhelm the user (by the number and/or complexity of the parameters).⁷⁹

Recall (from our earlier discussion of **CTN Actions**; see page 192) that the CTN can choose to take a variety of actions, ranging from issuing orders to CT-agents to “move towards” specific “targets,” to eavesdropping on open communication channels, to capturing suspected T-agents, and/or destroying selected links.

The CTN's *Tactical Action Plan Logic* (TAPL) is based on what it believes to be the TN's current state, and consists of taking whatever actions it decides are necessary to reduce the TN's projected ability to function. There are two general classes of possible actions:

- *CT-agent movement/reconnaissance*, and
- *Targeted TN attacks*.

Both classes include actions that are based on INTEL-based assessments of the “value” of specific nodes, groups of nodes, and/or links. Since there is no objectively “best” measure of relative importance, all of the different network metrics introduced earlier in the section **Complex networks: behavioral and structural metrics** (including *characteristic path lengths*, *clustering coefficient*, *centrality measures*, *betweenness metrics*, *efficiency*, *etc.*) are—a priori, at least—useful for discriminating among potential targets. It is up to the user to determine which specific metrics are the most useful for specific scenarios, terrorist/counterterrorist characteristics and mission objectives and parameters.

79. This same criterion has successfully guided the development of both the EINSTEIN combat model [16] and the multiagent-based variant of the war-game SCUDHunt [17].

An important goal of developing SOTCAC, is to provide analysts with a logical inference engine to explore alternative measures of TN efficiency. *Are the TN's goals best thwarted by eliminating its most highly "connected" T-agents? Or is it better to target its cell leaders? Or perhaps to eliminate the links that operatives have with support agents?* Different contexts obviously require different counterterrorist options and discriminatory metrics. While generic metrics are useful for tuning the actions of the CTN during a given run, their real purpose is in providing the analyst with a "basis set" of primitive measures with which an extended series of SOTCAC-determined TN:CTN coevolutions can be used to reveal an important class of *composite network metrics* that more tightly correlate local CTN action with global TN efficiency.

Any strategy to defeat the TN that is based solely on optimizing primitive efficiency metrics, without regard for either the properties that characterize the *specific* TN the CTN is trying to defeat or the consequences that specific CTN *actions* have on the TN's own adaptive behavior, is doomed to fail; at the very least, such a strategy is likely to be neither robust nor optimal. The fundamental problem is not to determine which components need to be eliminated at what time—the answer to this problem is provided by a straightforward application of the same value-calculus used by EINSTEIN's agents in deciding what enemy agents to "shoot" at time t [16]—but to better understand the deeper relationship between the entire scope of actions that the CTN can take (in all possible contexts) and the consequent ability of the TN to perform its assigned mission. Stated even more succinctly, the fundamental question that SOTCAC is being designed to help intelligence analysts answer is,

"Given that the TN's ability to perform its mission is measured by some well-defined function, F_{TN} , what INTEL about the TN must the CTN acquire, and what actions based on that information must the CTN perform under what (local and/or global circumstances), such that the value of F_{TN} is minimized?"

Using its "belief matrix" database of fused intelligence, the CTN may choose to either target certain nodes for disruption or deletion—with the intention of crippling one or more components of the TN—or reconnoiter a set of T-agents that, because of their local ego-maps,

the CTN has identified as being the most likely *future* sources of intelligence data.

CT-agent movement/reconnaissance

The CTN issues CT-agent movement and/or reconnaissance orders as a function of its current belief matrix and its (user-defined) asset allocation personality. Although the CTN has control over both *physical* CTN-agents that collect HUMINT data and *virtual* (COMINT, ELINT, and SIGINT) resources, the logic by which members of either set are targeted at components of the TN is formally the same, and is derived from a user-defined weighed fitness function.⁸⁰

Initially, of course, all beliefs (regarding all T-agents, i and links l_{ij}) are set equal to zero: $b(i; t=0) = b(i, j; t=0) = 0$. To simplify some calculations, let us introduce the *dogma threshold*, $0 \leq \beta_{Dogma} \leq 1$, which is the threshold belief strength at which the CTN believes (“dogmatically”) that either a physical agent suspected of being a terrorist is definitely a T-agent, or, in regard to social network links, that a link between T-agents i and j definitely exists. The consequence of having any belief value exceed β_{Dogma} is that the CTN is not compelled to use any of its resources to obtain information about the T-agent (or link); thus no movement orders will be issued.

In general, of course, the CTN will harbor indefinite belief values, which must all be properly (i.e. objectively) weighed before issuing CT-agent movement and/or reconnaissance orders. Consider some basic motivations that the CTN may weigh before deciding how to direct its assets:

- *Maximize coverage* (on given turn, w_{MCov}): In the absence of other kinds of information (such as when starting out in a zero information state at time 0), the CTN wants to maximize the

80. The design borrows heavily from SCUDHunt’s *sensor placement logic* [17]. SOTCAC’s *belief matrix* is the analogue of SCUDHunt’s battlefield. Thus, where in SCUDHunt, the human user (or agent) allocates resources on a physical playing board to gather raw data (in the hopes of deducing where the actual SCUD missiles are located), in SOTCAC the CTN allocates resources on both a physical space and abstract (link) space, in the hopes of deducing where the T-agents are, what they are doing and what links exist among them.

number of suspected (but not yet confirmed) terrorists that are covered by available HUMINT assets (on a given turn). At later times, this motivation will be combined with other motivations listed below.

- *Maximize total (cumulative) coverage (w_{MCov}):* the CTN wants to minimize the number of suspected (but not yet confirmed) terrorists that have not yet been searched.
- *Maximize future coverage (w_{FCov}):* Taking into account CT-agent stationing and vulnerability, the CTN wants to maximize coverage for future turns.
- *Maximize global belief (w_{GBel}):* the CTN generally wants to maximize the value of his belief matrix for as many suspected terrorist (and/or links) as possible. This is to be weighed against, for example, simple coverage, which does not take the CTN's estimate of an agent's reliability of returning valid information that would tend to maximize belief. ("Is it better to have relatively unreliable HUMINT that covers a lot of territory or very reliable HUMINT that returns data from a smaller physical area?" The answer, of course, is "let's see" by exploring alternative options.)
- *Maximize local belief (w_{LBel}):* As a complement to the CTN's general desire to maximize belief over the entire playing field, is its motivation to "home in on" terrorists (and links) that already entail high belief. In particular, the CTN wants to exceed its *threshold belief strength* ($b_{threshold}$) value and come as close as possible to his *dogma threshold* (b_{Dogma}), for as many terrorists and links as possible, so that (consistent with the CTN's personality) it has a sound informational basis on which to identify a suspect as T-agent (or suspected link as definitely existing). By weighing this motivation greater than, say, coverage alone, the CTN would tend to associate HUMINT resources with suspects that have already been (or are currently being) observed, but whose cumulative belief value is less than $b_{threshold}$; rather than indiscriminately ordering CT-agents to collect data about other suspects (that, while also yielding an expected information gain, would likely still not push any site beyond the threshold belief).
- *Minimize redundancy (w_{Red}):* a terrorist suspect or link that has developed a belief that exceeds b_{Dogma} is deemphasized in the

CTN's movement/reconnaissance fitness function. This reflects the human behavior of minimizing the needless expenditure of resources to obtain information that already exists.

If a CT-agent is assigned to move toward (or reconnoiter) a suspected T-agent or link, the CTN will not issue new move orders unless either (1) the required threshold of belief regarding the suspected T-agent or link has been achieved (either through the given CT-agent's prior efforts or via the HUMINT reports of other CT-agents); (2) a higher-valued TN component (that requires immediate attention) is nearby; (3) or a component of the TN has been targeted for attack (see below).

Motivations can be either fixed, applied in succession at each turn but otherwise remaining unchanged throughout a game, or, more generally, updated, turn-by-turn. In the latter case, it is possible to model a CTN that starts a run with one personality and completes it using a different set of weights. Dynamic, state-dependent weights are also possible to design. In any case, on a given turn t , the CTN considers all possible options of allocating each of the HUMINT and COMINT resources under its control, and calculates the appropriate allocation fitness function.

For example, in the case of using virtual assets to identify and/or reconnoiter communication and social links, the CTN uses the COMINT *sensor allocation fitness function*, $F_{COMINT}(i,j; t)$, for each possible link (i,j) :

$$F_{COMINT}(i, j; t) = \begin{cases} w_{MCov} \cdot (\text{number of sites covered at time } t) \\ + w_{CCov} \cdot (\text{total No. of sites covered at least once for times } \tau \leq t) \\ + w_{FCov} \cdot (\text{minimal No. of sites covered at time } t+1) \\ + w_{GBel} \cdot (\text{belief gain throughout physical \& social spaces at time } t) \\ + w_{LBel} \cdot (\text{local belief at time } t) \end{cases} \quad (73)$$

Suppose the CTN controls virtual assets $CTa_1, \dots, CTS_N$, and that CTa_i covers a patch P_i of the TN's social network space. One asset may only be able to cover a single link; others may cover a small swath, albeit less reliably. The CTN sums the fitness function over all sites of a given assets patch to determine which of his assets maximize overall fitness:

$$AssetFitness(P_k) = \sum_{(i,j) \in P_k} F_{CTa_k}(i, j; t) . \quad (74)$$

Of course, the form of $F(i, j; t)$ is very flexible, and other terms can easily be imagined. For example, we could add a “randomizer” term, w_{Ran} , which when coupled with a $(1-w_{Ran})$ term multiplying the expression above, adjudicates between the likelihood that the CTN makes effectively random (i.e. mostly “dumb”) assignments and the likelihood that the CTN makes more intelligent decisions (consistent with its personality).

It is also simple (in theory) to include another form of information sharing among either the CTN’s HUMINT assets or components of the CTN itself, as an overarching “intelligence agency.” For example a *trust matrix* can be used to define how different components of the CTN (charged with analyzing information from different sets of HUMINT and COMINT resources) filter each other’s communicated fitness function evaluations, before combining that information with their own updated belief matrix. In other words, a given component of the CTN—say, $CTN_x \subseteq CTN$ —decides where to place its sensors partly as a function of its own fitness function, F_x (and therefore as a function of its own “personality”) and partly as a function of what other components, $\{CTN_y\}$, with whom it shares its analysis, are telling CTN_x about their own asset allocation plans (as defined by their fitness functions, $\{F_y\}$; filtered by the degree to which one component trusts the information and analysis that other components are communicating to it, as defined by the trust matrix, $T_{x,y}$):

$$AssetFitness(P_k) = \sum_{(i,j) \in P_k} \left[F_{CTa_x}(i, j; t) + T_{x,y} \cdot F_{CTa_y}(i, j; t) \right] , \quad (75)$$

where CTa_x and CTa_y are assets that belong to CTN_x and CTN_y , respectively.

Targeted TN attacks

The second set of TAPL weights focuses on various social network criteria that the CTN uses to assess the “value” that a given T-agent represents to the TN (and thereby to determine what component of the TN to “target” for attack).

Suppose the “attack” consists of removing a T-agent from the TN. There are two basic criteria that the CTN may use for targeting (and, indirectly, for judging whether one T-agent, X , merits a greater or lesser assigned “value,” from the point of view of its elimination from the terrorist network, than some other T-agent, Y):

- *Expected fragmentation*, and
- *Expected link-time increase*.

The first criterion consists of assigning greater (or lesser) value to those T-agents whose removal from the TN the CTN expects will cause greater (or lesser) fragmentation to the TN. Although the basic idea behind this criterion is intuitive—i.e., “fragmentation” refers to the degree to which the TN is expected to become disconnected, relative to its state of connectedness prior to the deletion of one or more of its T-agents—its precise definition, which we also want to include a set of tunable parameters so that it may be adjusted by the analyst, is subtle and will thus be defined only loosely here.

One simple approach is to define a fragmentation index, F , that simultaneously accounts for both the *number of connected components* that remain after the deletion and their *size* [113]:

$$F = 1 - \frac{2 \sum_i \sum_{j < i} r_{ij}}{N(N-1)} = 1 - \frac{\sum_k n_k(1-n_k)}{N(N-1)}, \quad (76)$$

where N is the number of agents, $r_{ij} = 1$ if agent i can reach agent j , and is equal to *zero* otherwise, and n_k is the number of agents in the k^{th} connected component; F thus effectively counts the number of pairs that are disconnected from one another. If the network is completely connected, $F=0$; at the other extreme, when there are no connected components, $F=1$.

The second criterion consists of measuring the degree to which existing lines of communication are *lengthened*. If the existing network is sparse enough, of course, then removing even a single T-agent can sever all connections and thereby immediately fragment the network; however, it will more often be the case that the “best” the CTN can do

by removing a T-agent is to make it more difficult for the TN to transmit messages by lengthening the time it takes for them, on average, to reach their intended receiver. The supposition is that networks with longer “shortest paths” transmit information less efficiently, and are therefore also more prone to discovery.

A simple metric—call it Lt , for “link time”—that accomplishes this is based on Latora’s and Marchiori’s [149,150] global “efficiency” metric, E_{global} introduced in an earlier section:⁸¹

$$Lt = 1 - E_{global} = 1 - \frac{1}{N(N-1)} \sum_{i,j} \frac{1}{Dist(i,j)}, \quad (77)$$

where $Dist(i,j)$ is the graph distance between agents i and j . As does F , Lt achieves its maximum value when the network is completely disconnected (i.e., all agents are completely isolated). Likewise, Lt ’s value matches the value of F when each of the k components of a network are cliques (i.e., complete k -graphs); since, in this case, $Dist(i,j)=1$ for all pairs. However, for all other cases, in which pairwise distances vary from agent to agent, and from component to component, Lt is able to resolve the internal inhomogeneities (whereas F is not).

For example, (1) *degree* (which simply counts the number of links a given node has with other nodes); (2) *betweenness* (which measures the extent to which a given node mediates, or plays the role of “information broker” between, any two other nodes); (3) *closeness* (which measures the extent to which a given node is “close to” other nodes in the network); and (3) *centrality* (which measures the degree to which a node plays an important role in the TN).

81. See **Complex networks: metrics**; page 83.

Conclusion

“The new terrorist networks are not like those of the past, which had a loose but identifiable hierarchy and structure...there are many autonomous cells we do not know about. They do not need orders from Osama bin Laden to carry out the jihad...The threat, even with Osama bin Laden gone, is very high. These groups are protean; they change their shape like the AIDS virus.”—Jean-Louis Bruguière⁸²

This paper proposes three major theses:

1. That terrorist networks are, fundamentally, *self-organized, emergent “virtual multicellular organisms”* that live as much in the physical domain as they do in an abstract information space;
2. That the topology, function and behavior of terrorist networks *coevolve with their enemy* (which, for purposes of the discussion, is generically called the “counterterrorist organization”); and
3. That the best approach to understanding how terrorist networks operate—how they form, how they grow and evolve, and how they adapt to both internal and external changes to their environment—is one that combines the precepts and methodologies of several heretofore related, but disparate, disciplines: *complex systems theory, network science, social network analysis, mathematical graph theory, and multiagent-based modeling.*

After reviewing the general theory behind existing mathematical modeling tools applicable to the study of dynamic networks, the paper introduces the conceptual design of a multiagent-based simulation called SOTCAC—that is currently in development—to facilitate a generative, exploratory study of the self-organized emergent dynamics of terrorist networks. SOTCAC’s novelty, as an analytical exploratory tool, rests on how it generalizes the conventional interpretation of “agents” to leverage the strengths of several interrelated disciplines that have not heretofore been combined in one model.

