Allocation of Resources to Defend Spatially Distributed Networks Using Game Theoretic Allocations

by William M. Kroshl

B.A. in Economics, June 1975, Northwestern University
M.S. in Operations Research, March 1988, United States Naval Postgraduate School

A Dissertation submitted to

The Faculty of
The School of Engineering and Applied Science
of The George Washington University
in partial satisfaction of the requirements
for the degree of Doctor of Philosophy

January 31, 2015

Dissertation directed by

Thomas Mazzuchi
Professor of Engineering Management and Systems Engineering
and
Shahram Sarkani
Professor of Engineering Management and Systems Engineering

The School of Engineering and Applied Science of The George Washington University certifies that William Mark Kroshl has passed the Final Examination for the degree of Doctor of Philosophy as of November 13, 2014. This is the final and approved form of the dissertation.

Allocation of Resources to Defend Spatially Distributed Networks Using Game Theoretic Allocations

William M. Kroshl

Dissertation Research Committee:

Shahram Sarkani,
Professor of Engineering Management and Systems Engineering
Dissertation Co-Director

Thomas Mazzuchi,
Professor of Engineering Management and Systems Engineering & of Decision Sciences
Dissertation Co-Director

Jason Dever
Professorial Lecturer in Engineering Management and Systems Engineering
Committee Member

Pavel Fomin
Professorial Lecturer in Engineering Management and Systems Engineering
Committee Member

E. Lile Murphree
Professor Emeritus of Engineering Management
Committee Chair

Dedication

I dedicate this research to the members of my family: our children, Jennifer, Heather, William, and Mark; our two sons-in-law, Matt and Ahmad; and our daughter-in-law, Kara. They were a great source of encouragement throughout this entire process. I also wish to mention my six grandchildren: Natalie, Claire, Bradford, Liam, Ryan, and Ethan. I hope they take away from this two very important life lessons: to never give up striving to achieve your goals, and to remember that you are always learning.

Most of all, I want to dedicate this effort to my wife Tina and thank her for years of unflagging support, encouragement, optimism, and her uncanny knack to ask the right questions to get me moving past obstacles. She has been a constant source of strength and wisdom for 40 years, whether that was finishing my BA, working on my MS, and especially the last few years while completing my PhD. She was a wellspring of motivation and a guiding light when I thought I was lost in the maze and would never find my way out. Tina, I would never have come this far without you.

Acknowledgements

I would like to acknowledge the efforts of my advisors, Dr. Thomas Mazzuchi, and Dr. Shahram Sarkani, for their guidance and mentoring during this entire process. I would also like to acknowledge my colleagues at the Applied Physics Laboratory for their support, insights, and encouragement over the past few years. While *in toto* they are too numerous to mention, I would like to make a special mention of Dr. Dave Botto, Dr. Joe McDonough, Dr. Dan Katz, Dr. Alan Zimm, and Dr. Bryan Gorman, who were unfailing sources of encouragement, sage advice, and support. The management team at the Applied Physics Laboratory and within my department deserves special mention for supporting me in this endeavor over the last few years. I would indeed be remiss if I did not acknowledge the contribution of the camaraderie of my classmates who have shared this journey with me. I would like to thank my unknown peer reviewers who contributed greatly to the final form of my article in *Risk Analysis*. I would also like to thank the countless professors who have taught the classes I attended over the years at Northwestern University, The Naval Postgraduate School, and George Washington University.

I would also like to make a special mention of one of my professors at the Naval Postgraduate School, the late Professor Rick Rosenthal, who first planted the seed of pursuing a PhD into my thinking: a seed which took 25 years to finally reach full bloom with this research.

Abstract of Dissertation

Allocation of Resources to Defend Spatially Distributed Networks Using Game Theoretic Allocations

This dissertation presents research that focuses on efficient allocation of defense resources to minimize the damage inflicted on a spatially distributed physical network such as a pipeline, water system, or power distribution system from an attack by an active adversary. The allocation methodology recognizes the fundamental difference between preparing for natural disasters such as hurricanes, earthquakes, or even accidental systems failures and the problem of allocating resources to defend against an opponent who is aware of and anticipating the defender's efforts to mitigate the threat.

Conceptualizing the problem as a Stackelberg "leader–follower" game, the defender first places his assets to defend key areas of the network, and the attacker then seeks to inflict the maximum damage possible within the constraints of resources and network structure. The approach is to utilize a combination of integer programming and agent-based modeling to allocate the defensive resources. The criticality of arcs in the network is estimated by a deterministic network interdiction formulation, a maximum-flow linear program (LP), or a combination of both of these methods, which then inform an evolutionary agent-based simulation. The evolutionary agent-based simulation is used to determine the allocation of resources for attackers and defenders that results in evolutionarily stable strategies in which actions by either side alone cannot increase their share of victories.

These techniques are demonstrated on several example networks using several different methods of evaluating the value of the nodes and comparing the evolutionary agent-based results to a more traditional, Probabilistic Risk Analysis (PRA) approach. The results show that the agent-based allocation approach results in a greater percentage of defender victories than does the PRA-based allocation approach.

Table of Contents

List of Figures

List of Tables

Glossary of Terms

1. Attackers: Resources applied to the network in an attempt to damage it by the individual who is seeking to cause harm to the network.
2. Adversarial Risk Analysis: Risk analysis in which one is confronted by a thinking, planning adversary who is trying to maximize the damage he causes.
3. Allocation Model: The agent-based simulation developed in this research, which is used to determine an Evolutionarily Stable Solution (ESS) and thus the Game Theoretic (GT) allocation of resources. In this model, both the defenders and attackers evolve their strategies in response to one another.
4. Attacker–Defender: A type of two-person game in which one player chooses a strategy first, without any knowledge of the second player's strategy. The second player then chooses a strategy with full knowledge of the first player's strategy.
5. Classic Risk Analysis: A method of assigning Risk in which Risk is defined as the probability of an occurrence times the consequences of that occurrence.
6. Consequence: The result of risk being realized.
7. CPLEX: A solver for mathematical programming
8. Defenders: The resources applied by the owner of the network to protect nodes from damage by the opponent.
9. Evaluation Model: The agent-based model developed in this research, which allows attackers to evolve their tactics over time to respond to a fixed defensive strategy. This model was used to compare the effectiveness of various defensive resource allocations.

10. GUROBI: A solver for mathematical programming.

11. Leader–Follower: Another name for an Attacker–Defender game.

12. Probabilistic Risk Analysis: A risk analysis technique in which the probability of a specific risk occurring is estimated and used to calculate risk.

13. Resilience: The ability of a network to maintain operation after it is damaged. This includes redundancy, additional capacity, design for graceful degradation and recovery, and other design features that improve the ability of a network to continue to function when damaged.

14. Resource: In the context of this research, a resource is a fungible measure of capability that is assigned to a particular node for either attack or defense.

15. Stackelberg Game: An attacker–defender or leader–follower game used in economic modeling.

16. Strategy: In the context of this research, a strategy is an allocation of resources to either attack or defend nodes of a specific network. One specifies a strategy by listing each node and then listing the number of attacker or defender resources allocated to each node in the network.

17. Tactics: In the context of this research, tactics refer to the specific rule set that governs how any particular agent (unit or resource) chooses which node to attack or defend.

18. Terrorists: Synonymous with attackers to the network: Individuals who desire to harm the network or interfere with its operation in some manner

Chapter 1: Introduction

1.1 Background

Those individuals responsible for the security of networks and systems such as pipelines, power grids, and water or information systems are faced with the problem of protecting those systems from varying threats. These threats can be both those resulting from natural events and from human design. One of the most pivotal decisions that must be made is the allocation of protective resources to various portions of the network. These resources could be purely passive, such as "hardening" a node against physical intrusion. The resources could be utilized for the gathering of intelligence or for the active defense of a portion of the network. The resources could also be allocated to active kinetic or nonkinetic actions to render an attack ineffective against a particular node. Although this research concentrates on the nodes of a network in this effort, the ideas could be equally applicable to arcs that connect the nodes. Ideally, there would be sufficient resources to prevent the success of any possible attack against the network, allowing it to maintain its functionality.

One can theoretically develop an allocation plan to defend against almost any conceivable threat. However, in practice resources are never "unlimited," and decisions must be made regarding how they should be allocated. For this analysis, "resources" are considered to represent a finite, fungible, and unitary measure of capability that could be applied to increase the defensive capability or "hardness" of a particular node. Resources could be the effort applied to gathering intelligence about a particular part of the network or surveillance to provide information about those attacking a particular part of the network. Resources could just as easily be physical hardening of a key node (shipping terminal in a pipeline or transformer station on a power distribution line) or individuals assigned to defend or target a particular node. When allocating resources in the real world to protect the network, the network defender is often at a disadvantage in that resources must be allocated without direct, specific knowledge of the attacker's plans (J. L. Cox, 2009). The attacker's goals, values, and objectives might be known in a general sense, but the specifics of force allocation, timing, and points of attack are generally not known with certainty. Additionally, many resources devoted to defense of network nodes are not easily moved once in place. Sensors, fences, and other fixed systems, once installed, generally are not repositioned. These resources also have a limited range for their effects.

The attacker also has limited resources. Thus, the attacker must determine the best method of allocating attack resources against the network to accomplish his stated goal, which might involve limiting the effectiveness of the network by reducing the flow of "product" through the network. In this context, "product" is considered to be whatever flows through the network. It could be a physical product such as fluids or gas through a pipeline, or energy such as electricity through a distribution system, or even information through a data network. The attacker's goals could also be to gather publicity for a particular cause, affect political decisions, or induce fear in a population. Many such goals are possible, and any system that recommends resource allocation should include the ability to consider these multiple goals. The attackers, however, have a great advantage because their decision regarding which nodes to attack, and how to allocate resources, will be made with at least some knowledge of the defensive capability and resource allocation of the targeted network. In essence, while the defenders must allocate resources without knowing the plans of the attacker, the attackers can see what the defender has done and can then plan accordingly.

The situation under discussion is different from the situation in which a planner decides on resource allocation in the face of natural, “non-thinking” events resulting from random occurrences in nature such as weather, random failures of equipment, or other natural disasters. The network “attacker” in this case is not making decisions of any kind. The events are considered random and would occur regardless of any actions of the defender. While many of the tools, techniques, and terminology used for analysis of courses of action against forces of nature are similar to those used for the allocation of resources against a determined, thinking opponent, the problems have a fundamental difference. The forces of nature do not change their points of attack or allocation of resources on a network node based upon the defenses of that node, but a terrorist will gather as much intelligence as possible when deciding on targets and resource allocations. A hurricane does not change its path based on whether or not you have installed hurricane shutters and emergency power, but a terrorist may pass on attacking a node that is heavily defended by multiple layers of defense and sensors. In this work, several different methods that can be used for the resource allocation decision are discussed, focusing on the situation in which an active and malevolent “opponent” exists who is attempting to inflict the maximum damage to the network. Without loss of generality, this opponent and the opposition forces controlled by this individual are referred to as “terrorists.”

Most common models or methods for resource allocation are based on some measure of “Risk.” Risk analysis in the form of equating risk with a probability of occurrence and severity of consequence, such as the classic "risk cube," is a concept very familiar to engineers and decision makers. The Department of Homeland Security (DHS) has repeatedly maintained that the concept of risk must be used in establishing priorities for the department (Dillon, Liebe, & Bestafka, 2009). Many risk models are based on the classic formula:

$$\text{Risk}=(\text{Threat})(\text{Vulnerability})(\text{Consequence}) \tag{0.1}$$

As reported by Cox (L. A. T. Cox, Jr., 2008), the DHS standard for risk analysis in the chemical industry beginning in 2007 is based upon this method. In his article, Cox goes on to discuss how this model, while very useful, is limited when confronting an active opponent. The problem he highlights is that the uncertainty represented in this model is fundamentally different if it is a result of acts of nature or acts of man. Golany et al. (Golany, Kaplan, Marmur, & Rothblum, 2009) discussed this at some length in 2009, characterizing the two types of uncertainty as either caused by random acts of nature or insufficient knowledge (Probabilistic Uncertainty) or caused by the lack of knowledge of the plans of an active adversary (Strategic Uncertainty).

1.2 Research Question

The primary question for this research is to determine an approach that, under Strategic Uncertainty, provides a better allocation of resources than does a straightforward application of the classic Probabilistic Resource Allocation (PRA) method, which is often used (and quite successfully) for the allocation of resources to counter the effect of various natural disasters. It is desirable that the proposed method meet the following criteria:

1. The method must be usable for actual real-world problems, not just small textbook examples (Scalable).
2. The method must be better than, or at least as good as, the classic PRA-based approach (Efficient).
3. It should assume a thinking, plotting adversary who will seek to maximize the effectiveness of his attack against the network (Responsive).
4. It should be applicable to many different means of node evaluation (Flexible).

Specifically, the research hypothesis is that, when facing an opponent who has knowledge of the defender's resource allocation and can plan accordingly, an allocation approach based on a Game Theoretic (GT) allocation of resources will perform as well as, or better than, an allocation based upon a pure PRA-based approach, which only takes the value of the nodes under consideration without considering the potential actions of the attacker.

1.3 Overview of Technique

Using the taxonomy of uncertainty described by Golany and Kaplan (Golany et al., 2009) in the preceding section, an approach that only looks at the value of each node and makes an allocation of resources based upon the value of the node is focused on probabilistic uncertainty. The approach developed in this research for allocating resources under strategic uncertainty is one that combines deterministic network interdiction for node evaluation with an agent-based approach to finding the solution of the underlying GT model that is most favorable to the defender. The situation is considered as a two-person game of the type generally known as "leader–follower," "attacker–defender," or Stackelberg games. In this formulation, one person (the leader) will place his resources without the knowledge of the opponent's actions. The second player (the follower) then allocates resources with the full knowledge of how the leader has allocated his resources.

The method of solution for these "leader–follower" games is through the use of an evolutionary agent-based simulation, with the agents representing individual units of resources to be allocated by the attackers and defenders. The method of evaluation of the value of the nodes is flexible and allows for multiple categories of "value" or importance of any particular node. Each node of the network is considered to have two distinct "values": a public relations value, which represents the value of the publicity gained by an attack against the node, and a network value, which represents the contribution of that particular node to the mission of the network (the delivery of some product or service). The network values were calculated in one of two ways, either through a maximum-flow linear program (LP) (max-flow LP) or by a combination of a max-flow LP with a mixed-integer network interdiction program.

It is not really feasible to use data from actual attacks to evaluate the results of this modeling. The actual data are far too sparse, and what might be available is usually unavailable for unclassified research due to security considerations. Therefore, simulation is used to compare the effectiveness of the method's allocation method with a PRA-based method. By comparing simulation results from a GT allocation and a PRA-based allocation, a relative comparison of the success of each method is calculated and then compared.

1.4 Organization

This paper is organized into eight major chapters. Chapter 1 is the introduction, which provides a brief introduction to the question and frames the discussion for the literature review, which is Chapter 2. After a brief overview, the literature review devotes a section to each of the major topical categories in the review. This chapter completes with a short synthesis section that serves to blend the threads from the literature review topics together into a unified conclusion.