82. French Magistrate Jean-Louis Bruguière, quoted in “A Powerful Combatant in France’s War on Terror,” C. Hedges, *New York Times*, November 24, 2001.

While agent simulations of complex adaptive systems commonly use notionally physical agents to represent the local dynamical components of the systems being modeled, simulations built directly upon *nonphysical* components are less common. Rarer still, are simulations that combine physical and nonphysical agents into a broader class of *semiotic agents*, that both “live” in coupled physical and information spaces, and actively and adaptively “transform” their local topology.

SOTCAC’s agents live, and act, as much in the physical domain, as they do in an abstract realm of communication ties and invisible bonds of friendship, common grievance and trust. However, it is the terrorist’s dynamic, adaptive “social network” that lies at the heart of the model. The social network represents the space in which agents create and sever connections (and relationships) with other agents; the space in which material and nonmaterial resources are sought, fought over, acquired, and exchanged; and the space in which the agents’ physical activity is coordinated.

As the analysis of Al Qaeda’s pre-9/11 “social network” illustrates,⁸³ its patterns of communication links, and the meta patterns of evolving link patterns, provide important insights into the terrorist network’s operations. Real-world analysis, of course, is complicated by the fact that intelligence data is typically incomplete, uncertain, and inaccurate. Furthermore, since it takes time to extract meaningful information from raw data, and to derive defensible inferences from processed information, the intelligence analyst is often chasing an amorphous moving shadow; even as the analysis of existing data steadily proceeds, the data source itself—i.e., the properties of the terrorist network—is constantly changing.

Obviously, nothing can substitute for reliable intelligence data that describes a real network; and SOTCAC is in no way intended to replace the human analyst in drawing inferences from such data. Nonetheless, much can be learned about the structure and behavior of real networks by using mathematical graph theory and social network analysis to study the properties of notional structures that emerge in a multiagent-based dynamic graph model of terrorist networks.

83. See **Appendix 2: Mapping Al Qaeda**; page 221.

Just as *artificial life*, in complex systems theory, is predicated on the basic supposition that life “as it is” may be better understood by examining the dynamical possibilities of life “as it could be” [51], so too the agent-based approach advocated in this paper is predicated on the supposition that insights into the behavior of real terrorist networks may be gained by exploring the dynamical possibilities of terrorist networks “as they could be.” In particular, SOTCAC is designed to provide a conceptual scaffolding on which the *multidimensional space of possible terrorist network topologies* may be mapped and the relationship between agent dynamics on the micro-level and network behaviors on the macro-level may be systematically explored.

The structure and function of terrorist networks emerge, in SOTCAC, on three interrelated levels: (1) *dynamics on networks*, in which notional terrorist agents are assigned missions to strike physical targets and process/interpret information, search and acquire resources, and adapt to other agents’ actions; (2) *dynamics of networks*, in which topology of the evolving network itself is a fully dynamic, adaptive entity; and whose agents—i.e., nodes—build, maintain, and modify the network’s local (and therefore, collectively, its global) structure; and (3) *dynamics between networks*, in which the terrorist network and complementary counterterrorist network mutually *coevolve*: the terrorist network’s “goal” is to achieve the critical infrastructure (of manpower, weapons, financial resources, and logistics) required to strike, while the counterterrorist network’s mission is to prevent the terrorist network from achieving its goal.

As an example of the kinds of insights that the model offers the counterterrorist analyst, in principle, consider the role that “structural holes”⁸⁴ play in the information flow and social dynamics of a network. Since structural holes sit on the boundaries between flows (among otherwise separate cliques of knowledge structures), agents spanning these holes may be expected to wield enormous influence over the network’s local and global functioning and performance; they represent locations within the network from which other areas of the network can be reached with a minimal number of direct ties.

84. Structural holes represent only one byproduct of applying network theory to the operational characteristics of arbitrary social networks. See discussion in section **Complex networks: metrics**; page 96.

Analytical observations of this kind can assist intelligence analysts, operationally, in at least two ways: (1) *by strengthening vulnerability assessments*—by pinpointing the nodes and cliques of a network that are vital to information flow, “key” agents can be identified, whose removal can be expected to significantly degrade the network’s ability to “command” and/or to “control” its agent/cell-infrastructure; and (2) *by helping optimize data collection and resource allocation*—since the existence of structural holes may be inferred, indirectly, from the effect that as yet unobserved agents have on other parts of the network, the analysis can help focus the attention of HUMINT, COMINT or other INTEL data collection assets onto the most promising components of the network for further reconnaissance, investigation, and/or attack.

Using known (but typically incomplete and imprecise) data about a terrorist network, generative models like SOTCAC can be used to predict the properties and locations of critical components of the network that *are likely to exist, but have not yet been detected*. Structural holes are but one example of a “critical” network property, that is both objectively defined and measurable; however, it is by no means the only such measure, nor even, perhaps, the most relevant.

An important goal of developing SOTCAC is to provide analysts with a logical inference engine to interactively explore the efficacy of alternative measures of network “criticality.” *A priori, any* of the primitive network metrics introduced in this paper—*characteristic path length, clustering coefficient, centrality, betweenness, density, efficiency, etc.*—may be used to gauge the relative importance of selected components to a network’s overall ability to function. The deeper question is, “*Which of these metrics—or, better—which combination of metrics, and in what dynamical context, provides the counterterrorist organization the optimal basis on which to base its actions?*”

Is a terrorist network’s mission best thwarted by eliminating its most highly “connected” terrorists? Or is it better to target its cell leaders? Or perhaps to leave its most highly connected terrorists untouched (albeit with continued covert surveillance), but eliminate the links between mid-level operatives and supporting agents?




Figure 50. Schematic depiction of SOTCAC's superspace of TN-CTN coevolutions

Any strategy to “defeat” a terrorist network that is based solely on optimizing primitive efficiency metrics, without regard either for the properties that characterize the *specific* terrorist network being attacked, or for the dynamical consequences of how a given terrorist network *adapts* to the actions taken against it, is doomed to fail. At best, such strategies will merely be shortsighted and suboptimal; at worst, they may inadvertently inject a dynamics into the system that does far more harm than good. For example, although a simple-minded, albeit intuitive, strategy predicated on the assumption that removing a strong, charismatic leader will result in a systemic collapse of a network, *may* succeed; it is also at least as likely to not just fail, but fail *catastrophically*, by providing the remaining components of the “decapitated” network an “energy” that both stimulates and strengthens it.

A successful strategy can only emerge by understanding—on a fundamental level—how the set of *all possible actions that a counterterrorist organization can take* (in all possible contexts) is dynamically related to the set of *all possible measures of how well a terrorist network functions*. Such an understanding can only come about by systematically exploring the multidimensional space that contains all possible outcomes of terrorist ↔ counterterrorist coevolutions (figure 50 shows a schematic depiction of this multidimensional “superspace”). SOTCAC is being developed precisely for this reason.

In general, a complex adaptive network cannot easily be “defeated” by removing a few (or even many) of its *pieces*; rather, one must find ways to disrupt the autopoietic web of *interactions* that sustain and nurture it. Paraphrasing the quote that appears at the top of the preamble (see page *i*):

Osama bin Laden is a phenomena, not a person.

Therein lies Al Qaeda’s strength *and* weakness; but it is a weakness only if Al Qaeda is understood, and dealt with, as a complex adaptive network, and not as a set of simple, conventional “targets.”

Appendix 1: *Social network analysis*

“Social structure becomes actually visible in...the movements and contacts one sees...We should also be able to see structure in the life of an American community if we had a sufficiently remote vantage point, a point from which persons would appear to be small moving dots...We should see that these dots do not randomly approach one another, that some are usually together, some meet often, some never...If one could get far enough away from it human life would become pure pattern.”

—R. Brown, *Social Psychology* (1986)

Social Network Analysis (SNA) studies the properties of relationships among social entities. More precisely, SNA applies a broad set of mathematical methods (most notably those derived from mathematical graph theory) to study the—often latent or even invisible—patterns of interactions among individuals (which we will from now on refer to generically as *agents*).

Consider a typical business company and its ubiquitous hierarchical organization chart consisting of the President at the top, who sits on top of a group of Vice Presidents on the second tier, who are followed on down the chain by Division heads and their group managers, senior and junior analysts, and so on. Does this chart reflect how the organization really conducts its business? Or does the organization depend more on hidden, behind-the-scenes, communication links and self-organized groups of people whose ties are not formally recognized by the organizational chart, but through which the organization’s vital important flows? SNA addresses these questions by probing for, and examining the properties of, these hidden networks within networks. Of course, the analysis extends to a much broader set of real-world networks and phenomena than simply networks that represent business organizations.

Dating back to the 1930s—and Jacob Levi Moreno’s pioneering work using hand-drawn “sociograms” to trace communication lines among acquaintances [170]—SNA is today widely used in economic, sociological, ethological, and anthropological contexts. Its fundamental

assumption is that in order to understand how a network functions, *as a whole*, one must first understand how the network's constituent elements are all interrelated. While SNA has traditionally focused almost exclusively on *static* properties of the systems it studies—and thus has not (until very recently) examined self-organized emergent properties of *dynamic* nets—there is obvious and considerable overlap between SNA's basic philosophy and that of complex adaptive systems theory; certainly, both disciplines recognize the importance of understanding how the parts of a system are mutually dynamically coupled. Moreover, as the complexity, size, and diversity of the social networks that SNA studies all increase, the broader class of methodological tools of complex network science becomes increasingly important (see, for example, the discussion in the next section of network metrics, many of which are derived from the more general analysis of complex networks).

A critical goal of this paper is to demonstrate how conventional SNA may be generalized and enhanced—specifically, via the application of multiagent-based simulation techniques—to enable researchers to explore the self-organized, emergent behaviors in adaptive, coevolving dynamic networks. A prime example of why it is vital to begin developing analytical and modeling tools to study dynamic graphs is the emergence of global terror networks, such as *Al Qaeda*, about whose structure we know little; and even less about how it reacts, adapts and evolves over time.

SNA's central concept is that of applying a *set of relations* (and identifying *information flows*) to the agents making up a network; while de-emphasizing, often eliminating, a focus on the set of *attributes* of agents (a focus which is dealt with by other branches of social science). Some common relations are: agent X “communicates with” agent Y; X “meets face-to-face with” Y; X “gets information from” Y; X “coordinates with” Y; X “knows about but has never met” Y; X “provides resources for” Y; and so on. A given relation may also be assigned a set of associated properties, such as directionality, strength, reciprocity, frequency and duration. All of these properties are used in the design of SOTCAC.

SNA studies typically make the following additional assumptions [41]:

- *Agents and their actions are interdependent and autonomous.*
- *Relational links between agents are channels that transfer (or allow the transfer) of resources and information.*
- *The connectivity along a network's constituent agents is, by itself, an essential dynamical ingredient underpinning the network's function and behavior, and both provides opportunities for, and imposes constraints on, individual action.*
- *Network models conceptualize abstract topology (social, economic, political, etc.) as self-organized, emergent patterns of relations among its agents.*

While agents may be the focus of attention for a given problem (node-centric measures such as centrality, prestige, betweenness, for example, are integral metrics in many studies; see discussion below), SNA is usually concerned with understanding the properties and behaviors of larger topological entities consisting of a collection of agents and the links among them: *Dyads* (two agents and their links, at which level distance and reachability, structural and other notions of equivalence, and tendencies toward reciprocity are prevalent), *Triads* (three agents and their links, at which level balance and transitivity become important), or larger systems (subgroups of individuals, or entire networks, and networks of networks; at this level concepts such as cliques, centralization, connectedness, cohesion, diameter and density come into play).

The kinds of questions and issues that SNA typically addresses are, not unexpectedly, similar to the basic questions one asks of complex networks in general, although SNA-related questions more strongly emphasize processes and relationships:

- *Who are the most important “players” in a network, and what are the important structures?*
- *Who is the “star” of the network (i.e., most highly connected individual)? Who plays the role of mediator, or liaison? Which members of a network act as “bridges” between otherwise disconnected subgroups?*

- *How does information flow throughout a network? Are there bottlenecks? What kinds of topological changes can be made to enhance the information flow locally? What about globally?*
- *What kinds of endogenous and exogenous dynamics are responsible for the networks evolution and sustainability?*
- *How does the network adapt to local structural failures?*
- *What is the feedback mechanism between the network's dynamics and underlying topology?*

Example

Both as an example of the kinds of insights SNA can provide into how a social networks functions, and as a motivation for the more technical discussion of general network metrics in the next section, consider an imaginary company that consists of ten workers: *Andy, Bill, Claire, Drew, Eliot, Frank, Gary, Heather, Ilya, and Jane.*

Figure 51. Sample SNA “deconstruction” of an imaginary business that consists of ten workers; the columns show extracted *relationships*, graph *structure*, and graph *visualization* of the company

Information		Graph Structure		Graph Image
Name	Works with...	Node	Links	
Andy	Bill, Claire, Drew, Frank	n ₁	l ₁₂ , l ₁₃ , l ₁₄ , l ₁₆	
Bill	Andy, Drew, Eliot, Gary	n ₂	l ₂₁ , l ₂₄ , l ₂₅ , l ₂₇	
Claire	Andy, Drew, Frank	n ₃	l ₃₁ , l ₃₄ , l ₃₆	
Drew	Andy, Bill, Claire, Eliot, Frank, Gary	n ₄	l ₄₁ , l ₄₂ , l ₄₃ , l ₄₅ , l ₄₆ , l ₄₇	
Eliot	Bill, Drew, Gary	n ₅	l ₅₂ , l ₅₄ , l ₅₇	
Frank	Claire, Drew, Gary, Heather	n ₆	l ₆₃ , l ₆₄ , l ₆₇ , l ₆₈	
Gary	Drew, Eliot, Frank, Heather	n ₇	l ₇₄ , l ₇₅ , l ₇₆ , l ₇₈	
Heather	Frank, Gary, Ilya	n ₈	l ₈₆ , l ₈₇ , l ₈₉	
Ilya	Heather, Jane	n ₉	l ₉₈ , l _{9,10}	
Jane	Ilya	n ₁₀	l _{10,9}	

Figure 51 shows a schematic of the kind of raw information that SNA uses to evaluate the internal relation patterns of this system.⁸⁵ Two

85. The graph image shown in this figure is called a “Kite Network” and was introduced by David Krackhardt, a leading social network analyst and developer of a widely used social network visualization package called *KrackPlot*: <http://www.andrew.cmu.edu/user/krack/>.

nodes are connected in the graph image (shown on the right-hand-side of figure 51) if the people they represent “work with” each other. For example, *Andy* (n_1) works with *Bill*, *Claire*, *Drew* and *Frank*, but not with any of the other workers. Thus, n_1 is connected only to the nodes n_1 , n_3 , n_4 and n_6 .

This example illustrates the meaning of three basic measures of centrality in a network: *degrees*, *betweenness*, and *closeness*.

Consider the dynamic role that some of these individuals play, in the context of their position in the social network of the company, as depicted in the graph image. *Drew* is clearly the worker with the most links (6) to others in the organization; i.e. *Drew* is the best *connected* (sometimes also referred to as most *active*).

Note that this does not imply that *Drew* is the most “important” worker, for two reasons: first, because we cannot draw this conclusion without first learning more about what this company does and what kind of information is more or less vital to its function; and second, because it is generally true in social networks that what is of significance—*topologically*—is not so much how many links are attached to given node, but where the links anchored on that node lead, and how they connect parts of the network that would otherwise be disconnected [134].

Observe that all of *Drew*’s links are with individuals in the same clique; links that are to some degree also obviously redundant, since many of those he is connected to are already connected to each other.