Chapter 3 is a detailed discussion of the approach, including the elements of the problem, the solution approach, and the approach to the evaluation of the solution. Chapter 4 is a discussion of all the various formulations and models used in the research, along with a discussion regarding the implementation of the models and the techniques used in the actual analysis of the results.

Three different networks were used in the actual evaluation of the technique. Chapter 5 discusses a small-scale, 13-node notional model, which was chosen as an initial test of the technique. Chapter 6 extends this to a 30-node model, which represents a simplification of major portions of the Irish Electrical Power Grid. Chapter 7 extends the analysis to an 85-node model, which is a simplification of the Plantation Pipeline that runs from Gulf Coast ports up the East Coast of the United States to Baltimore-Washington International (BWI) and Dulles International airports. Each of these three chapters is organized in the same manner: a description of the network to be analyzed, the results of the GT and PRA-based resource allocations, and a comparison of the results.

Chapter 8 contains conclusions and recommendations for further research.

Chapter 2: Literature Review

2.1 Overview

The initial literature review for this research cut across a wide range of disciplines and topics. Initially, the review focused on gathering information on various methods of allocating resources and assessing risk for systems subject to attack by thinking opponents or other natural events. Over time, this became more focused, encompassing various aspects of adversarial risk analysis, network design, terrorism characterization and analysis, and probability estimation. As the solution space narrowed down to a GT approach, research concentrated on various aspects of game theory and the solution of specific types of games, finally settling upon the determination of Evolutionarily Stable Solutions (ESSs) to certain types of games. Several hundred scholarly articles and standard reference works were reviewed and consulted. Many were either duplicative or not germane to the ultimate solution method. However, the topics that proved to be most useful in developing this research are presented in visual form in Figure 1. This organization breaks down the topics into three major categories: Classic Risk Analysis, Adversarial Risk Analysis, and Networks.

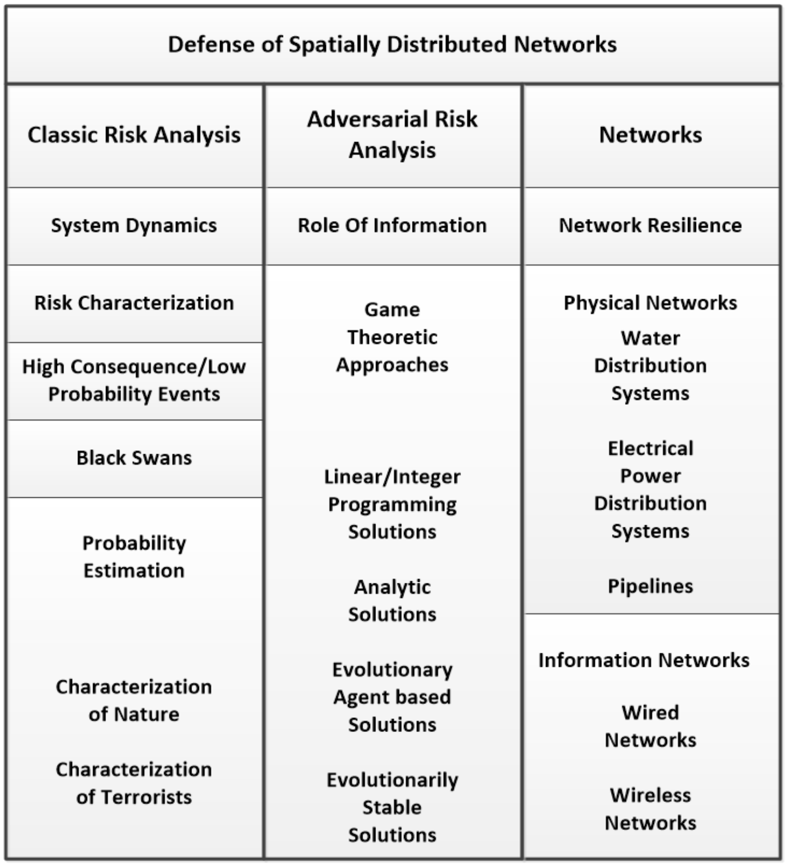

Defense of Spatially Distributed Networks

Classic Risk Analysis

System Dynamics

Risk Characterization

High Consequence/Low Probability Events

Black Swans

Probability Estimation

Characterization of Nature

Characterization of Terrorists

Adversarial Risk Analysis

Role Of Information

Game Theoretic Approaches

Linear/Integer Programming Solutions

Analytic Solutions

Evolutionary Agent based Solutions

Evolutionarily Stable Solutions

Networks

Network Resilience

Physical Networks

Water Distribution Systems

Electrical Power Distribution Systems

Pipelines

Information Networks

Wired Networks

Wireless Networks

Figure 1 Literature Review Topics

Classic Risk Analysis involves the use of estimates of probability of occurrence and the consequence of the occurrence as key factors in the estimation of risk. This approach requires that the probability of an occurrence and its impact upon the system (consequences) must be estimated or measured in some manner. The estimation of these parameters highlights the differences between characterizing terrorists, who are thinking, planning entities, and characterizing natural events, which occur with no conscious thought, design, or plan. Many different techniques, both analytic and simulation-based, have been proposed to characterize these risks. Of special interest in this area is the fact that some highly improbable (very low probability) events can have extremely devastating consequences. These high-consequence, low-probability events, under certain circumstances, have a dramatic (some would argue distorting) effect upon resource allocation methods. Related to these types of events are so called “black swans,” i.e., events that are almost totally unforeseen.

The approach used for the allocation of resources in this research is based in game theory, in which two opposing “players” make decisions that affect the outcome of the “game.” Each player is attempting to maximize his own payoff. In this research, a specific type of game called a “leader—follower,” “attacker – defender,” or “Stackelberg” game is used as a basis for the allocation of resources. There are many different variations on these types of games, and multiple approaches to both defining and determining their solution, which can be thought of as the optimal decisions by either player. Pure analytic techniques are suitable for small games, but they lose their utility on most “real-world” sized problems. Other methods, such as linear or mixed-integer linear programming approaches, are used along with evolutionary agent-based approaches.

Networks themselves have a rich literature. They are becoming increasingly important and correspondingly the subject of greatly increased research over the past few years. Two major branches of network risk analysis were reviewed in this work: physical networks such as power grids, water or gas distribution systems, or pipelines for the delivery of various liquid products, and information networks, in which packets of information vice physical product are distributed to nodes through the structure of the network. Another network specific topic that was explored was the concept of resilience, which simply means the network’s ability to continue to function when damaged. This ability is usually the result of the configuration of the connections between nodes in the network. A more resilient network is one that offers multiple redundant paths and nodes to ensure functionality.

For the purposes of this literature review, four major research lanes are discussed. **Classic Risk Analysis** includes general discussions about the quantification of uncertainty, resilience, and redundancy. This literature generally focuses on the ability of a network to withstand damage from a non-malevolent opposition such as weather or other natural events. **Adversarial Risk Analysis** includes analysis of risk in which there are two (or more) opponents, each of which is trying to maximize his own return. **Computational Approaches** focuses on methods to solve the Stackelberg "Leader–Follower" games, including analytic solutions, mathematical programming solutions, and simulation methods. Other useful references are discussed in a section on **Related Discussions**. In each of these research lanes, the references are discussed in chronological order, starting with the earliest reference in each lane, grouped by general topics within the category.

2.2 Classic Risk Analysis

Resilience is a concept that is often used to ensure that a network maintains functionality even if portions of it are damaged. Rose, Benavides, et al., in a 1997 article looked at resiliency as a key component of minimizing the effect of earthquakes in Memphis, Tennessee (Rose, Benavides, Chang, Szczesniak, & Lim, 1997). Through the use of an input/output model of electrical power mapped to economic output, they evaluated various resource reallocation strategies through a linear programming model to assess changes in output. Although their work focused on earthquakes, the method is extensible to other natural and manmade disruptions of a region. Their linear program was designed to minimize the economic disruption through the reallocation of resources. They included the effect of resiliency through modifying the estimates for electrical power generation after the incident and for reducing the need for electrical power after the incident, reflecting alternative sources of power or techniques of reducing demand while maintaining output. In their analysis, a magnitude 7.5 earthquake in Memphis resulted either in a 49.6 or 78.7 percent reduction in regional gross domestic product, depending upon the production assumptions used. However, their linear program model looking at reallocation (i.e., taking resilience into account) showed a 13.5 percent reduction, which was a significant savings. This work focused on the need to capture resiliency and the efficient allocation of assets after an accident, but it has important implications estimating the consequences of severe disruptions of electrical power.

In 2005, Rose and Liao used a Computable General Equilibrium (GCE) analysis instead of the Input-Output models just mentioned to estimate the regional economic impact of the loss of water supply to Portland, Oregon (Rose & Liao, 2005). The primary advantage of this technique is that it made explicit use of "behavioral content" of decision makers, showing a change in behavior of producers and consumers to market conditions. This sensitivity makes it easier to introduce a systematic analysis of resiliency by substitution of required goods and services. The authors further characterize resiliency as either inherent, which is "… the ability of individual firms to substitute other inputs for those curtailed by an external shock," or adaptive, which is "increasing input substitution possibilities…by providing information to match suppliers without customers to customers without suppliers." (Rose & Liao, 2005) In this research, resiliency has been considered in the techniques used to evaluate the network flow and, in some cases, the interdiction of key nodes in the network.

In 2004, Beitel et al. developed a model to consider the selection of targets by terrorists (Beitel, Gertman, & Plum, 2004). This article focused on a model that was developed to consider the selection of targets by terrorists as a part of a larger modeling effort. Employing a "classic" risk model approach, which uses the probability of occurrence and the consequence of occurrence, they sought to develop a methodology of assessing the probability that a terrorist would attack a specific target. Their target model used four investment measures (number of terrorists, terrorist resources, terrorist schedule, and the likelihood of success) and six return-on-investment measures (loss of life, primary economic loss, national economic stress and inconvenience, decreased Western presence, increased Islamic presence, and the opportunity to leverage with other terrorists). These measures were then combined in a multiplicative manner. Beitel et al. found that this model was very close to predictions based on historical data. This work was key in showing that the methods developed needed to be capable of dealing with multiple (nonhomogeneous) terrorist goals.

Recognizing the need to consider the behavior of the adversary in developing defensive strategies was a central component of Bier and Nagaraj's 2005 paper, in which they stated that "good defensive strategies must consider the adversaries' behavior" (V. Bier, Nagaraj, & Abhichandani, 2005). They maintained that the goals and motivations of the attackers must be considered, to determine if they are opportunistic or determined. If the attacker is opportunistic, he will be looking for easy targets. If the attacker is determined, it can be much more difficult to counter. To protect a specific target against opportunistic attackers (such as vandals) it is only necessary to provide a defense that is a little better than the rest (relatively). However, for a determined attacker this will not help appreciably.

The authors developed a series of equations looking at both series and parallel systems of two targets, and then they developed equations more generalizable to many components or targets. For a parallel system, the probability of a successful attack is the probability of successfully attacking all the components. However, in a series system, the probability of a successful attack against the system (assuming perfect knowledge) is the probability of success against the weakest component of the system. They maintain that, if the defenses are known to the attacker, the defender should seek to make all targets equally difficult, illustrating the importance of secrecy in a series system: For a fixed budget, the defender will achieve better results if the attacker has no knowledge of the defensive preparations. They recommended continued work on sequential attacks and on attacks with Bayesian updating of the probabilities. They stated, "…it would be worthwhile to extend our models to include a time dimension, rather than the current static or snapshot view of system security. This would allow us to model imperfect attacker information (including Bayesian updating of the probability that an attack will be successful, as estimated by both attacker and defender) and the possibility of multiple attacks (either attacks against multiple components, or multiple attacks against the same component)" (V. Bier et al., 2005).

A slight variation of the standard DHS risk approach of "Threat–Vulnerability–Consequence," as discussed in Equation 1.1, was proposed by Ayyub et al. in 2007, who applied it to the study of infrastructure in Maryland (Ayyub, McGill, & Kaminskiy, 2007). They emphasized that there was an inherent difference between natural disasters and terrorist attacks because the deliberate attacks were the result of intelligent human adversaries. These people (the terrorists) are able to adapt and change depending upon the situation. Because the actions of the attacker are inherently unpredictable, the effectiveness of any countermeasure is more difficult to measure because of the uncertainty in the capabilities and intentions of the attacker. The authors therefore started with the most certain (the consequences) and then proceeded to the least certain (the attacker actions). They developed a five-phase approach to risk analysis for systems, which resulted in an all hazard risk expression very similar to the standard DHS consequence, vulnerability, and risk model.

In 2008, Bier, Haphuriwat, et al. approached the problem by relying on the minimization of expected loss from the attacker through proper allocation of the defender's resources, constrained by a budget and considering "cost effectiveness" of the improvements to infrastructure (V. Bier, Haphuriwat, Menoyo, Zimmerman, & Culpen, 2008). An interesting part of their work was the observation that the more uncertainty in the defender about an attacker's valuation of targets leads to the more efficient policy of spreading resources across multiple targets. The same year, Cox postulated that the traditional product approach embodied in Equation 1.1 (the DHS-based model) was limited due to limitations in the product rule, the potential correlation between components, and the subjectivity and ambiguity of the categories.(L. A. T. Cox, Jr., 2008). The DHS-based model was based on the Risk Analysis and Management for Critical Asset Protection (RAMCAP) model, which used a "reasonable worst-case" approach. This model contains rules for both quantitative and qualitative risk assessment.

In this paper, Cox demonstrated that intelligent attackers can optimize their attack method in such a way as to achieve the maximum damage by changing their attack plans as needed. He reviewed problems with that model due to aggregating attack scenarios, non-additive vulnerabilities, estimating the expected value of a product of random variables, risk ranking inadequacies, and the difficulty of assessing threat vulnerability from event trees. He suggested that attackers' actions should be modeled in such a way as to achieve the maximum damage by changing their attack plans as needed. Methods he recommended included decision tree analysis, probabilistic activity AND-OR networks, planning models of attacks, and hierarchical optimization and GT approaches.

Golany and Kaplan provided an excellent working taxonomy regarding the nature of uncertainty in their 2009 paper (Golany et al., 2009). Their structure, briefly discussed in the preceding chapter, partitioned uncertainty into two categories: probabilistic uncertainty, which is caused by states of nature, and strategic uncertainty, which is caused by not knowing the plans of an adversary. For probabilistic uncertainty, they proposed a knapsack type optimization problem called PRAP (Probabilistic Risk Assessment Problem). For strategic uncertainty, they proposed SRAP (Strategic Risk Allocation Model). The forms of the objective function were slightly different. However, their results can be summarized with some simple rules. When dealing with probabilistic risk, they recommended the allocation resources evenly across all targets. When dealing with strategic risks, they recommended protecting the highest value target first, until it reaches the level of the second target. Move both of those down together until you reach the third target level, and continue to "add" targets to be protected, keeping them all "level" as you continue down the list until you run out of resources.

That same year, Dillon and Liebe reported on the problems of defining acceptable levels of risk in the standard DHS models, based on Equation (1.1) (Dillon et al., 2009). They focused on the problems of aggregating different consequences and properly defining baseline risks. The use of multi-attribute utility theory and risk analysis recognizes the problem and contributes to the solution of these challenges, which leads to the ARDA model (Anti-Terrorism Risk-based Decision Aid). Using this model involves a process of defining the system, assessing the baseline risk, assessing the risk mitigation alternatives, prioritizing the risk mitigation alternatives, and making a risk-based decision with feedback back into the process. Consequences are defined in terms of Mission Impact, Personnel Impact, and Economic Impact and then combined in a linear function to determine total consequence through a utility score. The authors propose a model for looking at portfolios, which eventually evolved into a risk cube approach weighting risk in the traditional manner. The utility theoretic approach and multi-attribute model tend to show that strategies that address a wide range of risks are preferred over those that counter a single strategy. However, the development of these utility functions is often problematic.

Brown and Cox, in a recent (2011) article, highlighted the problems of using a typical systems engineering approach using PRA-based tactics against a thinking adversary. They specifically criticized the assumption that the "…same type of conditional probability assessment applies as well to terrorism risk analysis as to probabilistic risk analysis of natural hazards and engineered systems…" (G. Brown & Cox, 2011). They further claim that applying these techniques may increase the risk of terrorist attack when applying the traditional Threat Vulnerability Consequence (TVC) concept. They give several examples of how PRA estimates may be self-defeating (attackers deliberately not attacking targets that PRA directs them to attack because they know the defender will be defending against them). Their key points include the following, which are all taken from (G. Brown & Cox, 2011):

- "…the probabilities of alternative attacker actions…assessed by the defender should depend on what research opportunities are available to the attacker." (p. 199)
- "…the probabilities of different attacks (based on the attacker's information) are uncertain and are not uniquely predictable by the defender." (p. 200)
- "Additional research may have zero information value to the attacker and defender." (p. 200)
- "Better information for attackers may reduce threats." (p. 200)
- "…knowledge of threat and vulnerability data… (together with consequence data), does not allow us to predict how alternative risk management interventions will affect risk. Other factors are essential, such as what types of attacker resources … produce how much damage when deployed in alternate ways." (p. 202)
- "PRA, as currently practiced, does not represent the function of the infrastructure at all." (p. 203)

The points highlighted in this article provide a succinct and useful summary of the problems inherent in using the classic PRA-based approach when looking at the defense of infrastructure.

2.3 Adversarial Risk Analysis

Insua and Rios discussed the key difference between Adversarial Risk Analysis (ARA) and PRA in 2009 (Insua, Rios, & Banks, 2009). Beginning with the fact that there are two decision makers in an ARA formulation and only one decision maker in a PRA problem, they demonstrated that these problems are a natural fit for GT solutions. These ideas were developed further in a later work (Rios & Insua, 2012) in which they further explore the role of GT formulations for ARA problems. Game theory is frequently used to describe and solve ARA problems. This section includes some references regarding the use of game theory in these problems.

Game theory has been discussed as a means of looking at thinking adversaries for many years. However, it has often been, as Hall indicated in his 2009 letter to the editor of Risk Analysis, "… the elephant in the room." (Hall Jr, 2009).

Before reviewing recent articles on the application of game theory to the defense of networks, the following discussion highlights several useful standard references regarding game theory and points out elements of game theory especially germane to this issue. One of the early and very useful examples focused on how cooperation can evolve in real-world systems (Axelrod, 1984). Building on the simple example of the Prisoners' dilemma, which is an analysis of whether two prisoners are better off cooperating against their jailor or competing against each other, Axelrod discussed many different conditions that can give rise to situations in which competition or cooperation can be better off for the individuals involved.

The essence of the Prisoner's Dilemma as discussed by Axelrod is the following: There are two prisoners, each of whom must choose to either cooperate with the other prisoner or defect. If both prisoners cooperate, each gets a payoff whose value is equal to 3. If one defects and one cooperates, the defecting player gets a payoff of 5 and the cooperating player gets a payoff of 0. If both players defect, each one gets a payoff of 1 (Axelrod, 1984).

Myerson's 1991 book on Game Theory (Myerson, 1991) provides a complete treatment of many of the elements of game theory used in this research. As he discusses on page 2, a game is defined as a social situation involving two or more individuals. These individuals, called players, are considered to be rational and intelligent. Because they are rational, they seek to maximize the benefit to them. This benefit is called the utility of the game. In the simple form of a game, two players are opposed to each other. In the example of the previous paragraph, the players must make a decision as to which strategy to follow. In the context of the Prisoner's Dilemma game, a strategy is the decision as to whether to cooperate or defect. The players make this decision without knowing the decision of the other player.

As in many decision theory models, the "best" strategy in some sense may not be a "pure" strategy, or one in which the same decision is made every time the game is played. A mix of pure strategies may be the best, where the mixed strategy is a defined probability that each pure strategy is chosen on any one turn or play of the game. At a conceptual level, consider a game in which one player must decide which of two possible routes to take to reach an objective. One route is faster—but more dangerous due to natural hazards—than the other. The second player must place assets in such a way as to defend one of the two paths. If each player knew what the other player would be doing, the problem would be simple. However, by appealing to common sense, it is readily believable that some variation in strategies would be best (sometimes taking one route, and sometimes the other). Game theory serves as a means to quantify this intuitive insight.

Many of the solutions to these game theory problems are characterized as finding the Nash equilibrium of the game. The Nash equilibrium is a combination of strategies in which both players are at their optimal payoffs and neither player can increase his payoff by his own actions alone. As brought out in Chapter 3 of *Game Theory: Analysis of Conflict* (Myerson, 1991), a randomized strategy vector is one which, for each pure strategy, specifies a probability that players would choose that pure strategy. As stated on page 93, the condition for a Nash equilibrium is "…a randomized strategy profile is a Nash Equilibrium if and only if no player could increase his expected payoff by unilaterally deviating from the prediction of the randomized strategy profile." (Myerson, 1991). Note that there can be multiple Nash equilibria for a game, each with different payoffs. Myerson goes on to present an analytic approach to the solution of finding a Nash equilibrium but warns that this method may be unworkable as the size of the problem increases.

Another very useful reference for general game theory problems and their solutions is *Game Theory Evolving* (Gintis, 2009). Of particular importance to this research is that Gintis proposed and discussed an evolutionary approach to the solution of game theory problems. In particular, he showed how, "… in a process where successful strategies drive out unsuccessful ones over time, stable stationary states are always Nash

Equilibria." (Gintis, 2009). Formally, he stated his Nash Existence Theorem as "If each player in an n-player game has a finite number of pure strategies, then the game has a (not necessarily unique) Nash equilibrium in (possibly) mixed strategies." (Gintis, 2009).

Often GT approaches were used to evaluate terrorist's goals, or the effects of negotiating with terrorists, such as described in 1988 by Sandler and Lapan (Lapan & Sandler, 1988). In this article, they developed a theoretical framework for looking at the question of bargaining between terrorists and governments, specifically on the wisdom of having a "no negotiation" policy with terrorists. The authors framed their argument from an economic standpoint. They treated the problem as a game theory problem and developed an expected value approach to determine when, based on the goals of the terrorist, it might make sense to negotiate. They also considered the reputation effects from negotiating and how changes in reputation over the long haul might affect the decision. They recommended modeling reputation as uncertainty of some aspect of the government information set.

Also in 1988, the authors published other research analyzing how terrorists determine their targets, and how "over deterrence" or "under deterrence" could result in transnational vice national terrorism (Todd Sandler & Lapan, 1988). In this article, the authors took a GT look at how terrorists determine their targets. They began with a review of then current terrorist attacks and pointed out that terrorists are not random but rather have specific goals in their attacks. Their basic model was a two-target, single-attacker model in which only one target can be attacked. Each target was assumed to be in a different country. The terrorists can either not attack any targets, or attack Country 1, or Country 2. They developed a decision tree analysis for payoffs to the terrorists and costs to the respective governments. Their decision rule was that the terrorist group will not attack if the expected payoff from attacking is negative. They held that the terrorists would attack the country with the highest net expected payoff.

Sandler & Lapan then developed payoff models, with costs to each country involving the deterrence expenditures, the expended cost of domestic attack, and the expected costs of an attack on another country (Country 1's assets located in Country 2). They characterized the cooperative solution as efficient because all costs are included. For domestic terrorists (military or civilian targets, for example) a non-cooperative solution (one in which the targets do not cooperate) will imply over deterrence, as compared with a cooperative solution. Interestingly, more information may cause the non-cooperative solution to depart further from the cooperative solution. The fallout from this is that sharing intelligence between terrorist targets may worsen efficiency rather than improve it when these governments do not also coordinate their deterrence expenditures. They also found that piecemeal policies to thwart terrorism may be detrimental.

On the domestic side, the choice of targets was modeled between government (hard) and private (soft) targets. In this case, non-cooperation leads to over-deterrence. In the transnational case, however, if the attack against Country 1 imposes equal costs on both countries, then the non-cooperative solution leads to under-provision of deterrence. An increase in information may worsen inefficiency when targets do not cooperate when choosing deterrence. Sandler & Lapan then move on to adapt the model to suicidal terrorists who are not affected by deterrence.

Their major findings were that over-deterrence can be identified in domestic terrorism, and that either over-deterrence or under-deterrence can result in transnational terrorism. This article was a relatively early example of the use of game theory in planning to counter terrorist attacks, and although the example was small, it was an excellent example of how game theory could be applied to the problem.

Overgaard examined the problem of imperfect information between terrorists and the government (Overgaard, 1994). The article posits a two-person game environment in which the terrorist has perfect information about the government, but the government may not have perfect information regarding the capabilities of the terrorist: if they are a high-resource group or a low-resource group. This situation was called a "signaling game" as discussed in (Lapan & Sandler, 1988) and others. The authors proposed that if an opponent terrorist group is rich in resources, then it was to the benefit of the government to bargain. They also brought out differences between the single-period and multi-period models. They attributed these differences primarily to the storability of terrorist resources.

Sandler and Arce provided several reasons game theory was an appropriate vehicle to analyze terrorism (Todd Sandler & Arce, 2003). On page 4, they listed several specific reasons for the appropriateness of game theory:

1. It captures the strategic interactions between terrorists and the targeted government (interdependent actions of both players);"
2. "It models strategic interactions between rational actors;"
3. "Each side issues both threats and promises to the other in order to gain strategic advantage (signaling in game theory);"

4. “Both terrorists and governments seek to maximize a good under constraints (resource and constrained actions);”
5. “Game theoretic bargaining analysis is applicable to hostage situations, which are frequently encountered in terrorism analysis;” and
6. “Both uncertainty and learning are relevant in the analysis of terrorism.” (Todd Sandler & Arce, 2003)

They gave several examples of two-player and three-player games, looking at active deterrence or passive hardening of targets and the role asymmetric information. Another interesting and informative use of game theory was their analysis as to whether or not a terrorist should attack either a hard commercial target or a softer tourist target. They analyze this situation as a game with four possible outcomes, either successful or unsuccessful for each of the two target types. Their analysis showed that the soft (tourist) target actually is a more attractive target, which they claim explained why so few military and government targets have been attacked, relative to commercial targets.

These same authors expanded on these ideas in a 2005 article, in which they specifically looked at a two-person zero-sum game examining the actions of two countries adopting the strategies of active, pre-active, or defensive action against terrorists (Arce & Sandler, 2005). The authors looked at a 3x3 matrix of actions between two different countries, looking at the potential of preemption, deterrence, or doing nothing. These are analyzed as a series of 2x2 games, first looking at decisions between preemption and doing nothing, then looking at doing nothing versus deterrence, and finally at the 3x3 example. Their major result was that the policy of deterrence dominates over preemption, even though it may not be the most optimal policy.

Electronic networks of sensors or computers have been the subject of several articles postulating the value of the GT approach. Agah and Basu examined a Denial of Service (DOS) attack against a network of wireless sensors (WSAN) (Agah, Basu, & Das, 2005). They concentrated their analysis on sensor-to-actor and actor-to-actor coordination within the network. A DOS attack was characterized by how it affects the ability of sensors to communicate with each other via the network. They consider two types of attacks: nodes not forwarding control messages or data and nodes falsifying route messages or causing error messages. These attacks are modeled as a two-player, non-zero sum, non-cooperative game. They rated nodes on cooperation (a node that has more power or energy forwards more messages) and reputation (whether or not a node forwards messages from its neighbors). Two different types of routing strategies for information were used: One is a linear combination of reputation and cooperation; this is a utility-based model. The second one is a watch list, where each actor monitors its neighbors, evaluates their trustworthiness, and routes accordingly. They compared two different algorithms for efficiency.

Also in 2005, Lye and Wing examined intrusion detection and defense of networks as a general sum game (Lye & Wing, 2005). They utilized a nonlinear program to solve for the Nash Equilibrium points of the game, taking as an example a network with a web server and private server. They consider the combinations of attacker and defender strategies for possible interactions. The game was a general sum game because the attacker and defender payoffs were not symmetric. They considered a multiple time period model of infinite time duration, discounting future benefits at varying parametric rates to consider both the possibility of favoring far future payoffs or near future payoffs.

Their analysis assumed that both the attacker and defender have perfect knowledge of each other's actions.

Machado and Tekinay looked at the problems of power management, reliability, and the preservation of wireless sensor networks in the face of both natural problems and terrorist attacks by using tree diagrams to examine these relationships (Machado & Tekinay, 2008). They described several different applications of Pursuit Evasion Games (PEG), which provided an excellent theoretical construct for modeling sensor networks in which one player is seeking to avoid surveillance and the other is trying to maintain surveillance. Of particular note to this research, they made a distinction between classic game theory, in which the players have decision making abilities for obtaining payoff maximization strategies, and evolving game theory, in which the most successful strategies are used with greater frequency by various agents. The method presented for solving the resource allocation problem in this research involves the use of evolutionary game theory.

Anderton and Carter looked at the problem of modeling terrorism as a microeconomic problem. They examined choices for terrorism or other activities as a maximization model in which one seeks to set the marginal rate of substitution as equal (Anderton & Carter, 2006). They examined income consumption curve based models as well as price consumption curve models, which examined the sensitivity of terrorists to changes in costs imposed by governments. They also examined cooperation between governments from a game theory perspective. They examined cooperation between governments, as discussed previously by (Todd Sandler & Arce, 2003) which looked at terrorism as a prisoner's dilemma game.

In an excellent 2007 article, Bier expanded on the two-person zero-sum approach by allowing the attacker's and defender's valuation of targets to differ (Vicki Bier, 2007). She began with a basic model of defending two targets and allowed for the defender to have uncertainty as to the attacker's preferences. In modeling this effort, it became clear that the ratio of the probabilities is the most important in determining which target should be defended and how many resources should be optimally allocated to its defense. After reviewing how this game behaved with allocations under different conditions (equal worth, unequal worth, etc.), the article showed that this game formulation behaved in a manner consistent with real-world expectations. The author also spent a great deal of time illustrating how centralized decision making for allocation of defensive assets was more efficient than decentralized allocation, which may defend targets that should not be defended, leading to wasted resources. The discussion then flowed into the problem of per capita allocation of anti-terrorism resources, which is wasteful when the probability of attacks in a particular area is very low. She also provided an example showing the value of intelligence and the value of deception for the defender.