As an example of how a node with fewer connections can be viewed as serving a more important “role” within the network, consider *Heather* (n_8). *Heather* has fewer links to others than *Drew* (she actually has fewer links than is the average for the entire company)—*but*—she is positioned prominently *between* two groups of the network: the clique on one side, and *Ilya* and *Jane* on the other. In the context of communications, *Heather* plays the role of “information broker,” and is the de facto go-between during exchanges between the two groups. In general, nodes that have a high *Betweenness* measure (see below), potentially can exert a strong influence over the information flow in the network. On the other hand, nodes with high *Betweenness*, such as

Heather, also represent some of the vulnerable components of a system. Without *Heather's* presence in the network in figure 51, for example, *Ilya* and *Jane* would effectively be cut off from the rest of the company.

Now consider *Frank* and *Gary*. Like *Heather*, each of them has fewer links than *Drew*, and so do not score high in *Closeness*; but these two individuals are closer, on average, to all other workers than is any other node in the network. From a SNA perspective, Frank and Gary therefore lie at the heart of the network, for they are in a position to “hear about” what might be happening elsewhere in the network first.

What of Ilya and Jane? Given their relative isolation at the tail end of the social network, one might, at first, be tempted to call these individuals only peripheral—even unimportant—players in the company. Neither of these workers has a high degree, neither is particularly close to anyone else, and both are too far from the others to act as “go-betweens.” However, *Ilya's* and *Jane's* solitude may only be apparent, and not indicative of the true role they play. They may appear to be separated from others simply because the data used in generated the graph in figure 51 was incomplete and/or erroneous (which is a problem that plagues all SNA deconstructions of terrorist networks; see discussion in next section); or they may play more important roles in other, overlapping, social networks that map relations other than the single “works with” relation shown here. This simple example reminds us that—while deconstructing social networks—we must always be mindful of the *fuzziness*, *incompleteness* and/or *erroneous* nature of the data used in generating social maps.

So far in this example, we have examined the differences among the dynamic roles that individual workers (seem to) play in the imaginary company depicted in figure 51; i.e., we have looked at some basic *local* properties. Equally important, of course, is a network's *global* structure; or, more precisely, the relationship among a network's local centralities. A highly centralized network, that has one or only a few central nodes—such as the one in this example—is highly vulnerable to attack [136]. If any of these highly central nodes are damaged or destroyed, the network may break apart into disconnected subnets, or become completely disconnected. A single highly connected node

represents a critical vulnerability of the system. On the other hand, less centralized networks are generally more resilient to attack. As the a network becomes less and less centralized, as a whole, it takes more and more targeted attacks on individual nodes to disconnect the system.

Appendix 2: Mapping Al-Qaeda

“Terrorist networks are not armies...today the world's most dangerous aggressors are not military organizations with divisions but self-organized networks of terror.”

—Albert-Laszlo Barabasi, *University of Notre Dame*

One obvious way to understand the structure (if not dynamics) of *real* terrorist networks, is to map them using existing data, and study the resulting structures as mathematical graphs. Valdis Krebs, a management consultant and developer of a social network analysis program called *InFlow*,⁸⁶ set out to do precisely this shortly after the tragic events of September 11, 2001 [193,195].


Krebs applied exactly the same SNA methodology his company successfully uses for mapping organizational structures for business firms (such as IBM, TRW, and Raytheon, among many others) to map the terrorist networks responsible for the attacks. Using public information gleaned from major newspapers such the *New York Times*, the *Wall Street Journal*, the *Washington Post* and the *Los Angeles Times*, Krebs began to “connect the dots.” Figure 52 shows a screenshot of an early information table containing known facts about the hijackers, published by the *Sydney Morning Herald* in Australia, on September 24, 2001.

Using the names of the 19 individuals identified as the attackers, Krebs used the WWW search engine Google to obtain links to other public information associated with them, including their backgrounds, possible associates and connections to others, financial ties, and further details regarding their flight training and other events they may have been associated with.

86. *InFlow* is a network analysis toolkit that maps and provides metrics for knowledge exchange, information flow, communities of practice, networks of alliances and other kinds of networks within and between organizations; see <http://www.orgnet.com/inflow3.html>.

By the middle of October 2001, Krebs (being careful to disentangle confirmed facts from those that were obviously false, such as erroneous stories that appeared about a cell in Detroit) had obtained enough information to begin mapping the links among the terrorists.

Figure 52. Screenshot of an early information-matrix of hijacker data



Trusted Prior Contacts


Figure 53 shows an early map produced by Krebs that shows only those connections that he found to exist among *trusted prior contacts*; or those forged between terrorists that have lived and trained together. Note that the terrorist nodes are color coded according to the flight they hijacked.

While the relationships shown in figure 53 may not all be correct, and are certainly incomplete at this early juncture (a deficiency that

would quickly be remedied as more reliable information was obtained; see below), this preliminary structure nonetheless already displays interesting structural properties.

For example, Krebs was immediately impressed by its *intelligent* sparseness; by how distant many of the hijackers on the same team were from each other [193]. This sparseness is arguably less an unfortunate artifact of incomplete information, and more a consequence of intelligent administration by the terrorist cell itself.

Figure 53. Social network of *trusted prior contacts* of 9/11 hijackers; taken from “Mapping Networks of Terrorist Cells” (Krebs [195])



The 19 agent network shown in figure 53 has an average path length of 4.75 steps; and a few hijackers are separated by more than six steps. The network clearly trades efficiency of information flow for secrecy.

By keeping cell members as distant from each other (and from members of other cells) as possible, damage to the network as a whole if one of its members is captured or otherwise compromised is minimized.

Krebs notes that this tradeoff is consistent with observations made by Osama bin Laden (as transcribed from a video released by the US Department of Defense on December 2001; [195]):

*Those who were trained to fly didn't know the others.
One group of people did not know the other group.*

While figure 53 provides a snapshot of Al Qaeda's cell compartmentalization, it cannot be telling the entire story. However great may be the need for secrecy, at some point terrorists must coordinate their resources and actions.


Meeting Ties

Figure 54 shows Krebs' update of the social network map of 9/11 hijackers, using both *trusted prior contacts* links (as in figure 53) and additional links indicating short-lived *meeting ties* (shown in yellow).

Superficially, there appears to be little obvious difference between the two networks shown in figures 53 and 54. Although Krebs' graph-layout program has obviously rendered the two networks differently from a purely visual standpoint, the only objective difference between the two graphs consists of the six yellow links representing transitory communication lines. However, the effect that these additional links have on the *dynamics* of the network is substantial.

Specifically, the six additional links provide temporary, but valuable, "shortcuts" to information flow and thus serve as transient aids for both *collaboration* and *coordination* [224]. For example, one can show (using Krebs' *Inflow* program or any other freely available SNA toolkit; see Appendix 3), that the six yellow links in figure 54 reduce the mean path length of the network shown in figure 53 by 40%. While the assertion that a relatively small number of links can make a dramatic difference in a network's topology has been made before in the context of small worlds graphs [74], Krebs' example illustrates that these ideas, and phenomena, are far from being "merely" academic.

Figure 54. Social network of 9/11 hijackers showing both *trusted prior contacts* (as in figure 53) and short-lived *meeting ties* (indicated in yellow); taken from “Uncloaking Terrorist Networks,” by Krebs [3]



While a terrorist network must constantly weigh its need to use its links to accomplish its objectives against maintaining secrecy, if it is to ultimately succeed in its mission it must periodically risk exposure to actively coordinate its assets. It is at these critical moments, of course, that the terrorist network (or at least a component of it) becomes especially vulnerable to eavesdropping by counterterrorist forces, disruption and/or destruction. An important component of the conceptual model SOTCAC, introduced in the next section, is to explore the dynamical consequences of obeying a local optimization rule (to be executed by a notional terrorist cell) of the form:

Use minimum number of active agents in cell C to maximize information-flow throughput in region R of network.

The last step of Krebs' evolving social network map of the connections among the 9/11 hijackers was to include their support group; the members of which, though not direct participants in the terror attacks, were nonetheless indirect accomplices, serving as channels for money transfers, and sources of training and other required skills and resources.

Direct & Indirect Links

Figure 55 shows Krebs' final map of the 9/11 hijackers, which adds all direct and indirect links with supporting agents to those already included in figures 53 and 54. The thickness of each link in figure 55 is proportional to the estimated strength of the connection between the agents it is representing a bond between.

What is immediately apparent from even a casual visual inspection of figure 55—and may be confirmed after performing a quantitative social network analysis—is that Mohamed Atta (who appears roughly at the center of the graph and whose node is colored green) plays a central role. Recall that Atta's role in the attacks, the importance of which as ring-leader is now well established, was unknown at the time that Krebs created this graph [3]. Of all the agents appearing in the network, Atta ranks highest on the lists of basic network centrality metrics, such as *degree*, *closeness*, and *betweenness* (see **Complex networks: metrics**; page 77). The figure thus strongly displays “scale-free” characteristics, with Atta as the dominant hub.

Recalling our earlier discussion of scale-free networks (see page 51), Krebs' snapshot view of (a part of) Al Qaeda is also consistent with the Barabasi-Albert model of scale-free network formation; namely, that such networks form, and grow, naturally whenever nodes (or agents) are added with preferential attachment. One may assume that recruits are predisposed to attach themselves to a terrorist organization's most influential members; and that recruiters look first to those they already know and trust to offer membership in the organization. While figure 55 does not, by itself, “prove” that Al Qaeda is scale-free,

it is consistent with what one expects to find for networks that have evolved via preferential link attachment.

Figure 55. Social network of 9/11 hijackers showing *trusted prior contacts* (as in figure 54), short-lived *meeting ties*, and associated *support network*; taken from “Uncloaking Terrorist Networks,” by Krebs [3]



An important reason for developing SOTCAC is to be able explore such basic issues as the dynamical relationship between how networks first form and grow, locally, and their eventual, emergent global structure.

Observations

Recall that “degree” measures *activity*; “closeness” measures *accessibility* to (and ability to monitor) others in the network; and “betweenness” measures the extent of *control* that an agent has over information flow in the network (i.e. an agent that ranks high in betweenness likely plays the role of information broker in the network). Ranking highly (or, as in Atta’s case, attaining the *highest* rank) in all three of these metrics strongly suggests a leadership role.

While these, and other, social network metrics have been developed in the context of communication and learning in business organizations, because they depend only on a network’s abstract structure, they apply equally well to any interconnected community of agents about whose internal structure, dynamics, and general properties of information flow the analyst is interested in studying. Krebs’ plots and business-organization-derived analyses thus reveal important properties and key individuals in the information flow of Al Qaeda’s internal dynamics (as best as they can be inferred from publicly available data).

A more careful analysis of figure 55 also uncovers less obvious patterns that hint at Al Qaeda’s underlying command structure. In “Six Degrees of Mohammed Atta,” for example, Thomas Stewart points out that...

...the greatest number of lines lead to Atta, who scores highest on all three measures, with Al-Shehhi, who is second in both activity and closeness, close behind. However, Nawaf Alhazmi, one of the American Flight 77 hijackers, is an interesting figure. In Krebs's number crunching, Alhazmi comes in second in betweenness, suggesting that he exercised a lot of control, but fourth in activity and only seventh in closeness. But if you eliminate the thinnest links (which also tend to be the most recent -- phone calls and other connections made just before Sept. 11), Alhazmi becomes the most powerful node in the net. He is first in both control and access, and second only to Atta in activity. It would be worth exploring the hypothesis that Alhazmi played a large role in planning the attacks, and Atta came to the fore when it was time to carry them out. [225]

Other latent patterns are discernible only in hindsight after key players of a network are identified in a social network map.

Voss and Joslyn, for example, during a study to apply advanced knowledge integration tools to large data sets of terrorist-related information (conducted at the *Los Alamos National Laboratory* [194]), used over 1000 open source reports correlating the identities of the hijackers with data on individual people, terrorist groups, events and expertise.

Figure 56. Map showing that 11 of the 19 hijackers came from a single stretch of Highway 15 of the Asir province of Saudi Arabia; screenshot from [194]



In a visual approach similar to Krebs', Voss and Joslyn first used a *VisioPro* chart to display the relationships among people, terrorist groups, terrorist cells, religious leaders, Mosques and terrorist events. As links were added from their database, theretofore "invisible" links emerged, such as the connection between Al Qaeda and Hizbollah.⁸⁷

The critical role that *prior trusted contacts* play in establishing and maintaining a resilient hidden core of a covert network is evidenced

by Voss's and Joslyn's discovery that *fifteen* of the nineteen hijackers not only came from Saudi Arabia, but *twelve* of them came from towns stretching out along a single highway in the southwestern Asir province of Saudi Arabia (see figure 56).

Ties between terrorists, forged in towns where they lived, nurtured by common grievances and life's experiences, and later strengthened while they trained together in Afghanistan, were all solidly—and *silently*—in place by the time the hijackers came to the United States. This invisible ambient field of links, glued by mutual *trust*, known only to those already within the core of the network and all but unknown to those outside of it, was rarely active; its strength is due to the fact that if the need arises, any or all of these “dormant” links can instantly be activated.

Lessons

There are several important observations worth emphasizing about what we have learned from this example:

1. While none of the individual “snapshots” of the 9/11 terrorist network is complete, or completely accurate, they collectively suggest that the real network possesses an obviously *intelligent design*, one that is highly tailored to efficiently carry out the terrorist operation.
2. This network does not result from a willful imposition of order, from a top master-mind on down the chain of command, but is instead a *self-organized, emergent entity* that is generated by a combination of Al Qaeda's mission requirements and the (decentralized and essentially autonomous) initiative of its agents.
3. The strong cohesion among the 9/11 hijackers was due to the trust that had been built up among them *years before the attacks*, dating back to their mutual participation in Al Qaeda's training

87. A link, dating back to 1993, emerged between Imad Fayeز Mugniyah (head of security of Hizbollah and possible Al-Qaeda member) and Osama bin Laden [224].

camps, and, in many cases, even farther back to their home towns.

The most important analytical tool that social network analysis adds to counterterrorism, is its ability to map the “*invisible dynamics inside*” [18] a terrorist network; both qualitatively (via graph visualization), and quantitatively (using its well-defined store of network metrics). In the case of networks that evolve over long periods of time, it is particularly vital to have an ability to map connections, as that is the only way in which the counterterrorist organization can gain any sense of the network’s emerging structure and potential strength.

For example, Krebs [193] points out that a map of the activities and contacts of two of the 9/11 hijackers (Nawaf Alhazmi and Khalid Almihdhar), that the Central Intelligence Agency had identified as “Al Qaeda suspects” when they met face-to-face in Malaysia in 2000, reveals that (a) all 19 hijackers lie within two steps of the original suspects, and (b) each has multiple links back into the larger Al Qaeda network.

4. The network, as a whole, is *relatively sparse*. The 19 members of Al-Qaeda’s 9/11 strike were relatively isolated from one another.
5. *There is clear leadership structure, despite the obvious lack of any formal command and control hierarchy*. Mohammed Atta leadership role is “obvious” only from a bird-eye view of the entire 9/11 network; it is Atta’s *dynamic role* in the network (as evidenced by centrality and closeness measures, among others) that testifies to his importance to the terrorist operation. Note also that the “head” of Al Qaeda—Osama bin Laden—is absent from the (perceived) communication channels and chain of command.

Osama’s role, dynamically, is not to manage the day-to-day activities of Al Qaeda’s many cells and operatives; rather it is to provide whatever financial and logistical resources are necessary for conducting Al Qaeda’s missions, and to engender and maintain a pervasive, system-wide, unified focus and vision.

The role of lower ranking members of the network is to nurture robust, but secretive, interconnectivity.