Both cooperative and non-cooperative games were reviewed by examining the budget tradeoffs between anti-terror actions and purely defensive actions under a fixed budget allocation in a 2009 article by Sandler and Siqueira (T. Sandler & Siqueira, 2009). They brought out the role of central authority in decision making, discussed the signaling games inherent in imperfect communications between governments and terrorists, and demonstrated how to analytically examine suicide bombers in a game theory analysis.

They also brought up how central authority results in more efficient decisions and resource allocations, agreeing with the previous article by Bier (Vicki Bier, 2007). They then examine the proactive – status quo – defensive game decision as a 3x3 game. In this area, they expand upon (Arce & Sandler, 2005). The authors also examined how to quantify the effect of a nation's independent choices ignoring the spillover effect of common benefit to all. These effects can be particularly important when considering the tradeoff between proactive measures and purely defensive measures. They then mentioned how leadership can reduce inefficient allocations of resources. They examined the time sequencing of these decisions and showed that altering of the sequence of choices did not affect the results. They also discussed the problem of games of asymmetric information: the signaling of intentions by terrorists. These games are called "signaling games." There was a discussion suicide terrorism, multi-agent and multiplayer games. The key idea they maintain with respect to signaling games is that interplay among targeted governments over time can cause schism of terrorist groups over time.

At the start of this section, I quoted Hall as saying that "…the elephant in the room is game theory" (Hall Jr, 2009). In a response to that letter, Cox (J. L. Cox, 2009) made the point that the games most relevant to terrorist analysis were the leader–follower or attacker–defender games in which first one "player" allocates resources, and then the second player allocates his resources. He claimed these were similar to Stackelberg games in economic theory and discussed how they were complementary to PRA-based approaches. He claimed that in many cases the full "game theoretic" approach was unnecessary because of perfect information and a turn-based logic, so a simple optimization approach looking for the Nash equilibrium was sufficient. However, he also observed that other approaches (such as a “min-max” linear program) are better if the defender allocates resources first.

Insua and Rios applied influence diagrams to the ARA approach in their 2009 article (Insua et al., 2009). They sought to maximize utility of the expected decisions of the two players. They examined both leader–follower games and games with simultaneous decisions by both players. After discussion about finding Nash equilibria for games using linear programming techniques (to be discussed in more detail shortly), they highlighted the problems inherent in this type of solution, which was the assumption of perfect knowledge of the opponent and the opponent’s actions. Insua and Rios also showed that the solution of a sequential game might not correspond to a Nash equilibrium for the game. They discussed simultaneous move games and developed a unified Bayesian framework for the ARA.

Other approaches were proposed by Ezell and Bennett in 2010, who examined many different techniques, including logic and fault trees, attack and success trees, influence diagrams, causal loop diagrams for system dynamics approaches, Bayesian network analysis, and GT analysis (Ezell, Bennett, Von Winterfeldt, Sokolowski, & Collins, 2010). They utilized an “expanded” version of PRA to include any probabilistic-based approach such as fault trees. While not explicitly taking issue with any GT approach, they did stress the importance and utility of tree diagrams and probabilistic models to describe the actions of the terrorists. Recognizing the importance of taking an intelligent, thinking adversary into account, they traced out the development of the DHS approach on bioterrorism, which began as a “pure” PRA-based approach in 2005 and evolved into an approach that recognized intelligent adversaries by 2008.

One of the very useful contributions of this article was the articulation of issues regarding analysis methods that assumed that the probabilities of various terrorist decisions and actions were an input to the model, and not developed as a result of the model. They believed that two major issues affected these a priori estimates: the incomplete nature of intelligence estimates and fact that such estimates were inherently static, not recognizing the dynamic nature of the terrorists (Ezell et al., 2010). These authors proposed four major families of approaches, three using decision trees and one using event trees. The difference between these decision trees was reflected in the identity of the player whose utility is maximized while the opposing player's choices were reflected as choices modeled on uncertain events. Ezell et al. discussed the need for decomposing large, complex problems into simpler components or sub-problems, and they proposed five major approaches to solving these problems: (1) Logic and Fault trees, (2) Influence diagrams, (3) Causal loop diagrams, (4) Bayesian Networks, and (5) Game Theory.

They found two main problems with many GT approaches: determining the proper Nash equilibrium to study out of the many that might exist, and the fact that the analysis, while it might give the best solution, may not adequately describe the terrorist's actual actions. This succinct summarization of the problems of game theory in the context of analyzing terrorism provided good guidance for the pitfalls of this approach.

The role of imperfect or limited knowledge of the adversary's actions or decision has been a thread running through several of these recently cited articles. This question was specifically addressed by Sorrentino and Mecholsky in the context of solutions to the classic "Prisoner's Dilemma" game, in which the agents have only imperfect knowledge of the other player's strategies (Sorrentino & Mecholsky, 2011). They considered the problem as an evolutionary network and sought analytic solutions for this game. Another article to use the Bayesian approach to developing probabilities appeared in 2011, but this article extended the development of subjective probability distributions to three-player games (Banks, Petralia, & Wang, 2011). They examined both turn-based and continuous games, providing additional examples of the ARA-based approach. Rios and Insua examined the game theory solution to the ARA problem under imperfect information and expanded this to the use of private information: information possessed by the defender that the attacker does not know the defender possesses (Rios & Insua, 2012). Although their analysis was based on attacker–defender games, they used influence diagrams in an essentially analytic approach. Another recent article extended the idea of imperfect information to imperfect observability, where the players may not be able to observe the actions of their opponent (Rothschild, McLay, & Guikema, 2012). Their approach used "level K game theory," in which a recursive algorithm is employed at each level of play to determine the prior probability map for the actions of opponents. This article highlighted the importance of information and the assumptions regarding the visibility of information in the formulation of the model.

2.4 Computational Approaches

Many of the proposed solutions using a GT approach consider the problem as a leader–follower, or Stackelberg game, in which one player first allocates resources without knowledge of the second player's force laydown. The second player then, with the full knowledge of the first player's laydown, allocates his forces. The major groupings of these approaches are linear programming models, analytic solutions, and evolutionary approaches.

One of the earlier approaches used was to solve this leader–follower game as a bi-level optimization problem. This mixed-integer programming approach was proposed by Moore and Bard in 1990 (Moore & Bard, 1990), but it is computationally difficult to solve. Moore & Bard developed several different types of heuristics to approach the solution of bi-level optimization problems in a tractable period of time. The basic problem was defined as having an attacker who sought to maximize the amount of damage done during an attack and a defender who sought to minimize this damage. The primary focus of their effort was on the heuristics, not on the actual formulation of the problem.

Related to these problems is the proper representation of the network, especially the problem of identifying the critical nodes and determining the "weak points" in the network whose loss would result in the greatest loss of total network throughput. Network Interdiction is the determination of the key arcs (or nodes) of a network to target for elimination that will result in the maximum disruption of the flow through the network. Perhaps the single most useful starting point for any discussion in Network Interdiction was written by Wood in 1993 (Wood, 1993), which examined the maximum flow in a network and sought to find a minimal cut set that would stop flow between the source and sink of a planar network. Initially, it looks at the max-flow/min-cut algorithm, then adds the complexity of limited resources for the cuts. The problem with this approach is that solution required an inefficient algorithm to a bi-level optimization problem using a Bender's decomposition. They begin by using the Bender's decomposition of the problem as a max-min problem. They then develop a generalization of integer cutting planes so as to be more efficient in the Bender's decomposition solution. They reduce the number of possible cuts through the use of Super Valid Inequalities (SVI), reducing the size of the feasible region. The authors implement their algorithms through the use of callable subroutines from CPLEX, and then show the efficiency. They also reference how this work is used for hardening a road. Although these two articles did not address the issue of resource allocation in great detail, taken together they provide a solid basis for a great deal of the body of work focused on solving these leader–follower games using linear programming methods.

Brown and Carlyle (Gerald Brown, Carlyle, Diehl, Kline, & Wood, 2005) developed a model called JOINT DEFENDER (JD), which used an attacker–defender framework for allocating defending intercepts to counter nuclear attacks at a theater level. The key to this work is the general model, which is based on an LP solution to leader–follower games. They began by formulating a min-max model, which had as its objective function the maximum value of target damage, subject to constraints for missile availability and available feasible target missile pairings. The goals were then to find the minimum of these possible maximum solutions. While this could be done by specialized decomposition algorithms, it is easier to look at the LP relaxation of the inner problem to create an inner maximization, using the dual variables, and then take the dual again to form a "min-min" problem, which can be solved via a standard solver. The solution of the JOINT DEFENDER MIP gives the optimal defense pre-positioning plan and the interceptor commitment plan. The mobile missile transport and attack plan are then discovered by entering these plans into the JD-min-max, which gives mobile missile transport plan and the attack plan.

The authors discussed the value of secrecy by allowing the attacker or defender to “hide” some of their assets from the applicable LP. They then presented a case study using Korea and discussed how the infrastructure might be evaluated for target values using standard army doctrine. In the appendix, they discuss the variations to the formulation that would allow for targets that can be damaged over time combined with point targets. The same sort of leader–follower formulation was applied to minimizing the penalties for a natural gas shipper, solving the problem as a mixed-integer bi-level programming problem with a reformulation to ease the computational difficulty of solving the optimization (Dempe, Kalashnikov, & Rios-Mercado, 2005).

The following year, Brown and Carlyle did further development of their ideas in three distinct “variants” of models (Gerald Brown, Carlyle, Salmerón, & Wood, 2006), which examined an attack on the U.S. Petroleum Reserve. This article presented several deterministic LP solution templates for solution of infrastructure defense problems. The authors began with a quick review of the 13 sectors critical to national defense and then discussed how “cutting sets” of traditional network analysis fell short of surviving a malicious attack, illustrating the problem of risk-based analysis as compared with GT analysis. They then discussed the ready availability of data from open sources and cited various writings of al Qaeda as evidence of how terrorists exploit the open nature of our society and information sources. They applied various bi-level and tri-level optimization models to examine the complete infrastructure system and its value to society.

They discuss three different variants of models: Attacker–Defender (AD), Defender–Attacker (DA), and Defender–Attacker–Defender (DAD). AD models have as an objective function the value of the infrastructure system. This formulation is a bi-level problem, where the defender chooses to minimize the losses, and the attacker tries to maximize the min losses, leading to a max-min problem. They then developed a method of taking the dual of the inner problem to turn this formulation into a straight maximization problem, which might be amenable to a Bender's decomposition for solution.

Next they motivated the DA (Defender–Attacker) model by means of the DAD (Defender–Attacker–Defender) models. These tri-level models are very difficult to solve, so they often resort to a DA model, of the form minimizing maximum costs. The problem with this approach is that the inner problem is often not linear. One advantage of the DA model is that it allows the inclusion of limited information.

The DAD model is a tri-level optimization of the form minimizing the maximum of the minimal costs. These can be solved (at times) using a Bender's decomposition, but work is still ongoing. They then give three examples in some detail: the use of an AD model of an attack on the U.S. strategic petroleum reserve, a DA model of border defense, and finally a DAD model on the defense of electrical power grids.

A Stackelberg game formulation was used to examine the interactions between the largest and medium sized firms in the electricity and emission allowance markets (Chen, Hobbs, Leyffer, & Munson, 2006). They solved this Mathematical Program with

Equilibrium Constraints (MPEC) through a smooth reformulation of some of the constraints.

In 2008, Scaparra and Church (Scaparra & Church, 2008) looked at both natural and manmade disasters but focused on manmade disasters formulated as a Stackelberg games and solved it as a bi-level min-max optimization problem, in a similar manner to earlier work by Brown and Carlye (Gerald Brown et al., 2005). They discuss how the problem is non-deterministic polynomial-time (NP) hard, and examined a tree search algorithm to solve the problem. They postulated three approaches to the problem: to do a decomposition, use duality (as done by Brown), or reformulation. Because their problem did not have a linear parameterized inner flow problem, they could not use duality, so they developed their own tree algorithm to solve the min-max problem.

These "leader–follower" games were used by Pita and Jain in their 2009 work examining the assignment of police to random patrols and check points at the Los Angeles airport using game theory (James Pita et al., 2009). They developed a software scheduling tool called ARMOR, which provided randomized strategies for the manning of these checkpoints and patrol activities. They sought to develop a method to overcome three specific challenges. The first was that potential attackers could, over time, observe the schedules of police at the checkpoints and patrols and plan their attacks with those schedules in mind. Second, the actual goals of the attacker could not be known with certainty. Finally, they wished that their randomization process would take the benefits of protecting different targets into account. After an extensive review of Bayesian Stackelberg games, they discussed a classic method of transforming such games of incomplete information into the normal form of the game through the use of chance nodes (the Harsanyi transformation). This method had the disadvantage of the game losing its compact structure, but it did allow the game to be solved. Their solution did not require the use of this transformation. It is based on several earlier works, most notably (Conitzer & Sandholm, 2006) and (Paruchuri et al., 2008; Paruchuri, Tambe, Ordóñez, & Kraus, 2006). Their ARMOR model used a maximization problem formulation, which considers the payoff matrices, probabilities of different types of attackers, and the possible strategies. The formulation as presented was a quadratic problem, which they then linearized through a change of variables.

This work was further developed as a scheduling tool in another paper by the same primary authors (J. Pita et al., 2009). They gave an excellent discussion of the existing methods of solving these sorts of leader–follower or attacker–defender games, but pointed out that one common element was the use of a priori assumption about the value function of the terrorists. They treated uncertainty as having three distinct dimensions: the uncertainty caused when the terrorist does not choose the optimal strategy, the uncertainty caused when the terrorist is unable to observe the defender strategy, and the uncertainty caused by the terrorists having different reward matrices than initially assumed. While the third type of uncertainty was incorporated into existing algorithms for solution of these types of games, they proposed new algorithms (based on MILPs) to deal with the first two types of uncertainty. These algorithms were discussed in some detail and then compared with human adversaries using “human-in-the-loop” experiments.

The Stackelberg game formulation was used in 2010 to develop a system for allocating U.S. Coast Guard security patrols in ports and harbors (An et al., 2012). One unique aspect of this work was that they did not assume perfect rationality on the part of decision makers. They also simplified the formulation to facilitate solution by grouping points for patrol areas to exploit dominance and equivalence. Other examples of using various methods to reduce the computational workload in solving bi-level optimization problems for the solution to Stackelberg games include new relaxation methods for mathematical program with equilibrium constraints (MPEC) (Steffensen & Ulbrich, 2010), and transforming the lower-level problem through a Schurs Decomposition of the MPEC to examine decisions relating to the natural gas market (Siddiqui & Gabriel, 2013).

A recent book on algorithmic game theory (Nisan, Roughgarden, Tardos, & Vazirant, 2007) has a very complete discussion of many different algorithmic approaches to solving games. This book focused less on the applications of games than the technical manner of determining analytic solutions to those games. Some of the methods discussed include the use of linear and nonlinear programming (as discussed earlier in this section). They presented a complete discussion as to the difficulty of solving these problems analytically. It provided a very complete review of many different approaches. Another recent book on multi-agent systems (Shoham & Leyton-Brown, 2009) provided outstanding examples on how to use agents in the modeling of many different situations, including various types of games described by game theory. The book provided practical examples that were coupled to various classes of games commonly seen in game theory applications.

Recently, Hong (Hong, 2011) also used network interdiction approaches to examine conflict on a network in which one player is delivering “goods” through a network and the other is delivering “bads” through the same network. He characterized this as a “flow game” in which the value of the flow is the net flow of “bads” into the target. Following this formulation, he recasts the problem into a discussion of min-cut algorithms, in a manner based on that originally used by Wood (Wood, 1993)

Another approach to finding solutions to Stackelberg games is founded in evolutionary biology. By casting a conflict between animals as a sort of “limited war” Smith and Price developed the concept of the Evolutionarily Stable Solution (ESS) (Smith & Price, 1973). This strategy can be simply expressed as “…a strategy that, if most of the members of the population adopt it, there is no mutant strategy which would give higher reproductive fitness.” (Smith & Price, 1973). One of the keys of this approach is that it involved the evolution of stable solutions over time through the development of strategies or behaviors. In this article, the strategies or behaviors were the types of tactics used (either lethal or non-lethal) by animals against each other. These ideas were further developed by Smith in *Evolution and the Theory of Games* (Smith, 1982). He defines a strategy as a specification of what a particular individual will do in a given situation. Populations are defined as groups of individuals who share a common strategy. He motivated the idea of the ESS as one in which, if a mutant utilizing a different strategy were to invade the population it would be at a disadvantage because the payoff of the prevailing strategy was greater than the payoff for the mutant strategy against all possible combinations of strategies.

These concepts were then taken from an infinite population to a finite population by Schaffer, who showed that an evolutionarily stable strategy is globally stable against invasion by any other mutant strategy (Schaffer, 1988). He also demonstrated the existence of both locally stable and globally stable equilibrium points. The ESS concept was applied to an economic context looking at profit maximizers under imperfect competition in markets (Schaffer, 1989). Further discussion of the ESS in an economic context appeared in *Evolution, Games and Economic Behavior* (Vega-Redondo, 1996). In particular, he discussed the centrality of the ESS in evolutionary theory, stating how it reflected a stationary situation a revolutionary process. He further showed that there was no available result that identified sufficient conditions to guarantee the existence of an ESS in general situations. He also provided examples of how an ESS may not exist (consider the game of "rock, paper, scissors," which has no stable ESS).

2.5 Related Discussions

The area of cyber infrastructure has become increasingly important, and the defense and analysis of information networks has been the subject of numerous articles in the past few years. In 2002 Haimes and Longstaff addressed this issue by their work on assessing the consequences of a terrorist attack (Haimes & Longstaff, 2002). They characterized the problem of analyzing attacks on infrastructure as highly complex due to both the spatial distribution of the physical components of these networks and the inherent nonlinearity of their interactions. They characterized interconnected infrastructures as structure-based (hardware, structures, and facilities) and human-based (institutions, organizations, culture, and languages).

They recommended the use of Hierarchical Holographic Modeling (HHM) as a way to capture most of the elements of risk. They also proposed a so-called Risk Filtering, Ranking and Management Method (RFRM) to manage a large number of risks and combine these using an eight-step process to filter and rank the risks prior to the application of a more conventional risk analysis method. They made specific mention of the fact that 90% of the national security communications in the U.S. are on the civilian Internet backbone, and they discussed in qualitative terms the risks inherent in this vulnerability.

One key element of many of the studies on terrorism has been the need to adequately describe the terrorists' motivations so that their actions can be predicted. Goldstein (Goldstein, 2006) proposed the use of agent-based models as a useful tool in the development of understanding about terrorists' actions. He maintained that better predictions would lead to better estimates as to the possible actions that terrorists might take. In particular, the article discusses the use of various Knowledge Management techniques such as named entity extraction and linguistic analysis of unstructured news input data to develop identities of terrorists, their roles and the relationships in the network. By intense study of the specific reactions of individuals in a given area or network, the author maintained that the ability to predict actions and results was greatly enhanced. This work has great potential to assist in the development of realistic probability of attack in different scenarios by allowing the evaluation of multiple "spawns" of the life of a terror cell. However, this article did not specifically address the use of agent-based models for infrastructure resource allocation, nor did it directly address game theory.

Another aspect of the analysis of terrorist attacks is the proper estimation of the impact of the attack, not only upon the actual infrastructure, but the secondary effects that the disruption of the services delivered by the infrastructure to the larger economy. In a

2007 article, Rose, Oladosu, and Liao examined the effects of the economic impacts of a major terrorist attack to the power distribution system of Los Angeles (Rose, Oladosu, & Liao, 2007). They concentrated on indirect effects and resilience. Within the context of this article, they defined resilience as including the ability of the customer (region) to adapt and react to the outage. The indirect effects increase the damage effects by disrupting the flow of the goods and services outside the area immediately affected by the outage. Within their effort, they focused on business interruption and its economic effects, not on the loss of life and damage to property that such an attack would entail. This explicit inclusion of the economic interdependence up and down the supply chain, radiating out from the affected area, was the heart of the contribution of this article. In order to accomplish this goal, they developed a regional economic model, based on a simulation that has been used for modeling the regional economic impact of other events and policies. They simplified the productive capacity of the area to 33 productive variables and made allowance for resilience (alternate energy production, backup generators, transfer of production to an unaffected area, etc.). Their research showed that the losses of a 2-week power outage in Los Angeles caused a loss of about 20.5 billion dollars, but that resilience could recoup about $2.8 billion of that loss. They also believed that this work had applicability to the reduction of damage from natural disasters as well.

Another recent article focusing on the use of game theory for network intrusion by published by Kantzavelou and Katsikas (Kantzavelou & Katsikas, 2010). They focused on the insider threat, which by a 2007 survey accounts for 34% of all intrusion events. They formulated their analysis as a game between the defender (the Intrusion

Detection System, or IDS), and the attacker. The attacker's actions were characterized as one of four different possible actions: Normal, Mistakes, Pre Attack, and Attack. The IDS also had four distinct actions: Continue, Give recommendation, Warning, or Stop User. The game is based upon a "two-player non-cooperative game" in which the defender has imperfect monitoring, because the IDS does not know what the defender is doing. They go on to show that the problem is NP complete, and very difficult to solve analytically.

2.6 Synthesis

After I reviewed this literature and examined the fabric of the combined analysis, several facts became quite clear. Most important was that there was a fundamental difference between risk analysis for natural disasters and risk analysis for terrorist attacks. In order to model events to some level of accuracy that would provide useful insight for the decision maker, the analysis had to take into account the interplay between the thinking, adapting attacker and the network. The attacker would not blindly charge in like an earthquake or tornado and attack indiscriminately; the attacker would seek to find weak spots and take advantage of them.

Second, the role of limited information had to be included in any realistic modeling of the situation. While the attacker would have ample time to study the network and make good (although not necessarily perfect) decisions based on the resource allocations of the defender, the defender did not have such flexibility. Once the allocation was made, it was essentially "fixed." While such things as changing schedules or rotating the location of reserves could have some effect upon the attacker, the defender did not have the flexibility to adjust as the attack was in progress.

These two facts led to the conclusion that a GT approach based on the "attacker–defender" model was highly desirable. However, for any problem of a size to be useful beyond a simple classroom example, the analytic solution of such a problem was extremely difficult. The deterministic solution approach using linear programming was very attractive, but implementation of such an approach required a well-defined objective function. The ability of such an approach to deal with multiple objectives could be problematic. The goal was to have a system that would allow for differing terrorist objectives. Some might wish to attack the target that would do the most damage to the network, which others might wish to attack one for maximum publicity. Some terrorists might tend to choose a lightly defended target to increase their chances of survival, while others might be deliberately planning on a "suicide attack" without regard for the chances of survival.

After reviewing multiple approaches, the most promising new approach seemed to be the use of an agent-based model to develop solutions in an evolutionary manner. This approach was consistent with the concept of the ESS discussed previously.

There was an excellent tutorial on agent-based simulation (Macal & North, 2010) that proved to be pivotal in both the decision to proceed with the agent-based model and with the general outline of the final form of the model. According to their structure, any Agent-Based Simulation (ABS) has three distinct elements: a set of agents, which includes their attributes and behaviors; a set of agent relationships that define the topology of their interactions; and the agent's environment, which includes how the environment affects the agents and how they interact with the environment. The agents only have local information available to them, and while they interact with other local agents, they may not interact with all the agents. They can adapt, have interactions with other agents, and sense the behaviors of other agents in their vicinity.

Based on this article, it was decided that a method that could identify favorable evolutionarily stable solutions would have the potential to provide good decisions for resource allocation under ARA. An evolutionary, agent-based approach offered an excellent way to model the actions of terrorists attacking a network. While many approaches had been suggested and applied to problems of varying sizes, the process of modeling the behavior of terrorists with multiple goals in larger problems was one that had room for exploration.

Chapter 3: Problem Delineation and Approach

3.1 Elements of the Problem

After I reviewed the available literature, it became clear that the problem of resource allocation under strategic uncertainty required that the actions of both the defender (the entity assigning resources that protected the network) and the attacker (the terrorist or the entity assigning resources that sought to damage or reduce the functioning of the network) be taken into account in a dynamic manner. The fundamental difference between an act of nature, such as a flood or storm, and a terrorist was that the terrorist had certain goals in mind and was able to take the defense measures of the network defender into account. After evaluating different approaches, the GT approach of casting the problem as an "attacker–defender" (also called "leader–follower" or Stackelberg) game was considered to be the best way to capture the adversarial relationship between the attackers (or terrorists) and the defenders.

From the reviews of recent literature examining the characterization of terrorist goals and objectives, two issues that needed to be considered became clear. The first was that terrorist organizations often had multiple goals. They might seek to inflict maximum damage upon their target with overwhelming force at that point, or they might attack with a weaker force that was indifferent about its own casualties. They might strike a weakly protected but non-vital target in order to "make a statement" with minimal risk to themselves, or they might seek to gain the maximum publicity and notoriety from an attack. They might seek in-depth damage at a critical node or widespread but more superficial damage at multiple nodes. The second issue to consider was that determining the actual probabilities for these different goals was very difficult and subject to much interpretation and error. Therefore, the technique developed had to be one that made as few assumptions about the prior goals of the attacker as possible. Further, the terrorists might not be homogeneous with regard to their goals, tactics, or tolerance to casualties, and the method chosen has to allow for this variability. These differences might lead the terrorists to concentrate their resources upon different locations in the network.

The valuation of network structures also needed to be addressed. For simplicity and ease of illustration, the examples focus upon values of the nodes of the network. There is nothing inherent in the technique that requires focusing upon the nodes: The arcs, or indeed both the arcs and the nodes, could just as easily have been evaluated. However, focusing on the nodes allowed the reduction of the physical location to a single point, rather than a locus of points in a line as an arc would have required. Each node of the network should have a value of some sort or another. However, there could be multiple values. For example, the destruction of the Liberty Bell or the Washington Monument would have little impact upon the functioning of the country but would have a great psychological and publicity impact upon the country. The 9-11 attack on the Pentagon had a direct (although short-lived due to redundancy) impact upon the operation of the country but a long-lasting impact upon our thinking. Continuing on in this same direction of thought, a successful attack that demolished a power switching substation might have a considerable economic impact upon an area but very limited public relations impact in its own right (not considering the impact due to loss of power). The solution system chosen had to be flexible enough to deal with multiple "values" for the nodes and be flexible with respect to how those values were determined. While the examples chosen had both attacker and defender holding consistent values for the nodes, the method does not depend upon that sort of reflexive commonality of valuation.

Another part of the problem with this analysis is that data are sparse and difficult to obtain. Further, no one would recommend letting terrorists repeatedly attack a target just to obtain experimental data, and any such data would obviously be tainted by the fact that the terrorists were not acting as they would under normal circumstances. Yet, it was essential to have a well-defined metric for the success of the allocation method and the ability to investigate the technique in some reasonably comprehensive and analytic manner. As in so many other conflict situation analyses, simulation seemed to offer the best mix of fidelity and rigor.

With these thoughts in mind, a working hypothesis was formed: When facing an opponent who has knowledge of the defender's resource allocation and can plan accordingly, an allocation approach based on a GT allocation of resources will perform as well as, or better than, an allocation based upon a pure Probabilistic Resource Allocation (PRA)-based approach, which only takes the value of the nodes under consideration without considering the actions of the attacker.

3.2 Solution Approach

The problem is postulated upon determining the best way for a defender to allocate resources to defend nodes of a spatially distributed network against a reasoning adversary. The attacker and defender are provided with a quantity of resources, which are then allocated to various nodes of a network. For maximum flexibility and clarity of presentation, these resources are considered to be nonspecific and fungible. They could be thought of as fiscal resources to be used to purchase offensive or defensive value at the nodes. Examples of these "purchases" might be fences, alarm systems, guards, intelligence, weapons, or explosives.

Each of the nodes has two numeric quantities associated with it that serve as expressions of the value of the node: one that reflects the importance of the node in maintaining flow through the system, called a node value, and a quantity called the public relations value, which reflects the value the terrorist would place on the publicity gained by an attack to this particular node. In the examples presented in this dissertation, the node value either reflects the quantity of flow through a node in the network when it is flowing at maximal capacity, or a combination of the maximum flow through the network and the value of lost flow if the node is interdicted, or removed from the network. It is highly desirable that the solution system be totally agnostic with respect to the method used to determine the value of the nodes. In a well-understood and studied network, these values will most likely be a mixture of historical usage, flow modeling, and engineering/operator judgment regarding the difficulty of replacing the functionality of each node. In order to keep this work as generic and non-sensitive as possible, although some of the networks were based upon real-world examples, the node capacities were chosen for illustration only. Two methods were used (sometimes in combination) for the determination of network flow using a combination of two linear programming models. The first is a network flow model, which determines the flow through each part of the network that will maximize the flow through the system. The second is a network interdiction model, which determines the identity of the nodes in the network that, if disabled, will cause the greatest reduction in the flow through the network. These two models provide input data for the determination of the values of each node of the network. They are discussed at length in Chapter 4.

The public relations value (PR value) of the node reflects the point of view of the terrorists. These values are only considered relative to each other, so the actual range they encompass is purely arbitrary. In the course of the modeling, the terrorists' choice of target node may be affected by the PR value of the node. Methodologies for the proper assignment of these values would depend upon available intelligence and profiles of terrorist motivation. The actual method used to determine the PR value the terrorists place upon a particular node is not germane to this research; the only impact upon current work is the actual PR value assigned to each of the nodes. In this dissertation, the PR value was arbitrarily assigned to the nodes and sometimes varied by the scenario under analysis. Within the context of this work, the PR value can best be considered as an alternative value for the nodes, which may affect some terrorists and defenders but not others.

Given the complexity of finding the solutions to even a small Attacker–Defender game using analytic approaches, the approach taken is an evolutionary approach to finding the solution (allocation of defenders to nodes) that is most advantageous to the defenders. This problem is solved through the use of an agent-based model, referred to as the allocation model. Within the context of this work, the term "strategy" refers to the overall allocation of resources: the quantity of resources the attacker or the defender allocates to each node in the network. The term "tactics" refers to the rules governing the behavior of the individual agents. This agent-based model provides combinations of strategies that coincide with the ESSs for the terrorists and defenders. Each of these strategy combinations (the number of agents located at a particular node) results in an allocation of resources to each node. By choosing the allocation of resources (number of resources assigned to each node) that result from the most favorable ESS for the defender, a resource allocation for the defender is developed. The overall approach is shown in Figure 2.