6. *Liability*—the overlap between Atta’s leadership responsibilities and his unique skills, as one of the trained pilots, made him (and other pilots who were also cell leaders) a liability; and represented a key vulnerability in Al Qaeda’s mission structure.
7. *Resiliency*—the structure of the 9/11 network was robust in that it was relatively impervious to the removal of the cell leader. If Atta were eliminated, Marwan al-Shehhi would have likely emerged as the new “leader” due to his high degree of connectivity (Marwan had the second highest ranking *degree*, and was ranked fourth in terms of both *betweenness* and *closeness* [195]).

Al Qaeda’s robustness extends to the highest levels of leadership. Because bin Laden’s primary function is being a “catalyst” that stimulates—and *energizes*—the system [80], the most likely consequence of removing bin Laden from the network (say, by assassination) is the creation of a “structural hole” [134], representing the critical system role bin Laden previously played. This hole will not only be immediately filled (so that, from Al Qaeda’s perspective, no real “harm” will have been done), but the dynamics of how the network *adapts* to bin Laden’s loss will only likely further stimulate and energize the system [marion].

8. The core set of 9/11 hijackers was surrounded by a *larger support network*, which both provided necessary resources and served as backchannel communication lines that, effectively, maintained otherwise sparse links among mission operatives connected.

Appendix 3: *Social network analysis and SOTCAC-related development resources*

SOTCAC is being developed using Microsoft's *Visual C++ Version 6.0* and *Developer Visual Studio 6.0*:

`http://msdn.microsoft.com/developer/`

However, because SOTCAC makes heavy use of existing graph algorithms and social network analytical tools and measures, the use and assistance of several different adjunct toolkits is essential to reduce development time and effort.

This appendix contains brief descriptions of (and WWW URL links to) some of the toolkits that have been tested for future use toward this end. Of these, four were chosen, although—because of the considerable overlap in the functionality of many of these programs, as a group—it is likely that several others will play some role in SOTCAC's development):

1. *Combinatorica* (which is collection of functions that are integrated into the *Mathematica* program);
2. *Pajek* (which is a freeware graph visualization and analysis program, and is widely used in the social network analysis community); and
3. *C++ Graph Analysis* and *MFC Graph Visualization* toolkits from *Tom Sawyer Software* (which, together, provide a complete set of customizable class libraries).

While the first two toolkits are used more for quick-look analysis and rapid prototyping of algorithms than software development, the latter two contain tested algorithms for virtually all of the social network metrics that SOTCAC requires, as well as provide the primitive visualization functions needed for SOTCAC's networks and GUI.

AGNA

AGNA (Applied Graph & Network Analysis) is a platform-independent freeware application, developed by Marius Benta, and designed specifically for social network analysis. In addition to providing basic visualization functions, AGNA automatically calculates various network properties and metrics (including diameter, density, cohesion, geodesics, four types of centrality coefficients, as well as many network-level aggregates and/or statistical descriptions).

Version 2.1.1 may be freely available on this website:⁸⁸

<http://benta.addr.com/agna/download.htm>

aiSee

aiSee is a commercial graph layout program that automatically calculates a customizable graph layout; data is provided using a proprietary (but straightforward) textual input format GDL (graph description language). The resulting layout is then displayed, and can be interactively explored, printed and exported to various graphic formats. *aiSee* contains 15 basic layout algorithms (including force-directed layout; see page 34), and is optimized to handle very large graphs (that contain up to 10^6 nodes).

aiSee is free for noncommercial use and may be downloaded from the *aiSee* homepage:

<http://www.aisee.com/>

Combinatorica/Mathematica

Combinatorica is a collection of over 450 algorithms for discrete mathematics and graph theory written in *Mathematica*.⁸⁹ Although *Combinatorica* is very comprehensive, it is not particularly well suited for analyzing large graphs. Its strongest asset is its tight integration within

88. AGNA requires *Java Virtual Machine* (Version 1.3 or above), which is available for free from Sun (<http://java.sun.com/getjava/>).

89. <http://www.wri.com>.

Mathematica. Codeveloped by Sriram Pemmaraju and Steven Skiena [95], it is included with the standard Mathematica distribution in the directory *Packages/DiscreteMath/Combinatorica.m*; it may also be downloaded from *Combinatorica*'s homepage:

www.cs.sunysb.edu/~skiena/combinatorica/index.html

Combinatorica Graph Editor

Combinatorica Graph Editor (CGE; developed by Levon Lloyd) is a Java-based graph editor that works with *Combinatorica*:

<http://www.cs.sunysb.edu/~lloyd/grapheditor>

GraphPlot

Mathematica v5.1 (released in November 2004) now includes an extremely powerful *GraphPlot* package that contains (among many other options), algorithms for spring embedding, the spring-electrical model, high-dimensional embedding, radial drawing, and layered drawing methods. It supports graphs saved using the *Combinatorica* package, and is designed to work efficiently for very large graphs.

Graphviz

Graphviz is a graph layout program that takes descriptions of graphs in a simple text language, and generates layouts (using one of six embedded graph layout algorithms) that can be saved in various formats (such as JPEG images and postscript). One caveat is that graphs are must be generated externally to the program; although there are many utilities that provide this function, such as *Grappa*, a JAVA-based “front-end” to *Graphviz*: <http://www.research.att.com/~john/Grappa/>.

Graphviz is open source licensed software:⁹⁰

<http://www.graphviz.org/>

90. A Mac OS X edition of *Graphviz* (that has won two 2004 Apple Design Awards), is available at this address: <http://www.pixelglow.com/graphviz/>.

JUNG

JUNG (*Java Universal Network/Graph Framework*) is an open-source JAVA-based software library that provides a common and extendible language for the modeling, analysis, and visualization of data that can be represented as a graph or network. JUNG supports a variety of graphs (including directed, undirected graphs, multi-modal, and hypergraphs); and includes implementations of a number of basic algorithms from graph theory, data mining, and social network analysis (including routines for calculating clustering, decomposition, optimization, random graph generation, statistical analysis, and calculation of network distances, flows, and centrality measures. JUNG also allows users to interactively explore network data by providing an embedded visualization package.

JUNG version 1.5 is available at this address (which also contains links to additional third-party libraries required to use JUNG):

<http://jung.sourceforge.net/download.html>

LEDA

LEDA is a general C++ class library for efficient data types and algorithms for graph- and network problems, geometric computations, and combinatorial optimization. Introduced in 1989 as an academic research project, it is currently distributed commercially by *Algorithmic Solutions Software GmbH*:

<http://www.algorithmic-solutions.com/>

Two related toolkits, that use LEDA, are *GDToolkit* (Graph Drawing Toolkit; <http://www.dia.uniroma3.it/~gdt/>), which is used for graph drawing and layout; and *AGD* (Algorithms for Graph Drawing). AGD (<http://www.ads.tuwien.ac.at/AGD/>) offers a broad range of existing algorithms for drawing 2D graphs as well as tools for implementing new algorithms.

Maple/Networks Package

Maple—which, like *Mathematica*, is a commercial symbolic programming language—includes a networks package consisting of about 100

command that include implementations of basic graph routines: network flows, connectivity, disjoint spanning trees, all-pairs shortest path, single-source shortest path, minimum weight spanning tree, Tutte polynomials, and characteristic polynomials. *Maple* also includes routines for visualizing graphs, although they are not as sophisticated as those included in Mathematica's Combinatorica package (see above). *Maple* version 9.5 is distributed by *Maplesoft*:

<http://www.maplesoft.com/products/maple/>

NetDraw

NetDraw is a free program for drawing networks. Developed by Steve Borgatti, it uses several different algorithms for 2D graph layout. *NetDraw* reads both UCINET system and *Pajek* text files (UCINET and *Pajek* are both described below). *NetDraw* saves graphs as EMF, WMF, BMP and JPG files; it can also save data to *Pajek*.

Features include visualizing multiple simultaneous relations (on a single graph), valued relations (in which, say, the strength of links is visualized by varying link thickness), various node attributes (by using colors, size, labels, etc.), and selective partitioning according to cliques or other user-defined class members. Built-in analysis is limited but includes identification of isolates, components, *k*-cores, cut-points and blocks.

The current version of *NetDraw* includes two graph layout algorithms: (1) *circular*, and (2) *spring embedding* (which is based on geodesic distance and includes options for exaggerating clustering, biasing toward equal-length edges, and turning on/off node-repulsion).

<http://www.analytictech.com/downloadnd.htm>

NetMiner

NetMiner is a powerful commercial software tool for exploratory network data analysis and visualization. Developed by Cyram Co., Ltd, *NetMiner* offers a robust exploratory data analysis system that combines a suite of social network analysis tools and graph drawing techniques:

<http://www.netminer.com/NetMiner>

Pajek

*Pajek*⁹¹ is a program analyzing and visualizing large networks containing on the order of ten or hundred of thousands of nodes. Implemented in *Delphi*, it runs in the *Windows* environment and is being developed by Vladimir Batagelj and Andrej Mrvar.

Pajek supports all conventional graphs (including directed, undirected, and mixed), as well as bipartite (valued) and dynamic graphs. *Pajek* also includes algorithms for simplifications and transformations (deleting loops, multiple edges, transforming arcs to edges etc.), calculating components (strong, weak, biconnected, symmetric), decompositions (symmetric-acyclic, hierarchical clustering), paths (shortest paths, and all paths between two vertices, critical paths), flows (such as the maximum flow between two vertices), and many social networks algorithms (such as centrality measures, hubs and authorities, measures of prestige, brokerage roles, and structural holes).

Pajek version 1.02 is available at its home page (and is free, for non-commercial use):

<http://vlado.fmf.uni-lj.si/pub/networks/pajek/>

PIGALE

PIGALE (Public Implementation of a Graph Algorithm Library and Editor) is a graph editor and a C++ algorithm library for planar graphs. It is developed by H. de Fraysseix and P. Ossona de Mendez. It is available under the terms of the *Free Software Foundation's* GNU general public license:⁹²

<http://pigale.sourceforge.net/index.html>

91. “Pajek” means “spider” in the Slovenian language.

92. *GNU Project* homepage: www.gnu.org/

SNA/RSCE

SNA (*Social Network Analysis*) is a fully documented collection of RSCE (*R Statistical Computing Environment*) routines for performing social network analysis. Utilities that are included range from hierarchical Bayesian modeling, plotting and transforming networks, along with various centrality and distance measures.

SNA is provided under the terms of the GNU general public license:

<http://legba.casos.ri.cmu.edu/R.stuff/>

RSCE is a language and programming environment for statistical computing and graphics; it is available for free under the GNU general public license, in source code form:

<http://www.r-project.org/>

UCINET

UCINET is a comprehensive commercial package for analyzing social networks with up to 32K nodes. Social network analysis methods include centrality measures, subgroup identification, role analysis, elementary graph theory, and permutation-based statistical analysis. UCINET also has strong matrix analysis routines, such as matrix algebra and multivariate statistics. NetDraw (see description above) is integrated with UCINET; UCINET can also export data to *Pajek*.

UCINET version 6.29 is available at its home page:

<http://www.analytictech.com/ucinet.htm>

Tom Sawyer Software

Tom Sawyer *Graph Analysis* (C++ Edition) and *Graph Visualization* (MFC Edition) toolkits provide a fully customizable and extensible set of class libraries with APIs for developing sophisticated graph and social network analysis applications. Other packages include graph layout components for Java and Linux applications:

<http://www.tomsawyer.com/home/>

Visone

Visone (*Visual Analysis of Social Networks*) is the product of a long-term research project⁹³ to develop models and algorithms to integrate the analysis and visualization of social networks. Visone includes an interactive graphical user interface, and—although it is tailored specifically to social networks (as witnessed by the fact that its suite of metrics focuses on social network measures)—supports the import and export of standard formats for general network data. Export is in publication-quality SVG (SCalable Vector Graphics), Postscript, and other formats.

Visone version 1.1 is available for Windows, Linux, and Mac OS X:

<http://www.visone.de/download/index.html>

INSNA is the professional association for researchers interested in social network analysis. The association is a non-profit organization incorporated in the state of Delaware. Founded by Barry Wellman in 1978, the current president is Bill Richards. A copy of its charter and by-laws are available here.

Finally, the following two WWW URL links provide a miscellany of additional resources focused on *complex network theory* and *social network analysis*, respectively:

Self-Organized Networks

This research site at the *University of Notre Dame* is maintained by Albert-László Barabási and Hawoong Jeong. It contains an extensive bibliography on books related to the study of complex network theory, a gallery of networks, links to on-line peer-reviewed research

93. Visone is developed by the *Algorithms & Data Structures Group* in the *Department of Computer & Information Science*, and the *Domestic Politics & Public Administration Group* in the *Department of Politics & Management*, both at the *University of Konstanz* and a network of collaborations, with members in several different universities.

papers, lecture slides, software for analysis and visualization, as well as links to other research groups and conferences.

<http://www.nd.edu/~networks/>

International Network for Social Network Analysis

The *International Network for Social Network Analysis* (INSNA) is a non-profit professional association for researchers in social network analysis. INSNA publishes *Connections* (which is an on-line periodical containing news, research articles, technical columns, and book reviews), sponsors the annual *International Social Networks Conference*, maintains various electronic services (such as the website listed below, and the discussion, SOCNET), and provides a portal to the *Journal of Social Networks*, published by Elsevier:

<http://www.insna.org/>

Appendix 4: World Wide Web URL links to resources related to terrorism, nonlinearity and complex adaptive systems

Shortly following the terrorist attacks on September 11, 2001, the author posted the first installment of a web page (updated monthly) containing resources related to *terrorism, nonlinearity* and *complex adaptive systems* that are available on the World Wide Web:

http://www.cna.org/isaac/terrorism_and_cas.htm

The following is a short extract of the resources listed on this page:

1. *Assessing Threats of Targeted Group Violence: Contributions from Social Psychology*, by Marisa Reddy Pyncheon and Randy Borum:

http://www.ustreas.gov/usss/ntac_pyncheon.pdf

2. *Complex Challenges: Global Terrorist Networks*, on-line weekly digest, edited by Gottfried Mayer-Kress:

<http://www.comdig.org/>

3. *Complex Societies: The Evolutionary Origins of a Crude Superorganism*, P. J. Richerson, University of California Davis:

[http://www.sscnet.ucla.edu/anthro/faculty/boyd/CrudeSuper/Complex for Human Nature III clean.htm](http://www.sscnet.ucla.edu/anthro/faculty/boyd/CrudeSuper/Complex%20for%20Human%20Nature%20III%20clean.htm)

4. *Complexity Targeting. A Complexity Based Theory of Targeting and its Application to Radical Islamic Terrorism*, K.B. Glenn:

www.au.af.mil/au/awc/awcgate/saas/glenn.pdf

5. *Complexity Theory and Al-Qaeda: Examining Complex Leadership*, R. Marion and M. U.-Bien:
http://www.isce.edu/site/Marion_Uhl-Bien.pdf
6. *Countering the New Terrorism*, I. O. Lesser, B. Hoffman, J. Arquilla, D. F. Ronfeldt, M. Zanini, and B. M. Jenkins, RAND Corporation:
<http://www.rand.org/publications/MR/MR989/>
7. *Cultural Barriers to Multinational C2 Decision Making*, by Helen Altman Klein, Anna Pongonis and Gary Klein:
http://www.dodccrp.org/events/2000/CCRTS_Monterey/cd/html/pdf_papers/Track_4/101.pdf
8. *Ethnicity and Self-Organization*, by Nils Zurawski:
<http://www.uni-muenster.de/PeaCon/zurawski/6.html>
9. *Formation of Economic and Social Networks*:
<http://www.econ.iastate.edu/tesfatsi/net-group.htm>
10. *Governance Under Fire: Organizational Fragility in Complex Systems*, by Louise K. Comfort:
http://www.maxwell.syr.edu/campbell/Governance_Symposium/comfort.pdf
11. *How Emotions and Personality Effect the Utility of Alternative Decisions: A Terrorist Target Selection Case Study*, by M. Johns, B. G. Silberman:
www.seas.upenn.edu/~barryg/emotion.pdf
12. *Identifying Potential Ethnic Conflict: Application of a Process Model*, T. S. Szayna, RAND Corporation:
<http://www.rand.org/publications/MR/MR1188/>

13. *International Network for Social Network Analysis:*

<http://www.sfu.ca/~insna/>

14. *International Organization Networks: A Complementary Perspective*,
A. Judge, Union of International Association:

<http://www.uia.org/organiz/ionet77.htm>

15. *Islam's War Against the West*, by H. Bloom:

<http://www.howardbloom.net/islam.htm>

16. *Lessons of the Virus*, by E. Nolin:

www.clickz.com/experts/archives/res/personal/print.php/895351

17. *Mapping Networks of Terrorist Cells*, by Valdis E. Krebs:

www.firstmonday.org/issues/issue7_4/krebs/

18. *Modeling and Simulating Terrorist Decisionmaking: A Performance Moderator Function Approach to Generating Virtual Opponents*, by B. Silverman:

www.seas.upenn.edu:8080/~barryg/terrorist.pdf

19. *Modeling Civil Violence: An Agent-Based Computational Approach*, J. M. Epstein, J.D. Steinbruner, and M.T. Parker, Brookings Institute:

<http://www.brook.edu/es/dynamics/papers/cviolence/cviolence.pdf>

20. *Modeling Terrorist Networks - Complex Systems and First Principles of Counter-Intelligence*, P.V. Fellman, D. Sawyer, and R. Wright:

http://www.snhu.edu/img/assets/3655/Modeling_Terrorist_Networks_Fellman_Sawyer_and_Wright.doc

21. *Modelling social systems as complex: Towards a social simulation meta-model*, C. Goldspink, Journal of Artificial Societies and Social Simulation, vol. 3, no. 2, 2000:

<http://jasss.soc.surrey.ac.uk/3/2/1.html>

22. *Modeling Terrorism and Complex Adaptive Systems*, Workshop, Santa Fe Institute:

<http://discuss.santafe.edu/terrorism/>

23. *Networks, Netwar & Information-Age Terrorism*, John Arquilla, David Ronfeldt and Michelle Zanini, RAND:

www.firstmonday.dk/issues/issue6_10/ronfeldt/

24. *Networks: Structure, Dynamics & Function*, Conference, Santa Fe:

<http://cnls.lanl.gov/networks/>

25. *Making the Nation Safer: The Role of Science and Technology in Countering Terrorism*, Committee on Science and Technology for Countering Terrorism, Division on Engineering and Physical Sciences, National Research Council, National Academy Press:

<http://www.nap.edu/html/stct/>

(Chapter 10: *Complex and Interdependent Systems*)

<http://www.nap.edu/html/stct/287-312.pdf>

26. *Robust Communication Dynamics in Complex Networks*, Workshop:

<http://www.research.ibm.com/nips03workshop/>

27. *Sentinel Threat Management System*, open source data network tracking system:

<http://www.trackingthethreat.com/>

References

- [1] S. Strasser, editor, *The 9/11 Investigations: Staff Reports of the 9/11 Commission*, Public Affairs, 2004
- [2] National Commission on Terrorist Attacks, *The 9/11 Commission Report: Final Report of the National Commission on Terrorist Attacks Upon the United States*, W.W. Norton & Company, 2004.
- [3] Valdis Krebs, "Uncloaking Terrorist Networks," *First Monday* (Peer Reviewed Internet Journal), Issue 7, Number 4, 2001, http://www.firstmonday.dk/issues/issue7_4/krebs/.
- [4] Thomas A. Stewart, "Six Degrees of Mohammed Atta," *Business 2.0*, December 2001, <http://www.business2.com/articles/mag/0,1640,35253,FF.html>.
- [5] S. A. Auyang, *Foundations of Complex-system Theories: In Economics, Evolutionary Biology, and Statistical Physics*, Cambridge University Press, 1999.
- [6] Y. Bar-Yam, *Dynamics of Complex Systems*, Westview Press, 2003.
- [7] K. Mainzer, *Thinking in Complexity: The Computational Dynamics of Matter, Mind, and Mankind*, 4th Edition, Springer-Verlag, 2003.
- [8] J. Ferber, *Multi-Agent Systems: An Introduction to Distributed Artificial Intelligence*, Addison-Wesley, 1999.
- [9] G. Weiss, editor, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, MIT Press, 2000.
- [10] M. Wooldridge, *Introduction to MultiAgent Systems*, John Wiley & Sons, 2002.

- [11] A. Ilachinski, *Artificial War: Multiagent-Based Modeling of Combat*, World Scientific Publishing, 2004.
- [12] Lt.Col. M.F. Beech, "Observing Al Qaeda Through the Lens of Complexity Theory: Recommendations for the National Strategy to Defeat Terrorism," *Journal for the Center for Strategic Leadership*, Volume S04-01, 2004: <http://carlisle-www.army.mil/usacsl/publications/S04-01.pdf>.
- [13] W. F. Wechsler, L. S. Wolosky, and M. R. Greenberg, editors, *Terrorist Financing: Report of Independent Task Force*, Council on Foreign Relations, 2002: http://www.cfr.org/pdf/Terrorist_Financing_TF.pdf.
- [14] Magnus Ranstorp, deputy director of the *Centre for the Study of Terrorism and Political Violence* at the University of St. Andrews in Scotland, quoted in "Danger Persists After Hobbling Of Al Qaeda," by Dan Eggen and Michael Dobbs, Washington Post, January 14, 2002.
- [15] J. Arquilla and D. Ronfeldt, *Networks and Netwars: The Future of Terror, Crime, and Militancy*, RAND Corporation, 2001.
- [16] A. Ilachinski, *EINSTein: An Artificial-Life Laboratory for Exploring Self-Organized Emergence in Land Combat*, Center for Naval Analyses Research Memorandum CRM 2239.A1, 2000.
- [17] Peter P. Perla, A. Ilachinski, C.M.Hawk, M.C.Markowitz, and C.A. Weuve, *Using Gaming and Agent Technology to Explore Joint Command and Control Issues*, Center for Naval Analyses Research Memorandum, CRM D0007164, Unclassified, October 2002.
- [18] Philip Vos Fellman and Roxana Wright, *Modeling Terrorist Networks: Complex Systems at the Mid-Range*, Complexity, Ethics and Creativity Conference, London School of Economics, September 2003: <http://www.psych.lse.ac.uk/complexity/Conference/FellmanWright.pdf>.
- [19] R. Goolsby, *Combating Terrorist Networks*, Briefing Slides, 8th International Command and Control Research and Technology Sym-

- posium* (ICCRTS), June 2003: www.dodccrp.org/events/2003/8th_ICCRTS/Pres/track_5/1_1430goalsby.pdf.
- [20] K. Kaplan, "The Sims Take-on Al-Qaeda, *LA Times*, Nov 2, 2001, on-line: <http://www.movesinstitute.org/Press/LATimes-Sims>.
 - [21] R. Marion and M. Uhl-Bien, "Complexity Theory and Al-Qaeda: Examining Complex Leadership," Paper Presented at *Managing the Complex IV: A Conference on Complex Systems and the Management of Organizations*, Fort Meyers, FL, December, 2002: http://www.isce.edu/site/Marion_Uhl-Bien.pdf.
 - [22] R. Marion and M. Uhl-Bien, "Complexity Theory and Al-Qaeda: Examining Complex Leadership," *Emergence*, Volume 5, 2003.
 - [23] C. Mesjasz, "How Complex Systems Studies Could Help in Identification of Threats of Terrorism?" , Paper presented at *International Conference on Complex Systems (ICCS)*, Nashua, NH, June 9-14, 2002: <http://kpz.ae.krakow.pl/mesjasz/necsi2002.html>.
 - [24] J. Raab and H. Milward, "Dark Networks as Problems," *Journal of Public Administration Research and Theory*, Volume 13, 2003: <http://jpart.oupjournals.org/cgi/reprint/13/4/413.pdf>.
 - [25] M. Sageman, *Understanding Terror Networks*, University of Pennsylvania Press, 2004.
 - [26] H. Arrow, J.E. McGrath, and J. L. Berdahl, *Small Groups as Complex Systems: Formation, Coordination, Development, and Adaptation*, SAGE Publications, 2000.
 - [27] James N. Rosenau, *Distant Proximities Dynamics beyond Globalization*, Princeton Univ Press, 2003.
 - [28] A. Ilachinski, *Multiagent-Based Synthetic Warfare: Toward Developing a General Axiological Ontology of Complex Adaptive Systems*, Center for Naval Analyses Research Memorandum CRM 7376.A1, 2003.

- [29] A. Ilachinski, *EINSTein: Release Version 1.1*, Center for Naval Analyses, Multimedia CD-rom, D0007498.A1, February 2003.
- [30] Claude Berge, *The Theory of Graphs*, Dover Publications, 2001.
- [31] E. P. MacKerrow, "Understanding Why: Dissecting Radical Islamist Terrorism with Agent-Based Simulation," *Los Alamos Science*, Number 28, 2003.
- [32] W.D. Casebeer and T. Thomas, "Violent Non-State Actors: Countering Dynamic Systems," *Strategic Insights*, Volume 3, Issue 3, 2004: (<http://www.ccc.nps.navy.mil/si/2004/mar/casebeerMar04.asp>).
- [33] W.D. Casebeer and T. Thomas, "Violent Systems: Defeating Terrorists, Insurgents, and Other Non-State Adversaries," *Institute for National Security Studies*, Occasional Paper 52, 2004: <http://www.usafa.af.mil/inss/OCP/OCP52.pdf>.
- [34] Paté-Cornell, M.E. and S.D. Guikema, "Probabilistic Modeling or Terrorist Threats: A Systems Analysis Approach to Setting Priorities Among Countermeasures," *Military Operations Research*, Volume 7, No. 4, 2002.
- [35] Kathleen M. Carley, Ju-Sung Lee and David Krackhardt, Destabilizing Networks, *Connections*, Volume 24, Number 3, 2001.
- [36] Stanislaw Raczynski, "Simulation of The Dynamic Interactions Between Terror and Anti-Terror Organizational Structures," *Journal of Artificial Societies and Social Simulation*, Volume 7, Number 2, 2004.
- [37] S. Borgatti and R. Cross, "A social network view of organizational learning: relational and structural dimensions of know-who," *Management Science*, Volume 49, 2003.
- [38] Rob Cross, Andrew Parker, and Robert L. Cross, *The Hidden Power of Social Networks: Understanding How Work Really Gets Done in Organizations*, Harvard Business School Press, 2004.

- [39] E. Rogers, *Diffusion of Innovations*, Free Press, 1995.
- [40] T. W. Malone, *The Future of Work: How the New Order of Business Will Shape Your Organization, Your Management Style and Your Life*, Harvard Business School Press, 2004.
- [41] S. Wasserman and K. Faust, *Social Network Analysis*, Cambridge University Press, 1994.
- [42] J. Scott, J, *Social Network Analysis: A Handbook*, Sage Publications, 1992.
- [43] W. Buckley, *Society: A Complex Adaptive System: Essays in Social Theory*, Taylor & Francis, 1998.
- [44] N. Luhmann, *Social Systems*, Stanford University Press, 1995.
- [45] I. de Sola Pool and M. Kochen, "Contacts and influence," *Social Networks*, Volume 1, No. 5, 1978.
- [46] M. E. J. Newman, "The structure and function of complex networks," *SIAM Review*, 2003.
- [47] M. Buchanan, *Nexus: Small Worlds and the Groundbreaking Science of Networks*, W.W. Norton, 2002.
- [48] Albert-Laszlo Barabasi, *Linked: How Everything Is Connected to Everything Else and What It Means*, Plume, 2003.
- [49] P. L. Krapivsky and S. Redner, "Organization of Growing Random Networks," *Phys. Rev. E*, Volume 63, 2001.
- [50] S. Malcolm and J. Goodship, *Genotype To Phenotype: Human Molecular Genetics*, BIOS Scientific Publishers, 2001.
- [51] C. G. Langton, editor, *Artificial Life*, MIT Press, 1997.
- [52] *Human Genome Project*: www.ornl.gov/sci/techresources/Human_Genome/
- [53] M. Eigen, "Macromolecular evolution: dynamical ordering in sequence space," pages 25-69 in [54].

- [54] M. Huynen and P. Hogeweg, "Pattern generation in molecular evolution; exploitation of the variation in RNA landscapes," *Journal of Molecular Evolution*, 1993.
- [55] S. Kauffman, *The Origins of Order: Self-Organization and Selection in Evolution*, Oxford University Press, 1993.
- [56] Stefan Bornholdt and Heinz Georg Schuster, editors, *Handbook of Graphs and Networks: From the Genome to the Internet*, Wiley-VCH, 2003.
- [57] M. Dodge and R. Kitchin, *Atlas of Cyberspace*, Pearson Education, 2002.
- [58] H. Bai-lin, editor, *Chaos*, World Scientific, 1984.
- [59] P. Cvitanovic, editor, *Universality in Chaos*, Adam Hilger, 1984.
- [60] H.G. Schuster, *Deterministic Chaos: An Introduction*, Second Edition, VCH Publishers, 1988.
- [61] G.A. Cowan, D.Pines and D.Meltzer, *Complexity: Metaphors, Models and Reality*, Addison-Wesley, 1994.
- [62] D.L. Stein, editor, *Lectures in the Sciences of Complexity*, Addison-Wesley, 1989.
- [63] M. Waldrop, *Complexity: The Emerging Science at the Edge of Order and Chaos*, New York: Simon and Schuster, 1992.
- [64] Serguei N. Dorogovtsev and Jose Fernando Ferreira Mendes, "Evolution of Networks," *Adv. Physics*, Volume 51, 2002:http://merlin.fae.ua.es/fvega/review_md.pdf.
- [65] Ronald Brelger, Kathleen Carley, and Philippa Pattison, editors, *Dynamic Social Network Modeling and Analysis: Workshop Summary and Papers*, National Academy Press, 2004.
- [66] R. Goolsby, "Combating Terrorist Networks: Current Research in Social Network Analysis for the New Warfighting Environment," *8th International Command and Control Research*

- and Technology Symposium (ICCRTS)*, June 2003: http://www.dodccrp.org/events/2003/8th_ICCRTS/pdf/044.pdf.
- [67] T. Blass, *The Man Who Shocked the World: The Life and Legacy of Stanley Milgram*, Basic Books, 2004.
 - [68] S. Milgram, "The small world problem," *Psychology Today*, Volume 2, 1967.
 - [69] F. Karinthy, "Chains," in *Everything is Different*, Budapest, 1929.
 - [70] P. Killworth and H. Bernard, "The reverse small world experiment," *Social Networks*, Volume 1, 1978.
 - [71] Duncan J. Watts, *Small Worlds: The Dynamics of Networks between Order and Randomness*, Princeton Univ Press, 2003.
 - [72] Duncan J. Watts, *Six Degrees: The Science of a Connected Age*, W.W. Norton & Company, 2003.
 - [73] J. Guare, *Six Degrees of Separation: A Play*, Vintage Press, 1990.
 - [74] Duncan J. Watts, 1999. "Networks, Dynamics, and the Small-World Phenomenon," *American Journal of Sociology*, Volume 13, Number 2, pp. 493-527.
 - [75] Stefan Wuchtya, and Peter F. Stadler, "Centers of complex networks," *Journal of Theoretical Biology*, Volume 23, 2003.
 - [76] A.L. Barabási, Z. Deszo, E. Ravasz, S. H. Yook, and Z. Oltvai, "Scale-free and hierarchical structures in complex networks," <http://www.nd.edu/~networks/PDF/Proceeding%20Sitges2004.pdf>.
 - [77] S. Strogatz, "Exploring complex networks," *Nature*, Volume 410, 2001.
 - [78] J. Kleinberg, "The small-world phenomenon: an algorithmic perspective," *Cornell Computer Science Technical Report 1776*, October 1999.