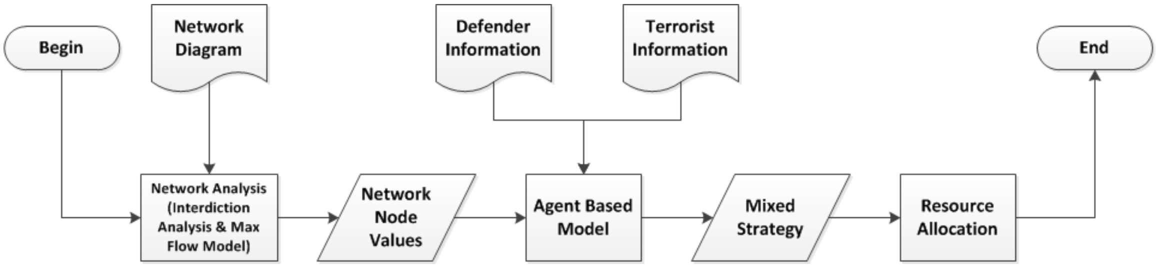

Figure 2 Overall Solution Approach

The network diagram and capacities of the nodes are the primary inputs to the network analysis, which determines the node values as discussed previously. Information on any existing non-movable resources already deployed to specific nodes, along with the number of resources to be allocated, are then used along with the network characteristics by the agent-based model. This model develops an ESS for attacker and defender, which consists of an allocation of resources by the attacker and the defender to each node. The resource allocation that provides the greatest number of victories by the defender (the largest number of nodes which remain in operation) is considered to be the defender's resource allocation from the model.

Note that the system is totally agnostic with respect to the method of evaluating the values of the nodes of the network. It does not require that the attacker and defender have the same values for each node. Indeed, the evaluation could be equally applied to arcs as well as nodes. It also is not limited to a particular number of values. Although this research is predicated upon two different values for each node (flow value and PR value), there is nothing in the technique limiting it to only two values per node.

3.3 Analysis Approach

The key to proving the usefulness of this technique was to show that, as measured by some meaningful metric, the agent-based evolutionary method performed as well as, or better than, a simple PRA-based approach. The available data sets in the public domain are very sparse, and those that were available were not amenable to this research for there was no way to determine if the "solution," or allocation of resources, actually represented the "best" that could be done or was just a single realization of a wide number of possible solutions, with varying degrees of success. After investigating the available data, it was determined that the best and most meaningful method to test the hypothesis was to employ simulation.

The method used to determine the various potential stable points (the ESS) of the systems involved using an agent-based simulation in which the agents on both sides represented resources applied to nodes. Two agent-based models were used in this research. In the first model, both sides (the attacker and the defender) evolved over time and then stabilized at a point at which changes by one side alone could not change or improve the outcome. This model is called the allocation model. The result of the allocation model is an allocation of defensive resources for each node in the network. In order to see how the allocation derived in this manner fared against a PRA-based allocation, another agent-based simulation (the evaluation model) was developed. In this second simulation, only the attackers (terrorists) were able to evolve and change over time. The allocation of resources was fixed at the start of the simulation run and was invariant during the run and between runs. In this manner, the full possible range of strategies by the terrorists could be examined, and a probability distribution of the potential outcomes as the terrorists evolved could be empirically derived from the results of multiple runs of the model.

The metric chosen for the final measure of success or failure was the number of nodes that were successfully defended (defender victories). The technique involved three different model runs. The allocation model was run to determine the best ESS from the perspective of the defenders. The resource allocation from that model was considered the GT allocation. A PRA-based allocation for defenders was determined using the node values, which were derived from analytic means. Then the evaluation model was run twice, first to determine the empirical probability distribution of the number of defender victories when using the defensive resource allocation from the PRA-based model. The model was then rerun to determine the empirical probability distribution for the number of defender victories using the GT allocation for defenders. The two different allocation models were also tested under several different conditions to examine the sensitivity of the results to certain assumptions regarding public relations values and the importance of public relations victories to the terrorists and different distributions of PR value with the same network flow value. For each of these runs, the empirical cumulative distribution of the number of defender victories was compared.

Chapter 4: Formulations and Models

4.1 Overview

Several different formulations and models are used throughout this research. Two linear programming formulations were used to assist in determining the value of the nodes to the network. The first of these was a network flow formulation, which sought to determine the flow through each node of the network when it was operating at its maximum capacity. The second linear programming formulation was a network interdiction model, which determined which network nodes would, if eliminated, cause the greatest reduction in the flow through the network. Linear programming methods were chosen as representative of the typical means of evaluating the importance of a network node to its overall functioning. The actual means of determining the value of the network nodes has no bearing on the agent-based simulation approach to resource allocation.

Both of these formulations were then implemented and used to determine network node value in one of the examples. In one of the three networks evaluated, the node values were chosen to be a combination of the results of the two network models (the network flow formulation and the network interdiction formulation). In the other two examples, only the network flow model was used. The methods chosen were dependent upon the topology of the networks. The use of multiple valuation approaches also showed that the results of the agent-based simulation were not dependent upon any one network evaluation methodology.

As discussed previously, two different agent-based simulations were developed. The first one, the allocation model, allowed the agents for both the attacker and defender to evolve and determine their optimal strategies. The second model, the evaluation model, was very similar except that it did not allow the defender agents to evolve. They were assigned to a node at the start of the model run and never varied. However, the terrorist agents did evolve and seek to do their best against the allocation. The structure of these two models will be presented. The final sections of this chapter discuss how the formulations, models, and analysis were implemented.

4.2 Network Flow Model Formulation

The network flow model is one that calculates the maximum flow through a capacitated network, as described by Bazaraa et al. (Bazaraa, Jarvis, & Sherali, 1990). The model formulation used in this work is presented below:

Indices

J	nodes
K	arcs

Sets

K_{ja}	set of arcs entering node *j*
K_{jb}	set of arcs leaving node *j*
J_t	set of terminal nodes in the network

Variables

x_k	flow through arc k

Parameters

c_k	capacity of arc k

Model

$Max\sum_{k} x_k \quad \forall k_{j,a}, j \in J_t$ flow into terminal arcs

Subject To

$x_\beta - x_\alpha = 0 \quad \beta = K_{jb},\ \alpha = K_{ja} \forall j \notin J_t$ flow balance constraints

$x_k \leq c_k \quad \forall k$ arc capacity constraints

$x_k \geq 0 \quad \forall k$ non-negativity constraint

The model seeks to maximize the flow into the terminal nodes of the network, ensuring that the flow into and out of every node is balanced, and ensuring that the capacity of each arc is not exceeded. By using the flow through each node as an input to the value of the node, the "bottlenecks," or nodes that are operating at capacity, are identified. These are the critical nodes that will affect the ability of the network to maintain flow in a damaged state.

4.3 Network Interdiction Model Formulation

Network interdiction is related to—but more computationally difficult than—a relatively simple maximum-flow model. Network interdiction models require the extensive use of binary variables, which increase the computational difficulty of the model. Without loss of generality, it is assumed that the network is a directed flow capacitated network with source and sink nodes located on the outer boundaries of the network. This assumption allows for a compact and relatively efficient network interdiction model. The following formulation, taken from Wood (Wood, 1993), is used to determine the most important arcs to interdict.

Indices

J nodes

K arcs

Sets

Ns nodes that are sources

Nt nodes that are sinks

A_{st} set of arcs incident to a node in either Ns or Nt

$\bar{A}_{st}$ set of arcs that are not incident to a source or sink node (A Complement)

Variables

α_J variable associated with nodes

β_k variable associated with arcs. *(Those whose value is 1 define the minimum capacity cut)*

γ_k variable associated with arcs**.** *(Those whose value is 1 define the interdicted arcs)*

Parameters

u_k capacity of arc k

r_k resource cost for cutting arc k

R total resources available for the interdiction

Model

$Min\sum_{k} u_k \beta_k$ capacity of unbroken forward arcs

Subject To

$\alpha_s - \alpha_t + \beta_{s,t} + \gamma_{s,t} \geq 0$ for all arcs that connect to a source or sink node

$\alpha_s - \alpha_t + \beta_{s,t} + \gamma_{s,t} \geq 0$ for all arcs that do not connect to a source or sink node

$\alpha_t - \alpha_s + \beta_{s,t} + \gamma_{s,t} \geq 0$ for all arcs that do not connect to a source or sink node

$\alpha_j = 0$ for all source nodes

$\alpha_j = 1$ for all sink nodes

$\sum_k r_k \gamma_k \leq R$ Resource constraint for interdiction activity

$\alpha_j, \beta_k, \gamma_k$ Binary

The results of this model were used as inputs to help determine the node values for some of the example networks. Due to the topology of a network, sometimes the removal of a single node did not affect the overall flow through the network because of the inherent redundancy in flow paths. However, this formulation can be used to determine the identity of the node that, if removed, will have the greatest effect upon the flow through the network.

4.4 Agent-Based Model Structure

As described by Gilbert (Gilbert, 2008), agent-based models differ from other equation-based computational models in the manner of focus. In "conventional" models, or models that are not agent-based, the entire system is defined from the top down with a series of equations that represent all the interactions. The interactions must be completely specified and probabilities assigned to all possible events. However, agent-based models concentrate on developing a simple set of rules that both prescribe and constrain the actions of individual agents. These simple rules also prescribe how the agents interact with each other and their environment. The complexity of the interactions at high levels of the model are not specified but are allowed to develop from the simple interactions of the agents.

Macal and North (Macal & North, 2010) provided a useful framework for the development of agent-based models. According to their construct, the three distinct elements in an agent-based model are defined as follows:

1. A set of agents, their attributes and behaviors.
2. A set of agent relationships and method of interaction: an underlying topology of connectedness defines how and with whom agents interact.
3. The agents' environment: agents interact with their environment in addition to other agents.

Each of these three elements is discussed in the following sections.

4.4.1 The Set of Agents

Agents can have differing characteristics and rule sets that govern their behavior. Agents with the same characteristics and rule sets (but not necessarily the same values for those characteristics) are called "breeds." The models for this research have three distinct breeds of agents: Terrorists, Defenders, and Nodes. Although the Nodes are agents, they only serve to describe features in the environment for the Terrorists and Defenders. They provide a location in the Cartesian grid that constitutes the agents' "world," focal points for interaction and data gathering. The behaviors of the defenders and terrorists are called "tactics." These strategies all relate as to how the agents pick their "target," which is the individual node of the network they choose to attack or defend. Each agent has a rule that defines its tactics. This rule determines how the particular agent picks its target.

Initially, these tactics are assigned randomly to all terrorists and defenders. The actual probability that any one agent has a particular tactic is a parameter that is established before the start of the model runs. The terrorists' and defenders' physical location on the Cartesian grid is also randomly assigned each turn. Over time and continuing repetitions, the population of less successful agents is reduced through attrition. Eventually, the percentage of the total agents that follow each of the target allocation strategies is the percentage of times each of these strategies would be employed in the optimal mixed strategy for the ESS of the game. The actual number of agents that cluster at each node also tends to stabilize as the strategies stabilize. The proportion of agents at each node thus represents the resource allocation that should be assigned to the node. These proportions of agents are referred to as strategies because they refer to the overall allocation of resources for the attack or defense of each node in the network.

This fact provides one of the key strengths of the model. The behavior of the terrorists (their tactics) can be described in general terms as to how they might pick their target or their goal. It is not necessary to determine the probability that they would strike at a particular target. In a similar manner, it is not necessary to estimate the probability that a given defender will defend any one single node of the network: The choice of strategy can be described more generally in terms of behavior and values. The model then allows a correspondence to be established between these behaviors and the actual decision the planner must make, i.e., how many resources to devote to attacking or defending each node.

The behaviors of the agents (terrorists or defenders) can be one of the following:

1. *Attack or defend the most valuable node in the network.* This choice reflects an "all or nothing" valuation by the agent. No consideration is given the physical location of the terrorist or the most valuable node.
2. *Attack or defend the closest node in the network.* Each node has been assigned an unchanging "location" in the XY plane. All attackers and defenders are initially distributed randomly over the entire XY plane. This choice reflects a random spatial distribution of effort, independent of the value of the node to the network.
3. *Attack or defend the least defended node within a specified distance.* This strategy reflects the desire to attack a network at its weakest point, where one is more likely to achieve success and survive the encounter. For the defender, it reflects the desire to evenly distribute resources without regard to the value of the node in question. For both sides, the specified distance reflects limited information available to the individual agent.
4. *Attack or defend the most valuable node within a specified distance.* This reflects the desire to do the most damage in the attack, or to defend the most important part of the network. For both sides, the specified distance reflects limited information available to the individual agent.
5. *Attack or defend a random node.* This strategy reflects a desire to attack or defend without regard to the value of a node, the number of defenders, or the relative strength or weakness of the node.
6. *Attack or defend the node with the highest public relations value within a specified distance.* This represents the desire to gain publicity for a cause through completing an attack, without regard for the actual value of the target node to the

network or the defenses of this node. For the defender, it represents the desire to counter this type of action.

These behaviors, or tactics, of the terrorists and defenders are considered representative of the actions of actual decision makers who control the terrorists or defenders. These tactics could be modified to represent different behaviors. Although the tactics in this version of the model are the same for both attackers and defenders, there is no requirement for them to be reflexive in this manner. It is also important to note that the "decision-making horizon" of the agents is the range that encompasses their field of vision to choose nodes. The horizon represents limitations on the information available to the agents.

4.4.2 The Set of Agent Relationships

The agents (terrorists and defenders) interact based on their location. Each terrorist or defender will be assigned to a single node based upon their geographic location within the Cartesian grid and their strategy. Once the strategy for each agent is determined, and their "target" node is defined, they move to that node. First, the defenders determine their target node and move to it. After the defenders have been assigned in this way (the resource allocation of the defenders fixed), the terrorists then determine their targets and move to their target nodes.

The sequential "movement" of agents to their respective target nodes is the key to finding the ESS to the "leader–follower" or "attacker–defender" game. The defender, or leader, must make an allocation decision without reference to the actions of the terrorist. The only information available to the defender is the value of the various nodes of the network and the history of how successful that particular agent has been in the past. The attacker (terrorist), or follower, on the other hand, has not only this valuation information but also knows the current disposition of defenders. By managing the visibility of knowledge in this manner between the terrorists and defenders, the allocations that result will represent the ESSs of the overall game.

After all the agents have determined their strategy and moved toward their target node, the number of terrorist agents and defender agents are counted at each node. The agent breed (terrorist or defender) with the largest number of agents at the node is considered the winner at that node. During the terrorists' allocation of resources, if the terrorists have numbers far in excess of what is required for a victory, the strategy of the agent that is attempting to move to the node is randomly reassigned, resulting in its movement to a different node. Once the number of terrorist agents at a node is greater than twice the number of attacker agents at that node, the strategy for any further agents attempting to move to that node for that iteration is randomly changed so that they will move to another node. Agents (terrorists and defenders) have the knowledge of the location of the network nodes, and terrorists have knowledge of the number of defenders at each node. Defenders know how many defenders are at any particular node. Both breeds know the value and the public relations value of all nodes. Note that defenders move to their nodes first, and they never have knowledge of the attackers.

This program seeks Evolutionarily Stable Solutions, so the terrorist and defender agents evolve over time. If an agent "evolves," its tactics change. After a number of turns, each agent is evaluated for the number of times it has been present at a node where it has been successful and the number of times it has failed. A percentage of the least successful agents then have their strategy randomly reset to a new strategy. This resetting reflects the evolution of the agent. In effect, the least successful agents are "killed off" and replaced by other agents. Over time, agents with the least successful tactics form a smaller percentage of the population of agents. All the success and failure history for each agent is then reset to zero, and the simulation continues until the next evolution evaluation occurs, based only on what has occurred with the set of newly evolved agent tactics from the previous evaluation. Additionally, at this time a certain percentage of all the agents, regardless of their success or failure record, have their strategy randomly reassigned. This randomization helped to ensure that all parts of the solution space were explored by representing totally random changes that were not driven solely by the fitness, or success, of any one agent.

4.4.3 The Agents' Environment

The agents' environment is a Cartesian plane. The nodes of the network are assigned coordinates on this plane that correspond to the physical locations of the nodes. The nodes are connected by arcs, which begin and end at specified nodes. In the current version of the model, the arcs do not affect the decisions of the terrorists or defenders. They only serve to affect the value of the nodes through their effect on flow, as detailed in the maximum-flow model and the interdiction model. All resources are directed at the nodes, not the arcs. Future work could address a formulation in which resources are deployed for the defense of the arcs.

4.5 Agent-Based Models Used in this Research

Two distinct agent-based models were developed for this research. In the first model (the allocation model), the behaviors (tactics) of both the defending agents and the terrorist agents are allowed to evolve over time. This allocation model is used to develop the solutions corresponding to the ESS for the game. The second (evaluation) model is used to compare a fixed (non-evolving) defensive strategy of resource allocation to an evolving terrorist strategy. The allocation of defensive resources, node by node, is assigned as an input to this model, and the terrorists evolve to fight it. There are no defender agents in this model: only terrorists. The terrorists behave exactly as just described for the allocation model. The only difference between the two models is that the defensive resources at each node are invariant over time in the evaluation model. The terrorists evolve, changing their strategies to inflict maximum damage upon the network. However, the number of defenders at each node never varies and is reflective of the input allocation being evaluated. This model is used to compare various resource allocations by defenders (solutions) to determine how well they protect the network as the terrorists evolve to attack the network and its resources. This evaluation model provides insight into how well the resource allocation performs against an informed and evolving terrorist threat.

Both of these models are built on a similar structure. The discussion follows the operation of the allocation model. The evaluation model follows exactly the same sequence except that the defender resources are fixed at each node, and this allocation never changes. Differences between the two models will be pointed out during the discussion. The sequence diagram, Figure 3, provides a visual representation of the flow of events in a single repetition of the model.

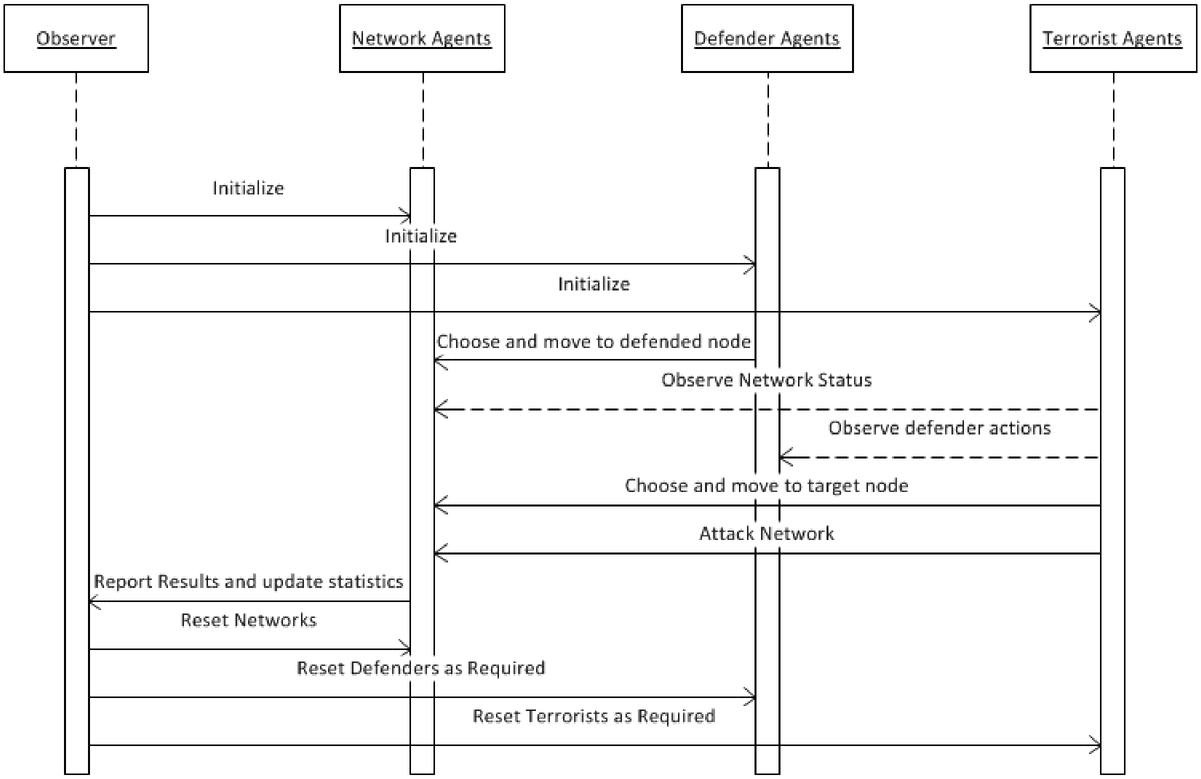

Figure 3 Sequence Diagram for Allocation Model

The observer tracks all events in the simulation, calculates statistics, and directs events as shown in the diagram. Network agents refer to the nodes. When the nodes are initialized, their values (location, value, and PR value) are set and do not vary. They also may have an inherent defense value that will not vary throughout the run of the simulation. For the allocation model, these values represent resources that are already deployed to the nodes and are not subject to reassignment based on the results of the model. For the evaluation model, this inherent defense will reflect the final allocation of resources under evaluation. These defenses could be the result of a PRA-based allocation or the result of an ABS analysis determining the GT allocation.

When the terrorists or defenders are first initialized, their tactics are randomly assigned, as is their location in the XY plane. A single cycle begins when the defenders move to the node that corresponds to their current tactics. Once this movement is complete, the terrorists move to the node that corresponds to their current tactics. Again, it is important to note that each terrorist or defender agent represents a single unit of resources. The resolution of the conflict at each node depends upon the number of terrorists and defenders at each node.

The resolution mechanism for the conflicts at each node assigned victory to the breed of agent (terrorists or defenders) that has the most resources at that node. After the conflict is resolved and the win/loss counters of each terrorist and defender are updated, a new cycle begins by randomly reassigning locations to every terrorist and defender. After a number of cycles, the total number of victories is calculated for each terrorist and defender. A percentage of the agents with the least victories have their strategies randomly reassigned to a new strategy. Both the percentage of the agents to have their strategies reassigned and the number of cycles between these evaluations can be set as parameters in the simulation. Over time, the number of agents with the less successful strategies will be reduced, and the percentage of agents with each tactic—and the overall numbers of agents at each of the nodes—will approach an ESS for the network.

One inherent problem with any evolutionary approach to finding the ESS is that it is impossible to determine if all of the potential points have been found. Therefore, in no sense is there a claim that this technique provides the “best” solution to the resource allocation problem. No matter how many simulation runs are performed, there is a chance (albeit small with a large number of runs) that a better allocation could exist.

However, it will be shown that in the three example networks presented this approach provides an improvement over a more traditional PRA-based approach when basing the allocation on the ESS from 400 runs, which provides for the successful defense of the largest number of nodes.

4.6 Model Implementation

The two linear programming models were implemented using the General Algebraic Modeling System (GAMS) (GAMS Development Corporation, 2009) and the GUROBI (Gurobi Optimization, 2010) solver. The two agent-based simulations were implemented in NetLogo Version 5.03 (Wilensky, 1999). NetLogo is a freeware, dedicated agent-based simulation platform. It has a rich environment of built-in functions and operators that allow for rapid development of agent-based models. Many guides for programming in NetLogo are readily available, but in this development work the book, *A Field Guide to NetLogo*, by Scott and Koehler (Scott & Koehler, 2011), was particularly useful in the design and implementation of the model code.

The allocation model was developed first. The evaluation model is the same as the allocation model except that resources assigned to defend each node are an input to the model and never vary. The terrorists behave as they do in the allocation model except that the conflict is resolved based on the number of terrorists and the resources assigned to defend each node. The details of the network configuration (node coordinates, node value, node PR value, etc.) are configured via easily modified text files. Most of the run data are entered via a customized Graphical User Interface (GUI), which also provides run time graphical status updates and diagnostics. The GUI, pictured in Figure 4, presents a visual representation of the network being analyzed, as well as diagnostic information regarding the state of the model. It is possible to "step" through the model, one repetition at a time, or to run it for a specified number of repetitions. For runs to be used for analysis, the NetLogo environment has a feature called "Behavior Space," which automates sensitivity analysis by varying parameters over a specified range for a series of runs. These results are then automatically entered into a Microsoft Excel spreadsheet.

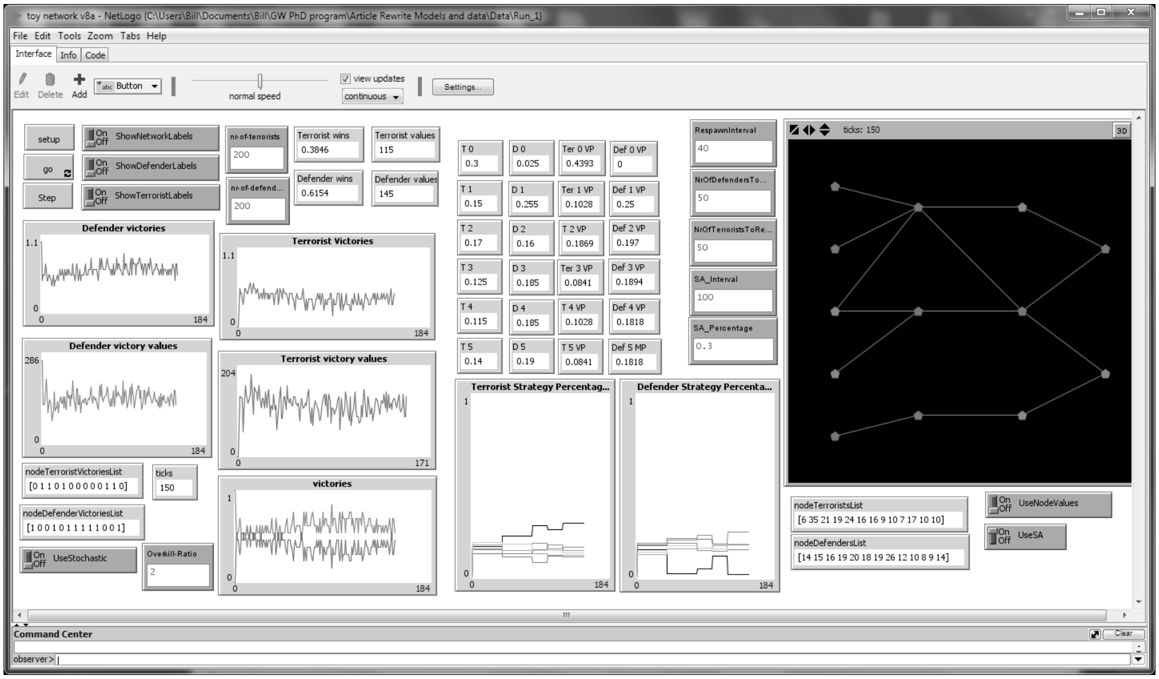
toy network v8a - NetLogo {C:\Users\Bill\Documents\Bill\GW PhD program\Article Rewrite Models and data\Data\Run_1}

☑ view updates: continuous

nr-of-terrorists: 200
nr-of-defend...: 200
Terrorist wins: 0.3846
Terrorist values: 115
Defender wins: 0.6154
Defender values: 145

T 0: 0.3	D 0: 0.025	Ter 0 VP: 0.4393	Def 0 VP: 0
T 1: 0.15	D 1: 0.255	Ter 1 VP: 0.1028	Def 1 VP: 0.25
T 2: 0.17	D 2: 0.16	T 2 VP: 0.1869	Def 2 VP: 0.197
T 3: 0.125	D 3: 0.185	Ter 3 VP: 0.0841	Def 3 VP: 0.1894
T 4: 0.115	D 4: 0.185	T 4 VP: 0.1028	Def 4 VP: 0.1818
T 5: 0.14	D 5: 0.19	T 5 VP: 0.0841	Def 5 MP: 0.1818

RespawnInterval: 40
NrOfDefendersTo...: 50
NrOfTerroristsToRe...: 50
SA_Interval: 100
SA_Percentage: 0.3

ticks: 150

nodeTerroristVictoriesList: [0 1 1 0 1 0 0 0 0 0 1 1 0]
ticks: 150
nodeDefenderVictoriesList: [1 0 0 1 0 1 1 1 1 1 0 0 1]
Overkill-Ratio: 2
nodeTerroristsList: [6 35 21 19 24 16 16 9 10 7 17 10 10]
nodeDefendersList: [14 15 16 19 20 18 19 26 12 10 8 9 14]

Figure 4 Allocation Model GUI

4.7 Analysis Implementation

The analysis of the model data runs used several different software tools. All model runs used for analysis used the “Behavior Space” option, which allowed selected output to be written to Excel-compatible Comma Separated Value (CSV or .csv) files. Depending upon the size of the network under analysis, anywhere from 4 to 16 of these files were required to capture 400 simulation runs of 2000 repetitions each. The disparity in the number of files required is due to the increasing amount of data as the number of network nodes increased. The software used in the analysis had limitations on the total size of the input files that could be imported. Summary tables of data results were prepared using Excel.

The results of the Behavior Space runs were combined using a custom MATLAB script, which read in all that data, combined the results from the various .csv files, and then organized them into a single large data structure. All subsequent analysis was done on this data structure via a series of custom MATLAB scripts. These scripts performed all the data manipulation, and presented the data in formats most amenable to analysis and interpretation. All of the histograms, Cumulative Distribution Function (CDF) plots, and Quantile-Quantile (QQ) plots that appear in the individual analysis sections, as well as the calculation of statistical tests such as the Lilliefors tests for normality, were done using MATLAB.

Initially, the allocation model was run under certain specified conditions (which were dependent upon the network under analysis), and the results (number of nodes where the defender was victorious) were evaluated. A detailed examination of the actual number of defenders at each of the nodes was then completed. Usually, there were very slight differences in the number of defenders at each node within the subset of results, which had the same total number of defender victories. These numbers were averaged and then normalized to the total number of defenders (amount of resources) available. This set of ordered pairs (node ID and number of defenders) was then considered to be the GT allocation for resource allocation under the specified set of conditions. This process was continued for each of the various conditions under consideration until the GT allocation was established for each of the sets of conditions under examination.