- [79] P. Erdos and A. Renyi, "On the evolution of random graphs," *Publ. Math. Inst. Hung. Acad. Scie.*, Volume 5, 1960.
- [80] Albert-László Barabási and E. Bonabeau, "Scale-Free Networks," *Scientific American*, Volume 288, 2003.
- [81] *Inflow*. <http://www.orgnet.com>.
- [82] *Internet Mapping Project*: <http://research.lumeta.com/ches/map/gallery/index.html>.
- [83] Duncan J. Watts and S. Strogatz, "Collective dynamics of 'small-world' networks," *Nature*, Volume 393, 1998.
- [84] Bela Bollobas, *Modern Graph Theory*, Springer-Verlag, 1998.
- [85] R. Diestel, *Graph Theory*, Springer-Verlag, 2000.
- [86] F. Harary, *Graph Theory*, Perseus, Cambridge, 1995.
- [87] B. Bollobas, *Random Graphs*, Cambridge University Press; 2nd edition, 2001.
- [88] Albert-Laszlo Barabasi and R. Albert, "Statistical Mechanics of Complex Networks," *Rev. Mod. Phys.*, Volume 74, 2002: http://merlin.fae.ua.es/fvega/review_ab.pdf.
- [89] Serguei N. Dorogovtsev and Jose Fernando Ferreira Mendes, *Evolution of Networks: From Biological Nets to the Internet and WWW*, Oxford University Press, 2003; complete text available on-line: www.fyslab.hut.fi/~sdo/evolution_of_networks.pdf.
- [90] Jonathan L Gross and Jay Yellen, editors, *Handbook of Graph Theory*, CRC Press, 2003.
- [91] K. Kohh, "Molecular Interaction Map of the Mammalian Cell Cycle Control and DNA Repair Systems, Molecular Biology of the Cell, Volume 10, August 1999: <http://www.genopole-lille.fr/fr/biblio/articles/Kohn.pdf>.
- [92] K. Thulasiraman and M. Swami, *Graphs: Theory and Algorithms*, John Wiley and Sons, 1992.

- [93] D.M. Cvetkovic, M. Doob, and H. Sachs, *Spectra of Graphs: Theory and Applications*, Wiley, 1998.
- [94] A. Ilachinski, *Cellular Automata: A Discrete Universe*, World Scientific Publishing, 2001.
- [95] Sriram Pemmaraju and Steven Skiena, *Computational Discrete Mathematics: Combinatorics and Graph Theory with Mathematica*, Cambridge University Press, 2003.
- [96] *WEBSOM: Self-Organizing Maps for Internet Exploration*, Teuvo Kohonen: <http://websom.hut.fi/websom/>
- [97] Ioannis G. Tollis, Giuseppe Di Battista, Peter Eades, and Roberto Tamassia, *Graph Drawing: Algorithms for the Visualization of Graphs*, Prentice Hall, 1998.
- [98] Michael Kaufmann, Dorothea Wagner, *Drawing Graphs: Methods and Models*, Springer-Verlag, 2001.
- [99] Petra Mutzel, Michael Junger, *Graph Drawing Software*, Springer-Verlag, 2003.
- [100] J. O'Rourke, *Computational Geometry in C*, Cambridge University Press, 2001.
- [101] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman, 1979.
- [102] S. Card, et.al., *Information Visualization*, Morgan Kauffman, 1999.
- [103] *Self-organized Networks*, University of Notre Dame, homepage: <http://www.nd.edu/~networks/database/>.
- [104] *Network Visualization*, Max Planck Institute, homepage: <http://www.mpi-fg-koeln.mpg.de/~lk/netvis/tech.html>.
- [105] Linton C. Freeman, "Visualizing Social Networks," *Journal of Social Structure*, Volume 1, 2000, <http://moreno.ss.uci.edu/freeman.pdf>.

- [106] C. Mitchell, "Situational analysis and network analysis," *Connections*, 17, 1995.
- [107] R. V. Sole and S. Valverde, "Information Theory of Complex Networks," pages 169-190 in *Complex Networks*, edited by E. Ben-Naim, H. Frauenfelder, and Z. Toroczkai, *Lecture Notes in Physics*, Springer-Verlag, 2004: <http://www.santafe.edu/research/publications/workingpapers/03-11-061.pdf/>.
- [108] Albert-Laszlo Barabasi, Z. Dezs, E. Ravasz, S-H. Yook, and Z. Oltvai, "Scale-free and hierarchical structures in complex networks," Preprint (to appear in *Sitges Proceedings on Complex Networks*), 2004: <http://www.nd.edu/~networks/PDF/Proceeding%20Sitges2004.pdf>.
- [109] M. E. J. Newman, D. J. Watts, and S. H. Strogatz, "Random graph models of social networks," *Proceedings of the National Academy of Sciences*, Volume 99, 2002.
- [110] P. Ball, *Critical Mass: How One Thing Leads to Another*, Farrar, Straus and Giroux, 2004.
- [111] Michalis Faloutsos, Petros Faloutsos and Christos Faloutsos, "On Power-Law Relationships of the Internet Topology," *SIGCOMM*, 1999: <http://www.cs.cmu.edu/%7Echristos/PUBLICATIONS/sigcomm99.ps.gz>.
- [112] R. Albert, H. Jeong, and A.L.Barabasi, "The diameter of the World Wide Web," *Nature*, Volume 401, 1999.
- [113] P. Crucitti, V. Latora, M. Marchiori, and A. Rapisarda, "Efficiency of scale-free networks: error and attack tolerance," *Physica A*, Volume 320, 2003.
- [114] S. Mossa, M. Barthelemy, H. E. Stanley, and L. A. N. Amaral, "Incomplete Information and the Growth of Scale-Free Networks: The 'Cost of Information'," *Phys. Rev. Lett.*, Volume 88, 2002: <http://polymer.bu.edu/hes/articles/mbsa02.pdf>.
- [115] M. E. J. Newman and Juyong Park, "Why social networks are different from other types of networks," *Phys. Rev. E* 68, 2003.

- [116] Petter Holme and Beom Jun Kim, "Growing Scale-Free Networks with Tunable Clustering," *Physical Review E*, Volume 65, 2002: http://arxiv.org/PS_cache/cond-mat/pdf/0110/0110452.pdf.
- [117] K. Klemm and V.M. Eguiluz, "Highly clustered scale-free networks," *Phys. Rev. E*, Volume 65, 036123, 2002.
- [118] K. Klemm, V.M. Eguiluz, R. Total and M. san Miguel, "Global Culture: A Noise Induced Transition in Finite Systems," *AIP Conf. Proc.*, Volume 661, 2003.
- [119] J. Kleinberg, "Navigation in a Small World," *Nature*, Volume 406, 2000: <http://www.cs.cornell.edu/home/kleinber/nat00.pdf>.
- [120] L.A. Adamic, R. Lukose, A. Puniyani, and B.A. Huberman, "Search in power-law networks," *Physical Review E*, Volume 64, 2001: <http://www.cs.unibo.it/babaoglu/courses/cas/papers/search-in-power-law.pdf>.
- [121] L.A. Adamic, R. Lukose, and B.A. Huberman, "Local Search in Unstructured Networks," Chapter 13 in *Handbook of Graphs and Networks*, edited by S. Bornholdt and H.G. Schuster, Wiley, 2003.
- [122] D. Watts, P. Dodds and M.E.J. Newman, "Identity and search in social networks," *Science*, Volume 296, May 2002.
- [123] H. Bernard, R. Killworth, P.D. Evans, M.J. McCarthy, and G. Shelley, "Studying social relations cross-culturally," *Ethnology*, Volume 2, 1988.
- [124] A. Koestler, *Janus: A Summing Up*, Vintage Books, 1979.
- [125] M.E. Newman and D.J. Watts, *Phys. Lett. A*, Volume 263, 1999.
- [126] K. Sneppen, A. Trusina and M. Rosvall, "Hide and Seek in complex networks," *Preprint*, 2004: <http://xxx.lanl.gov/pdf/cond-mat/0407055>.

- [127] K. Sneppen and M. Rosvall, "Modeling Dynamics of Information Networks," *Physics Review Letters*, Volume 91, 2003: <http://www.arxiv.org/abs/cond-mat/0308399>.
- [128] A. Ilachinski and P. Halpern, "'Structurally Dynamic Cellular Automata", *Complex Systems*, Volume 2, 1987.
- [129] S. Majercik, *Structurally Dynamic Cellular Automata*, Masters Thesis, Computer Science, University of Southern Maine, 1994: <http://www.bowdoin.edu/~smajerci/pubs/masters.ps>.
- [130] D. O'Sullivan, "Exploring spatial process dynamics using irregular graph-based cellular automaton models," *Geographical Analysis*, Volume 33, 2001.
- [131] S. Saidani and M. Piel, "Dynagraph: a Smalltalk Environment for Self-Reconfigurable Robots Simulation," paper presented at the *European Smalltalk User Group (ESUG) Conference*, 2004 Research Track: http://cubitus.info.unicaen.fr:8080/samir/uploads/DynaGraph,_a_Smalltalk_Environment_for_Self-Reconfigurable_Robots_Simulation.pdf.
- [132] S. Saidani, "Dynamic graphs as cellular automata," *Discrete Mathematics and Theoretical Computer Science*, 2003.
- [133] D.J. Brass, "A Social Network Perspective on Human Resources Management," *Research in Personnel and Human Resources Management*, Volume 13, 1995.
- [134] R. S. Burt, *Structural Holes: The Social Structure of Competition*, Harvard University Press, 1992.
- [135] Linton C. Freeman, "Centrality in Social Networks," *Social Networks*, Volume 1, 1979.
- [136] Petter Holme, Beom Jun Kim, Chang No Yoon and Seung Kee Han, "Attack vulnerability of complex networks," *Phys. Rev. E*, Volume 65, 2002.
- [137] Petter Holme and Beom Jun Kim, "Vertex overload breakdown in evolving networks," *Physical Review E*, Volume 65,

- 2002: <http://arxiv.org/ftp/cond-mat/papers/0204/0204120.pdf>.
- [138] P. Bonacich, "Power and Centrality: a Family of Measures," *American Journal of Sociology*, Volume 92, 1987.
 - [139] R. A. Horn and C. R. Johnson, *Matrix Analysis*, Cambridge University Press, 1990.
 - [140] N. Biggs, *Algebraic Graph Theory*, Cambridge University Press, 1994.
 - [141] S. Brin and L. Page, "The anatomy of a large-scale hypertextual search engine," *Computer Networks*, Volume 30, 1998.
 - [142] K. Stephensen and M. Zelen, "Rethinking Centrality: Methods and Applications," *Social Networks*, Volume 11, 1989.
 - [143] M.E.J. Newman, "A Measure of Betweenness Centrality Based on Random Walks," e-print: <http://arxiv.org/abs/cond-mat/0309045>.
 - [144] M.A. Beauchamp, "An Improved Index of Centrality," *Behavioral Science*, Volume 10, 1965.
 - [145] G. Sabidussi, "The Centrality Index of a Graph," *Psychometrika*, Volume 31, 1966.
 - [146] M. Jaaskelainen, "Centrality Measures and Information Flows in Venture Capital Syndication Networks," 2001: <http://www.sal.tkk.fi/Opinnot/Mat-2.108/pdf-files/ejaa01.pdf>.
 - [147] L. C. Freeman, S. P. Borgatti, and D.R. White, "Centrality in valued graphs: A measure of betweenness based on network flow," *Social Networks*, Volume 13, 1991.
 - [148] Linton C. Freeman, "A Set of Measures of Centrality Based on Betweenness," *Sociometry*, Volume 40, 1977.
 - [149] V. Latora and M. Marchiori, "How the Science of Complex Networks Can Help Developing Strategies Against Terrorism," *Chaos, Solitons and Fractals*, Volume 20, 2004.

- [150] V. Latora and M. Marchiori, "A Measure of Centrality Based on the Network Efficiency," e-print: <http://arxiv.org/abs/cond-mat/0402050>.
- [151] T. Carpenter, G. Karakostas, and D. Shallcross, "Practical Issues and Algorithms for Analyzing Terrorist Networks," *Telcordia Technologies, Inc.*, Invited paper at WMC 2002: <http://www.cas.mcmaster.ca/%7Egk/papers/wmc2002.pdf>.
- [152] S. Fortunato, V. Latora, and M. Marchiori, "Method to find community structures based on information centrality," *Physical Review E*, Volume 70, 2004: <http://axpfct.ct.infn.it/~latora/flm04.pdf>.
- [153] D. R. White and S. Borgatti, "Betweenness centrality measures for directed graphs, *Social Networks*, Volume 16, 1994.
- [154] S. White and Padhraic Smyth, "Algorithms for Estimating Relative Importance in Networks," *Knowledge Discovery in Data and Data Mining*, August, 2003.
- [155] J. Bruggeman, G. Carnabuci, and I. Vermeulen, "A note on structural holes theory and niche overlap," *Social Networks*, 2003: <http://users.fmg.uva.nl/jbruggeman/socnet03.htm>.
- [156] J. Moody, D. R. White, "Social Cohesion and Embeddedness: A Hierarchical Conception of Social Groups," *Santa Fe Institute Working Papers*, 00-08-049, <http://www.santafe.edu/sfi/publications/Working-Papers/00-08-049.pdf>.
- [157] J. Moody, *Racial Friendship Segregation*, Sociogram: <http://www.sociology.ohio-state.edu/jwm/race1.gif>.
- [158] M. Sageman, Understanding Terror Networks, *Foreign Policy Research Institute*, November 2004: <http://www.fpri.org/enotes/20041101.middleeast.sageman.understandingterror-networks.html>.
- [159] M. Sageman, "Understanding Al Qaeda Networks," *Briefing Slides*, National Institute of Standards and Technology

- (NIST), http://www.bfrl.nist.gov/PSSIWG/presentations/Understanding_al_Qaeda_Networks.pdf.
- [160] E.O. Laumann, *Prestige and Association in an Urban Community*, Bobbs-Merril, 1966.
 - [161] S.C. Johnson, "Hierarchical Clustering Schemes," *Psychometrika*, Volume 2, 1967.
 - [162] R.K. Ahuja, T. Magnanti, and J. Orlin, *Network Flow: Theory, Algorithms, and Applications*, Prentice-Hall, 1993.
 - [163] M. E. J. Newman, M. Girvan, "Finding and evaluating community structure in networks, *Phys. Rev. E*, Volume 69, 2004: <http://arxiv.org/pdf/cond-mat/0308217>.
 - [164] J.R. Tyler, D.M. Wilkinson, and B.A. Huberman, "Email as spectroscopy: automated discovery of community structure within organizations," in M. Huysman, E. Wenger, and V. Wulf, editors, *Proceedings of the First International Conference on Communities and Technologies*, Dordrecht, 2003.
 - [165] F. Radicchi, C. Castellano, F. Cecconi , V. Loreto, and D. Parisi, "Defining and identifying communities in networks," *Proceedings of the National Academies of Sciences*, 2004: <http://www.pnas.org/cgi/reprint/101/9/2658>.
 - [166] G.W. Flake, S.R. Lawrence, C. Giles and F. Coetzee, "Self-organization and identification of Web communities," *IEEE Computer*, Volume 35, 2002.
 - [167] M. E. J. Newman, "Detecting community structure in networks," *Eur. Phys. Jour. B*, Volume 38, 2004: <http://www.santafe.edu/~mark/papers/epjb.pdf>.
 - [168] M. E. J. Newman, "Fast algorithm for detecting community structure in networks," *Phys. Rev. E*, Volume 69, 2004: <http://arxiv.org/pdf/cond-mat/0309045>.

- [169] D. Lusseau and M.E.J. Newman, "Identifying the role that individual animals play in their social network," *Ecology Letters*, 2004: <http://arxiv.org/pdf/q-bio.PE/0403029>.
- [170] J.L. Moreno, *Sociometry, Experimental Method and the Science of Society. An Approach to a New Political Orientation*, Beacon House, Inc., 1951.
- [171] A. Clauset, M.E. Newman, and C. Moore, "Finding community structure in very large networks," *Physical Review E*, 2004: <http://www.eece.unm.edu/ifis/papers/community-moore.pdf>.
- [172] H. Zhou, "Distance, dissimilarity index, and network community structure," *Physical Review E*, Volume 67, 2003.
- [173] M.O. Ball, "Complexity of network reliability computations," *Networks*, Volume 10, 1980.
- [174] M.O. Ball, "Computational complexity of network reliability analysis: an overview," *IEEE Transactions on Reliability*, Volume 35, 1986.
- [175] Lucet, C. and J-F. Manouvrier, "Exact Methods to Compute Network Reliability," *First International Conference on Mathematical Methods in Reliability, Bucarest, Roumania, September 1997*:<http://citeseer.ist.psu.edu/cache/papers/cs/21324/httpzSzzSzwww.hds.utc.frzSz~manouvzSzmmr.pdf/exact-methods-to-compute.pdf>.
- [176] T. Marlowe and L. Schoppmann, "Polynomial-Time Computability of the Edge-Reliability of Graphs Using Gilbert's Formula," *Mathematical Problems in Engineering*, Volume 4, 1998: <http://www.hindawi.co.uk/open-access/mpe/volume-4/S1024123X98000817.pdf>.
- [177] J. A. Buzacott, "A recursive algorithm for finding reliability measures related to the connection of nodes in a graph," *Networks*, Volume 10, 1980.

- [178] D.D. Harms, *Network Reliability: Experiments with a Symbolic Algebra Environment*, CRC Press, 1995.
- [179] M.L. Shooman, *Reliability of Computer Systems and Networks: Fault Tolerance, Analysis, and Design*, Wiley, 2001.
- [180] Hans L. Bodlaender and Thomas Wolle, *A Note on the Complexity of Network Reliability Problems*, Institute of Information and Computing Sciences, Utrecht University, Technical Report UU-CS-2004-001: <http://archive.cs.uu.nl/pub/RUU/CS/techreps/CS-2004/2004-001.pdf>.
- [181] A.E. Motter and Y.-C. Lai, "Cascade-based attacks on complex networks," *Physical Review E*, Volume 66, 2002.
- [182] A.E. Motter, T. Nishikawa, and Y.-C. Lai, "Range-based attack on links in scale-free networks: are long-range links responsible for the small-world phenomenon?," *Phys. Rev. E*, Volume 66, 2002: http://arxiv.org/PS_cache/cond-mat/pdf/0206/0206030.pdf.
- [183] Y.-C. Lai, A.E. Motter, and T. Nishikawa, "Attacks and Cascades in Complex Networks," *Lecture Notes in Physics*, Springer-Verlag, Volume 650, 2004: http://chaos1.la.asu.edu/~yclai/papers/LNP_04.pdf.
- [184] Petter Holme, "Edge overload breakdown in evolving networks," *Physical Review E*, Volume 66, 2002: http://arxiv.org/PS_cache/cond-mat/pdf/0207/0207466.pdf.
- [185] D.J. Watts, "A simple model of global cascades on random networks," *Proceedings of the National Academy of Sciences*, Volume 99, 2002.
- [186] Reka Albert, Istvan Albert, and Gary L. Nakarado, "Structural Vulnerability of the North American Power Grid," <http://arxiv.org/pdf/cond-mat/0401084>.
- [187] J.D. Farley, "Breaking Al Qaeda Cells: A Mathematical Analysis of Counterterrorism Operations (A Guide for Risk Assessment and Decision Making), *Studies in Conflict and Terrorism*,

Volume 26, 2003: <http://dimacs.rutgers.edu/Workshops/Defense/article4.pdf>.

- [188] G. Woo, "Insuring Against Al Qaeda," *National Bureau of Economic Research Conference*, 2003: <http://www.nber.org/~confer/2003/insurance03/woo.pdf>.
- [189] G. Woo, "Adaptation in an Environment of Terror," Briefing Slides, *Duke Environmental Leadership Forum*, 2002: <http://www.env.duke.edu/forum02/woo.pdf>.
- [190] I. Peterson, "Splitting Terrorist Cells," *Science News Online*, Volume 165, Number 2, January, 2004: <http://www.science-news.org/articles/20040110/mathtrek.asp>.
- [191] Wei-Chang Yeh, "A simple algorithm to search for all minimalcutsets with unreliable nodes," *Complexity International*, Volume 8, 2001: <http://journal-ci.csse.monash.edu.au/ci/vol08/yeh01/yeh01.pdf>.
- [192] G. Woo, "The Al Qaeda War Game," *Swiss Military Review*, December 2002.
- [193] Valdis Krebs, "Mapping Networks of Terrorist Cells," *Connections*, Volume 24, Number 3, 2002: <http://www.orgnet.com/MappingTerroristNetworks.pdf>.
- [194] S. Voss and C. Joslyn, "Advanced Knowledge Integration In Assessing Terrorist Threats," *Los Alamos Laboratory*, <ftp://ftp.c3.lanl.gov/pub/users/joslyn/knowint.pdf>.
- [195] Valdis Krebs, "Social Network Analysis of the 9-11 Terrorist Network," <http://www.orgnet.com/hijackers.html>.
- [196] R. Smith, "Modeling and Simulation Aids Insight on Terrorism," *Signal*, 2001.
- [197] M.W. Nance, *The Terrorist Recognition Handbook: A Manual for Predicting and Identifying Terrorist Activities*, Lyons Press, 2003.

- [198] Rohan Gunaratna, *Inside Al Qaeda: Global Network of Terror*, Berkley Pub Group, 2003.
- [199] D.R. White and F. Harary, "The Cohesiveness of Blocks in Social Networks: Node Connectivity and Conditional Density," in *Sociological Methodology 2001*, edited by M. Sobel and M. Becker, Blackwell Publishers, 2002.
- [200] S. French, *Decision Theory: An Introduction to the Mathematics of Rationality*, Ellis Horwood Limited, 1986.
- [201] Peter R. Monge and Noshir S. Contractor, "Emergence of Communication Networks," in *Handbook of Organizational Communication*, Second Edition, edited by F. M. Jablin and L.L. Putnam, Sage Press, 1999.
- [202] N. S. Contractor, R. Whitbred, F. Fonti, A. Hyatt, P. Jones, and B. O'Keefe, "Self-organizing communication networks in organizations: validation of a computational model using exogenous and endogenous theoretical mechanisms," paper presented at the meeting of the *International Communication Association*, Jerusalem, Israel, 1998: <http://www.spcomm.uiuc.edu/users/nosh/manuscripts/Comp/comp.htm>.
- [203] Peter R. Monge and Noshir S. Contractor, *Theories of Communication Networks*, Oxford University Press, 2003.
- [204] D.E. Byrne, *The Attraction Paradigm*, Academic Press, 1971.
- [205] J.C. Turner, *Rediscovering the Social Group: A Self-Categorization Theory*, Oxford University Press, 1987.
- [206] Albert-Laszlo Barabasi and R. Albert, "Emergence of scaling in random networks", *Science*, Volume 286, 1999.
- [207] H. Ibarra and S. B. Andrews, "Power, social influence, and sense making: effects of network centrality and proximity on employee perceptions," *Administrative Science Quarterly*, Volume 38, 1993.

- [208] P.R. Monge, L.W. Rothman, E.M. Eisenberg, K.I. Miller, and K.K. Kirste, "The dynamics of organizational proximity," *Management Science*, Volume 31, 1985.
- [209] R.E. Rice and C. Aydin, "Attitudes toward new organizational technology: network proximity as a mechanism for social information processing," *Administrative Science Quarterly*, Volume 9, 1991
- [210] E.J. Bienenstock and P. Bonacich, "Network exchange as a cooperative game," *Rationality and Society*, Volume 9, 1997.
- [211] K.S. Cook and T. Yamagishi, "Power in exchange networks: a power-dependence formulation," *Social Networks*, Volume 14, 1992.
- [212] R. Axelrod, *The Evolution of Cooperation*, Basic Books, 1984.
- [213] R. Axelrod, "The Evolution of Strategies in the Iterated Prisoner's Dilemma," in *Genetic Algorithms and Simulated Annealing*, edited by L. Davis, Morgan Kaufman, 1987: <http://www-personal.umich.edu/~axe/research/Evolving.pdf>.
- [214] R. Axelrod, *The Complexity of Cooperation*, Princeton University Press, 1997.
- [215] G.C. Homans, *The Human Group*, Harcourt Brace, 1950.
- [216] H.A. Simon, *Models of Man*, Wiley, 1957.
- [217] N. Lin and M. Granovetter, editors, *Social Capital : A Theory of Social Structure and Action*, Cambridge University Press, 2002.
- [218] K.W. Back, "Influence through social communication," *Journal of Abnormal and Social Psychology*, Volume 46, 1951.
- [219] G.C. Homans, *The Human Group*, Harcourt, 1950.
- [220] S. E. Seashore, "Group cohesiveness in the industrial work group," *Institute for Social Research*, 1954.

- [221] N. J. Evans, and K. L. Dion, "Group cohesion and performance: a meta-analysis," *Small Group Research*, Volume 22, 1991.
- [222] N. E. Friedkin, "Social Cohesion," *Annual Review of Sociology*, Volume 30, 2004.
- [223] Durkin, *Expert Systems: Design and Development*, Prentice Hall, 1994.
- [224] Valdis Krebs, "Surveillance of Terrorist Networks," <http://www.orgnet.com/tnet.html>.
- [225] Thomas A. Stewart, "Six Degrees of Mohammed Atta," *Business 2.0*, December, 2001, <http://www.business2.com/b2/web/articles/1,17863,514212,00.html>.

Bibliography

- [1] R. Albert and A.-L. Barbási, "Statistical Mechanics of Complex Networks," <http://www.nd.edu/~networks/Papers/review.pdf>.
- [2] *Al Qaeda Training Manual*, United States Department of Justice, on-line: http://www.au.af.mil/au/awc/awcgate/terrorism/alqaida_manual/.
- [3] L. Amaral, A. Scala, M. Barthelemy, and H. E. Stanley, "Classes of Behavior of Small-World Networks," *Proc. Natl. Acad. Sci.*, Volume 97, 2000: <http://polymer.bu.edu/hes/articles/asbs00.pdf>.
- [4] A. Arenas, A. Cabrales, A. Diaz-Guilera, R. Guimera, and F. Vega-Redondo, "Search and Congestion in Complex Networks," Proceedings of the Conference *Statistical Mechanics of Complex Networks*, Sitges, Spain, June 2002: <http://arxiv.org/pdf/cond-mat/0301124>.
- [5] R. Axelrod, *The Complexity of Cooperation*, Princeton Univ Press, 1997.
- [6] A. Balcioglu, *An Algorithm for Enumerating the Near-Minimum Weight S-T Cuts of a Graph*, Thesis, Naval Post Graduate School, December 2000: http://library.nps.navy.mil/uhtbin/hyperion-image/00Dec_Balcioglu.pdf.
- [7] A.-L. Barabasi, Albert, and H. Jeong, "Mean-field theory for scalefree random networks," *Physica A*, Volume 272, 1999.
- [8] A.-L. Barabasi, R. Albert, and Jeong H., "The Internet's Achilles' heel: error and attack tolerance of complex networks", *Nature*, Volume 406, 2000.

- [9] A.-L. Barabási, E. Ravasz and T. Vicsek, "Deterministic scale-free networks," *Physica A*, Volume 299, 2001.
- [10] E. Ben-Naim, H. Frauenfelder, and Z. Toroczkai, editors, *Complex Networks, Lecture Notes in Physics*, Springer-Verlag, 2004.
- [11] H. Bloom, *The Lucifer Principle: A Scientific Expedition into the Forces of History*, Atlantic Monthly Press, 1997.
- [12] M. Boguna, R. Pastor-Satorras, A. Diaz-Guilera, A. Arenas, "Emergence of clustering, correlations, and communities in a social network model," *Preprint*, 2003: <http://arxiv.org/pdf/cond-mat/0309263>.
- [13] P. Bonacich, "Factoring and weighting approaches to status scores and clique identification," *Journal of Mathematical Sociology*, Volume 2, 1972.
- [14] P. Bonacich, "Simultaneous Group and Individual Centralities," *Social Networks*, Volume 13, 1991.
- [15] B. G. Buchanan and E.H. Shorti, *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*, Addison-Wesley, 1984.
- [16] W. E. Combs, *The Fuzzy Systems Handbook*, Academic Press, 1999.
- [17] R. Cross, S.P. Borgatti, and A. Parker, "Making Invisible Work Visible: Using Social Network Analysis to Support Strategic Collaboration," *California Management Review*, Volume 44, 2002.
- [18] M. Diani and Doug McAdam, editors, *Social Movements and Networks: Relational Approaches to Collective Action*, Oxford University Press, 2003.
- [19] B. Dutta, editor, *Networks and Groups: Models of Strategic Formation*, Springer Verlag, 2003.

- [20] P. Eades, "A heuristic for graph drawing," *Congressus Numerantium*, Volume 42, 1984.
- [21] D. Eggan and M. Dobbs, "Danger persists after hobbling of Al Qaeda U.S. officials fear attacks by Bin Laden supporters," *Washington Post*, January 13, 2002.
- [22] J. M. Epstein and Robert L. Axtell, *Growing Artificial Societies: Social Science from the Bottom Up*, MIT Press, 1996.
- [23] J. M. Epstein, John D. Steinbruner, Miles T. Parker, "Modeling Civil Violence: An Agent-Based Computational Approach," Brookings Institute, Center on Social and Economic Dynamics, Working Paper Number 20, January 2001, <http://www.brook.edu/dybdocroot/es/dynamics/papers/cviolence/cviolence.pdf>.
- [24] B. H. Erickson, "Secret societies and social structure," *Social Forces*, Volume 60, 1981.
- [25] Saad al Fagih, Interview, *Frontline*, Public Broadcasting Service, 1999: <http://www.pbs.org/wgbh/pages/frontline/shows/binladen/interviews/al-fagih.html>.
- [26] *Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations*, U.S.-Canada Power System Outage Task Force: <https://reports.energy.gov/BlackoutFinal-Web.pdf>.
- [27] L. R. Ford and D.R. Fulkerson, *Flows in Networks*, Princeton University Press, 1962.
- [28] L. C. Freeman, "Visualizing Social Groups" *American Statistical Association, Proceedings of the Section on Statistical Graphics*, 2000, <http://moreno.ss.uci.edu/groups.pdf>.
- [29] N. E. Friedkin, "Theoretical Foundations for Centrality Measures," *AJS*, Volume 96, 1991.

- [30] R. J. Gaylord and Lou D'Andria, *Simulating Society: A Mathematica Toolkit for Modeling Socioeconomic Behavior*, Telos Press, 1998.
- [31] N. Gilbert and Klaus G Troitzsch, *Simulation for the Social Scientist*, Open University Press, 1999.
- [32] N. Gilbert and Rosaria Conte, editors, *Artificial Societies: The Computer Simulation of Social Life*, UCL Press, 1995.
- [33] M. Gladwell, *The Tipping Point: How Little Things Can Make a Big Difference*, Back Bay Books, 2002.
- [34] R. Guimera, A. Diaz-Guilera, F. Vega-Redondo, A. Cabrales, and A. Arenas, "Optimal network topologies for local search with congestion," *Phys. Rev. Lett.*, Volume 89, 2002: <http://arxiv.org/pdf/cond-mat/0206410>.
- [35] R. D. Howard and Reid Sawyer, editors, *Terrorism and Counterterrorism: Understanding the New Security Environment, Readings and Interpretations*, McGraw-Hill, 2003.
- [36] B. Hoffman, *Inside Terrorism*, Columbia University Press, 1998.
- [37] R. Hudson, *Who Becomes a Terrorist and Why: The 1999 Government Report on Profiling Terrorists*, Lyons Press, 2002.
- [38] A. Ilachinski and P. Halpern, "Structurally dynamic cellular automata," *Complex Systems*, Volume 1, 1987.
- [39] A. Ilachinski, *EINSTEIN: An Artificial-Life Laboratory for Exploring Self-Organized Emergence in Land Combat*, Center for Naval Analyses Research Memorandum CRM 2239.A1, 2000.
- [40] A. Ilachinski, *Land Warfare and Complexity, Part I: Mathematical Background and Technical Sourcebook*, Center for Naval Analyses Information Manual CIM-461, Unclassified, 1996. Available on-line, in Adobe's Acrobat format, at URL address <http://www.cna.org/isaac/lw1.pdf>.
- [41] A. Ilachinski, *Land Warfare and Complexity, Part II: An Assessment of the Applicability of Nonlinear Dynamics and Complex Sys-*

- tems Theory to the Study of Land Warfare*, Center for Naval Analyses Research Memorandum CRM-68, Unclassified, 1996. Can be downloaded from the WWW at URL address <http://www.cna.org/isaac/lwpart2.pdf>.
- [42] A. Ilachinski, *A Mobile Cellular Automata Approach to Land Combat*, Center for Naval Analyses Information Manual CIM-482, Unclassified, 1996.
 - [43] A. Ilachinski, *Irreducible Semi-Autonomous Adaptive Combat (ISAAC): An Artificial-Life Approach to Land Warfare*, Center for Naval Analyses Research Memorandum CRM 97-61, 1997. Can be downloaded from the WWW at URL address <http://www.cna.org/isaac/crm9761.htm>.
 - [44] A. Ilachinski, "Irreducible Semi-Autonomous Adaptive Combat (ISAAC): An Artificial-Life Approach to Land Warfare," *Military Operations Research*, Volume 5, Number 3, 2000.
 - [45] A. Ilachinski, *EINStein: Version 1.0.0.4b*, Center for Naval Analyses, Multimedia CD-rom, D0003394.A1, February 2001.
 - [46] A. Ilachinski, *Multi-Agent-Based Synthetic Warfare: Towards Developing A General Axiological Ontology of Complex Adaptive Systems*, Center for Naval Analyses Research Memorandum CRM D0007376.A1, 2002.
 - [47] A. Ilachinski, *EINStein: Release Version 1.1*, Center for Naval Analyses, Multimedia CD-rom, D0007498.A1, February 2003. C. Isbell, M. Kearns, D. Kormann, S. Singh, P. Stone, "Cobot in LambdaMOO: A Social Statistics Agent," *AAAI*, 2000, <http://www.cc.gatech.edu/fac/Charles.Isbell/projects/cobot/>.
 - [48] H. Inose, "Communication networks," *Scientific American*, March, 1972.
 - [49] M. Johns and Barry Silverman, "How Emotions and Personality Affect the Utility of Alternative Decisions: A Terrorist Target Selection Case Study," <http://www.seas.upenn.edu/~barryg/emotion.pdf>.

- [50] K. Klemm and V.M. Eguiluz, "Growing scale-free networks with small-world behavior," *Phys. Rev. E*, Volume 65, 057102, 2002.
- [51] P. Klerks, *The Network Paradigm Applied to Criminal Organisations: Theoretical nitpicking or a relevant doctrine for investigators? Recent developments in the Netherlands*, Connections, Volume 24, 2001: <http://www.sfu.ca/~insna/Connections-Web/Volume24-3/klerks.pdf>.
- [52] T. Kohonen, T.S. Huang, and M.R. Schroeder, *Self-Organizing Maps*, Springer-Verlag, 3rd edition, 2000.
- [53] V. Krebs, "Surveillance of Terrorist Networks," <http://www.orgnet.com/tnet.html>.
- [54] L. Krempel and Thomas Plümper, "Exploring the Dynamics of International Trade by Combining the Comparative Advantages of Multivariate Statistics and Network Visualizations," *Journal of Social Structure*, Volume 4, 2002: <http://www.cmu.edu/joss/content/articles/volume4/Krempel-Plumper.html>.
- [55] A. Lomi and Erik R. Larsen, editors, *Dynamics of Organizations: Computational Modeling and Organizational Theories*, AAAI Press, 2001.
- [56] M. J. Mataric, "Designing and understanding adaptive group behavior," *Adaptive Behavior*, Volume 4, No. 1, December 1995,
- [57] H. McCarthy, P. Miller, and P. Skidmore, editors, *Network Logic: Who governs in an interconnected world?*, Demos, 2004: www.demos.co.uk/networklogic_pdf_media_public.aspx.
- [58] C. Mitchell, "Situational analysis and network analysis," *Connections*, 17, 1995.
- [59] PAJEK, Visualization and Analysis Package for Large Networks, <http://vlado.fmf.uni-lj.si/pub/networks/pajek/>.

- [60] P. P. Perla, A. Ilachinski, C.M.Hawk, M.C.Markowitz, and C.A. Weuve, *Using Gaming and Agent Technology to Explore Joint Command and Control Issues*, Center for Naval Analyses Research Memorandum, CRM D0007164, Unclassified, October 2002.
- [61] M. Prietula, Kathleen Carley, and Les Gasser, editors, *Simulating Organizations: Computational Models of Institutions and Groups*, AAAI Press, 1998.
- [62] M. R. Pynchon and Randy Borum, Assessing Threats of Targeted Group Violence: Contributions from Social Psychology, *Behavioral Sciences and the Law*, Volume 17, 1999.
- [63] W. Reich, editor, *Origins of Terrorism: Psychologies, Ideologies, Theologies, States of Mind*, Woodrow Wilson Center Press, 1998.
- [64] R. Renfro and R. Deckro, "A Social Network Analysis of the Iranian Government," paper presented at 69th MORS Symposium, Working Group 8, 2001 (<http://www.fas.org/irp/eprint/socnet.pdf>).
- [65] D. Ronfeldt and John Arquilla, "Networks, Netwars, and the Fight for the Future," *First Monday* (Peer Reviewed Internet Journal), Volume 6, Number 10, 2001, http://firstmonday.org/issues/issue6_10/ronfeldt/.
- [66] R. Rothenberg, "From Whole Cloth: Making up the Terrorist Network," *Connections*, Volume 24, 2002.
- [67] E. Rothstein, "A Lethal Web With No Spider," *New York Times*, October 20, 2001.
- [68] I. Sanders, "To Fight Terror, We Can't Think Straight," *Washington Post*, May 5, 2002.
- [69] T. C. Schelling, *Micromotives and Macrobehavior*, W.W. Norton & Company, 1978.
- [70] T. Siegfried, "Network science could provide patriot gains," *Dallas Morning News*, Sep 9, 2002.

- [71] H. A. Simon, *The Sciences of the Artificial*, MIT Press, 1996.
- [72] *Social Network References*: <http://www.socialnetworks.org/>.
- [73] M. Stella, editor, Complexity and Critical Infrastructure Vulnerabilities, Proceedings of a Workshop Sponsored by The Cyber Conflict Studies Association and The Center for Technology and National Security Policy, National Defense University, 2004: www.ndu.edu/ctnsp/complexity_book.htm.
- [74] T. A. Stewart, "Six Degrees of Mohammed Atta," *Business 2.0*, December, 2001, <http://www.business2.com/b2/web/articles/1,17863,514212,00.html>.
- [75] S. Strogatz, *Sync: The Emerging Science of Spontaneous Order*, Hyperion, 2003.
- [76] K. Supowit and E. Reingold, "The complexity of drawing trees nicely," *Acta Informatica*, Volume 18, 1983.
- [77] J. Surowiecki, *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*, Doubleday Books, 2004.
- [78] B. Tadic and V. Priezzhev, Voltage distribution in growing conducting networks," *Eur. Phys. Jour. B*, Volume 30, 2002.
- [79] United States Department of Defense, *Transcript of bin Laden Video Tape*, December 13, 2001, <http://www.defenselink.mil/news/Dec2001/d20011213ubl.pdf>.
- [80] Washington Post, 2001. "The Plot: A Web of Connections," 24 September 2001, http://www.washingtonpost.com/wp-srv/nation/graphics/attack/investigation_24.html.
- [81] G. Weimann, "How Modern Terrorism Uses the Internet," Special Report, *US Institute of Peace*, <http://www.usip.org/pubs/specialreports/sr116.pdf>.
- [82] M. B. West (Major), USMC, *From Metaphors to Models: Broadening the Lens of the Hunter Warrior Experiment with a Complex Adap-*

tive System Tool, Thesis, AY 98-99, Marine Corps Combat Development Command, Quantico, VA.

- [83] S. White and P. Smyth, "Algorithms for Discovering Relative Importance In Graphs," *Proceedings of Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Washington D.C.*, 2003: http://www.ics.uci.edu/~scott/rel_auth.pdf.
- [84] G. Woo, "The Art of Terror," *Risk Transfer*, September 2002.
- [85] G. Woo, "Quantifying Insurance Terrorism Risk," *National Bureau of Economic Research meeting, Massachusetts*, 2002.
- [86] M. Wooldridge, *Introduction to MultiAgent Systems*, John Wiley & Sons, 2002.
- [87] B. Yu, M. Venkatraman, and M.P. Singh, "An adaptive social network for information access: Theoretical and experimental results," *Applied Artificial Intelligence*, 2000.

List of figures

Figure 1.	A schematic overview of some of SOTCAC's components discussed later in this paper (T=Terrorist, CT=Counterterrorist)	4
Figure 2.	Social network of ties among 9/11 hijackers	6
Figure 3.	Schematic of how mathematical graphs can be used to capture arbitrary relationships among objects, and serve as conceptual anchors of multiagent-based models	11
Figure 4.	Schematic illustration of SOTCAC's coupled information and physical spaces; in contrast, the dynamics of the EINSTEIN combat model are confined solely to the physical domain.	14
Figure 5.	Examples of large complex networks	30
Figure 6.	A small portion of the molecular interaction map for the regulatory network responsible for mammalian cell cycles.	33
Figure 7.	Graph consisting of eight nodes and ten links	34
Figure 8.	All undirected, unlabeled graphs G of order 4	36
Figure 9.	Sample renderings of the same (order 25 and size 49) random graph using the five visualization algorithms discussed in the text	43
Figure 10.	An example of graph visualization using spring-embedding	44
Figure 11.	Schematic of the spectrum of all possible graphs, $G(N,M)$, of size N and order M	47
Figure 12.	Zoology of graphs	49

Figure 13. A schematic illustration of several important epochs during the evolution of a random graph.	54
Figure 14. An example Watts-Strogatz algorithm	56
Figure 15. Typical decays of characteristic path length and clustering coefficient using the Watts-Strogatz small-world random graph model.	57
Figure 16. Schematic illustration of a single chain in Watts, <i>et al.</i> 's hierarchical social network model	69
Figure 17. Comparison between the number of message chains of length L , as observed in Milgrim's original "small-worlds" experiment	71
Figure 18. Network schematic for calculating search information, target entropy and road entropy.	73
Figure 19. Schematic of partitioning a local neighborhood into three disjoint sets	78
Figure 20. Sample step of applying an SDCA transition rule to a 5-by-5 lattice of nodes	79
Figure 21. Sample SDCA evolution.	80
Figure 22. An example of how to calculate the local clustering coefficient.	87
Figure 23. Sample star, circle and line graphs used for comparing centrality metrics	89
Figure 24. Sample calculations of global and local link densities	91
Figure 25. Sample weighted graph	99
Figure 26. Illustration of two similar local topologies	105
Figure 27. An illustration of a network with "community structure"	108

Figure 28. Spring-embedding visualization	109
Figure 29. Schematic of the global Salafi network.	110
Figure 30. Illustration of the effects of node deletion on an initially connected network	116
Figure 31. Hypothetical terrorist cell.	123
Figure 32. Schematic of the decision-making process in EINSTEIN.	129
Figure 33. Schematic of the decision-making process in SOTCAC	132
Figure 34. Schematic of SOTCAC’s physical and information space representations of TN and CTN coevolutions	134
Figure 35. Schematic timeline of terrorist network and counterterrorist network coevolution	139
Figure 36. Schematic illustration of the basic components of a terrorist network	143
Figure 37. Schematic illustration of σ ’s local topology map at time t	151
Figure 38. Schematic illustration of the four basic categories of contacts in σ ’s ego-map	153
Figure 39. Schematic illustration of the different distance functions	157
Figure 40. EINSTEIN’s action selection.	173
Figure 41. Schematic illustration of an agent s ’s local neighborhood in social network	185
Figure 42. Schematic of link-creation and link-deletion in SOTCAC	187
Figure 43. Schematic illustration of how the CTN “infers” latent TN-structure using raw data filtered by a CT-agent, using method 1	192

Figure 44. Schematic illustration of how the CTN “infers” latent TN-structure using raw data filtered by a CT-agent, using method 2	193
Figure 45. Schematic of a typical scenario in which three T-agents are captured by CTa ₁	196
Figure 46. A schematic view of what the counterterrorist network believes the TN’s structure is at a given moment in time.	199
Figure 47. Graphical view of the CTN’s composition-belief matrix	200
Figure 48. Graphical view of the CTN’s structure-belief matrix	202
Figure 49. Behavior of Durkin-summation function	206
Figure 50. Schematic depiction of SOTCAC’s superspace of TN-CTN coevolutions	219
Figure 51. Sample SNA “deconstruction” of an imaginary business that consists of ten workers	224
Figure 52. Screenshot of an early information-matrix of hijacker data	230
Figure 53. Social network of trusted prior contacts of 9/11 hijackers	231
Figure 54. Social network of 9/11 hijackers showing both trusted prior contacts (as in figure 52) and short-lived meeting ties.	233
Figure 55. Social network of 9/11 hijackers showing trusted prior contacts (as in figure 53), short-lived meeting ties, and associated support network.	235
Figure 56. Map showing that 11 of the 19 hijackers came from a single stretch of Highway 15 of the Asir province of Saudi Arabia	237

List of tables

Table 1.	Terrorist networks as complex adaptive systems. . .	7
Table 2.	Some typical social network metrics that measure properties of links between nodes	84
Table 3.	Some typical network metrics that measure properties of individual nodes	85
Table 4.	Some typical social network metrics used o describe entire graphs	86