The number of defender resources was then allocated based on a straight PRA-based approach, with the percentage of resources allocated to each node equal to percentage of total network value assigned to each of the nodes. This was normalized to the total resources available. At the end of this process, for each set of conditions, there were two sets of resource allocations to each node: one based on the results of the evolutionary agent-based simulation (the GT allocation) and one based on the "fair share" allocation of resources considering node value (the PRA allocation). Each of these allocations was then used as the defender resource allocation for the evaluation model, when 400 separate runs of 2000 iterations were performed against each allocation under the specified conditions. These results (within each run) also tended to stabilize. The empirical distribution for the number of defender victories was developed for each allocation method (GT and PRA) for their respective conditions. The CDF for each of these allocations was developed from the available data, and the results compared. The data were combined using custom MATLAB scripts to create data structures of results (as previously described), and then other MATLAB scripts were used to compare the results and prepare QQ plots of the distribution of the number of defender victories.

4.8 Overview of Evaluation Networks

This analysis technique was applied to three different networks as examples. The first, small-scale example is for a 13-node network, which is a capacitated network with sources and sinks on the boundaries of the network. The network valuation model used was a combination of the network flow model (maximizing the flow through the network) and the network interdiction model (determining the most critical node to remove to have the greatest reduction in network flow).

The second network is a mid-size (31-node) network that is loosely based upon the Irish electrical power grid. The network valuation model used was a pure maximum-flow model based on assumed values for the demand and generating power available at various nodes of the network. The network represented a simplification of the actual power network down to one that contained twice the number of nodes of the first small-scale example. Values taken for capacity and demand were purely for illustrative purposes.

The third network example is a large-scale (85-node) network that is a slight simplification of the Plantation Pipeline, which runs from several Gulf Coast ports along the Eastern Seaboard of the United States, delivering liquid product to a variety of locations running up to Dulles International Airport in Dulles, Virginia. This network was essentially a main branch with several smaller branches extending off of it. It was assessed that an interdiction model would provide essentially the same information as a straightforward network flow analysis, so only the later formulation was used to determine the network value of the nodes. Again, this network was analyzed for several different conditions.

4.9 Overview of Statistical Tests Used in This Analysis

Two different statistical tests were used in this analysis: the Lilliefors test for normality and the two-sample t-test for the difference in means (Sheskin, 2000). Both tests were executed on test data using functions from the MATLAB statistical toolbox. The results are presented in tables corresponding to the test results in each of the three networks analyzed.

The Lilliefors test is a test for normality of data. In this test, the null hypothesis is that the sample in the test data comes from a normal distribution, against the two-sided alternative that the data does not come from a normal distribution. Therefore, if the null hypothesis is rejected, the data are not normal.

For the two-sample t-test, the null hypothesis is that the two data samples are independent random samples from two normal populations with equal means and unknown but equal variances, against the two-sided alternative that the means are not equal. If the null hypothesis is rejected, the test result is that the means are different.

Chapter 5: Small-Scale Example

5.1 Evaluation Network

The small-scale network, which consisted of 13 nodes, was used as a first test of the technique. This network, represented in Figure 5, has two source nodes and five sink nodes.

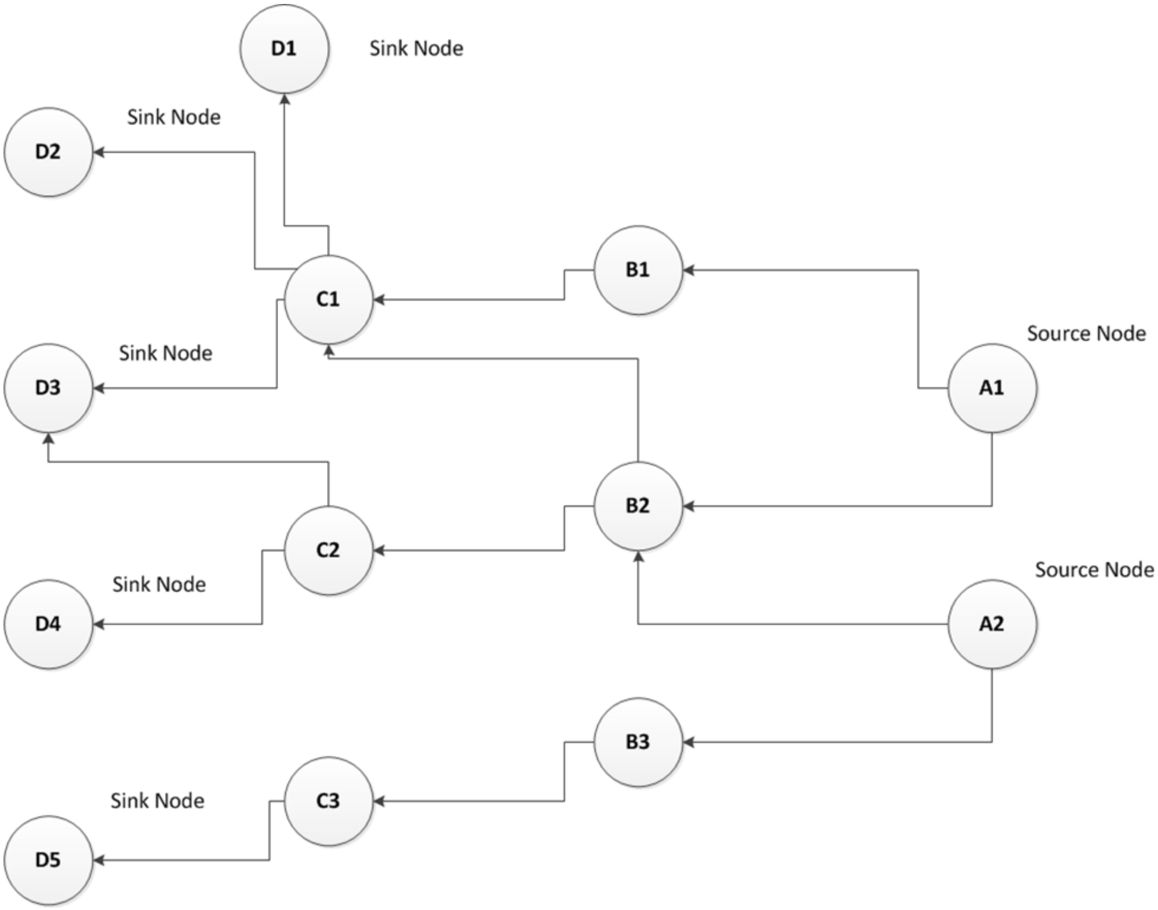

Figure 5 Small-Scale Network

In order to deal with the issue of node–arc equivalence, each node is split virtually into two nodes, one for all incoming arcs and one for all outgoing arcs, connected by a single virtual arc whose capacity defines the capacity of the node. The technique was followed in all the networks analyzed.

The network flow model and the network interdiction model were used in combination to determine the value of each node. First, the maximum-flow model was run and the flow through the nodes tabulated. Then the network interdiction model was run to identify the single node whose interdiction has the greatest effect upon the flow of the system. The flow through each node was determined under these conditions. The interdiction model was limited to selecting a single node by making the costs of interdiction and resources available for interdiction to drive the model to selecting a single node to interdict. This node was then temporarily removed from the set of possible nodes to interdict by making its cost of removal greater than the total removal budget available, driving the solution to the next, more important node. This process was repeated for several iterations, looking for other less-critical nodes. Interdicting some nodes of a network by setting their capacity to zero and then running a maximum-flow LP may result in some nodes having a greater flow through them than the original maximum-flow model. The final node value was determined by the maximum flow through the node considering all cases (both maximum-flow and interdicted). This method ensured that a node that was critical as a “backup” to another node was given the proper importance.

To examine the effect of other assumptions upon the results, six different scenarios were developed. The scenarios were characterized by the distribution of terrorist tactics and the PR values of different nodes. The distribution of terrorist tactics are either (1) uniform, which implied that the initial distribution of tactics for all terrorists is uniformly distributed among the possible strategies, or (2) favoring PR, in which case one half of the initial terrorists followed a tactic that attacks a node with a high PR value, and the rest of the of terrorists' tactics were uniformly distributed among the remaining tactics. The three sets of node PR values are (1) where one node has a higher PR value than the rest (Baseline Case), (2) the source nodes have a higher PR value (Source High), and (3) the sink nodes have a higher PR value (Sink High). The PR value was an arbitrary scale relating the relative importance of each node from a public relations standpoint. Note that at no time does this model require any sort of a transfer function to map PR values into network flow values or vice versa. They are completely separate valuations for the nodes. While these illustrations were done with only two values (PR and network flow), there is no reason that other "values" measuring other potential goals could not be included into this model. The key values for the nodes in this network are shown in Table 1.

Table 1 Small-Scale Network Node Data

Node ID	Node Value	Publicity Values		
		Baseline PR	Sink High PR	Source High PR
A1	15	1	1	7
A2	45	1	1	7
B1	15	1	1	1
B2	30	1	1	1
B3	15	1	1	1
C1	20	1	1	1
C2	20	1	1	1
C3	15	1	1	1
D1	20	1	7	1
D2	10	7	7	1
D3	20	1	7	1
D4	20	1	7	1
D5	15	1	7	1

The allocation model was run for the six scenarios (combinations of node values and terrorist strategies), and best resource allocation determined from those model runs. This allocation is called the Game Theoretic (GT) allocation. The comparison model was then run on the PRA allocation and on the GT allocation for all six scenarios and the results compared.

5.2 Determination of Resource Allocations

The resource allocation model was run for 2000 iterations for each of these six scenarios. The scenarios are characterized in Table 2.

Table 2 Small-Scale Network Scenarios

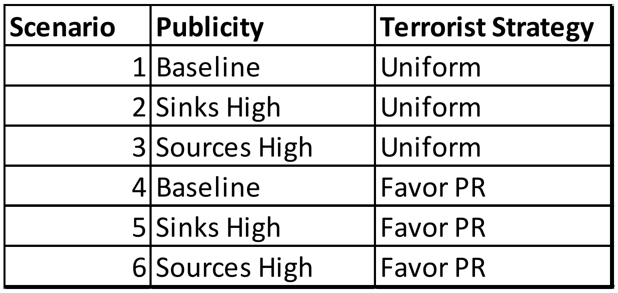

Scenario	Publicity	Terrorist Strategy
1	Baseline	Uniform
2	Sinks High	Uniform
3	Sources High	Uniform
4	Baseline	Favor PR
5	Sinks High	Favor PR
6	Sources High	Favor PR

Examination of the results showed that the distribution of agents (either terrorists or defenders) at each node tended to stabilize with little change after 1000 repetitions. Each scenario was run for 2000 repetitions. The number of defenders at each node was averaged for the last 500 repetitions, as was the number of defender victories. This process was then repeated for 400 simulation runs for each scenario to provide a distribution of results. The results for these cases can be seen in Figure 6 to Figure 11. The number of defender victories refers to the number of nodes that were successfully defended. Both terrorists and defenders had the same number of resources, which contributed to the seeming advantage held by the terrorists in application of resources. The average number of defenders at each node for the results that were most favorable to the defenders were then determined. This allocation became the Game Theoretic (GT) allocation of resources for the scenario. The Probabilistic Resource Allocation (PRA) was determined by allocating the resources according to the proportion of total value held by each of the nodes. This allocation was invariant across all scenarios. The resource allocations used for the defenders in the PRA case and each of the GT cases are summarized in Table 3. Two hundred terrorists and 200 defenders were used in these examples to minimize rounding error to determine a unitary number of defenders at each node.

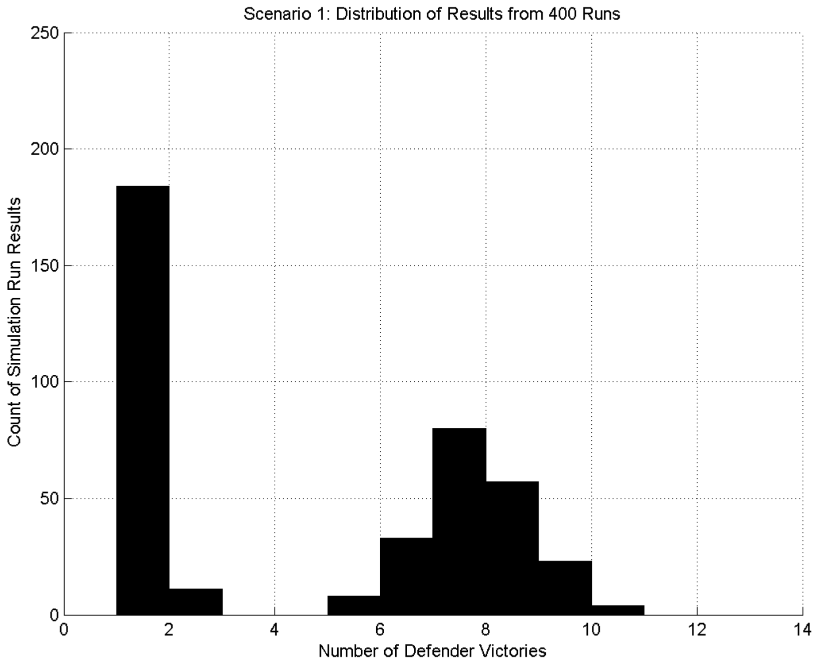

Figure 6 Small Network Scenario 1 Histogram

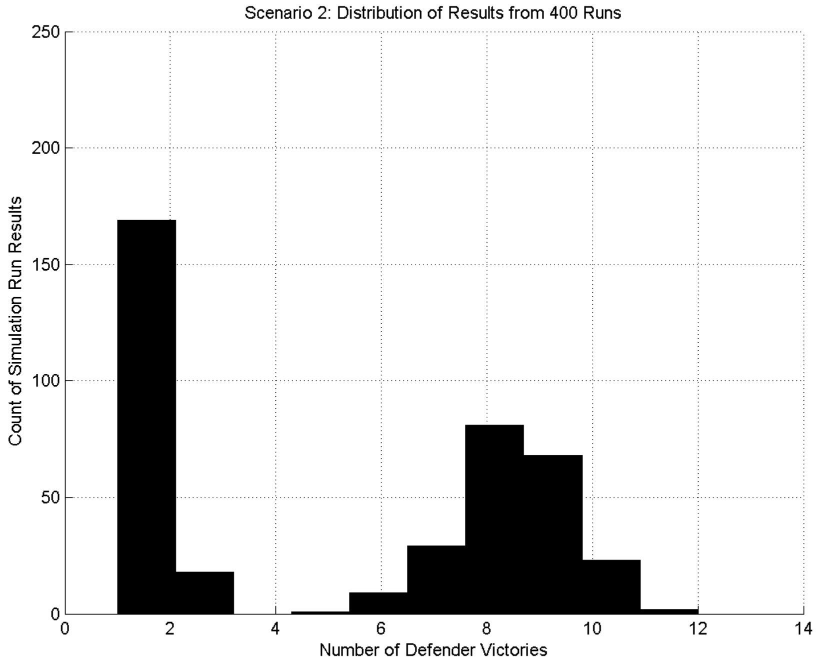

Figure 7 Small Network Scenario 2 Histogram

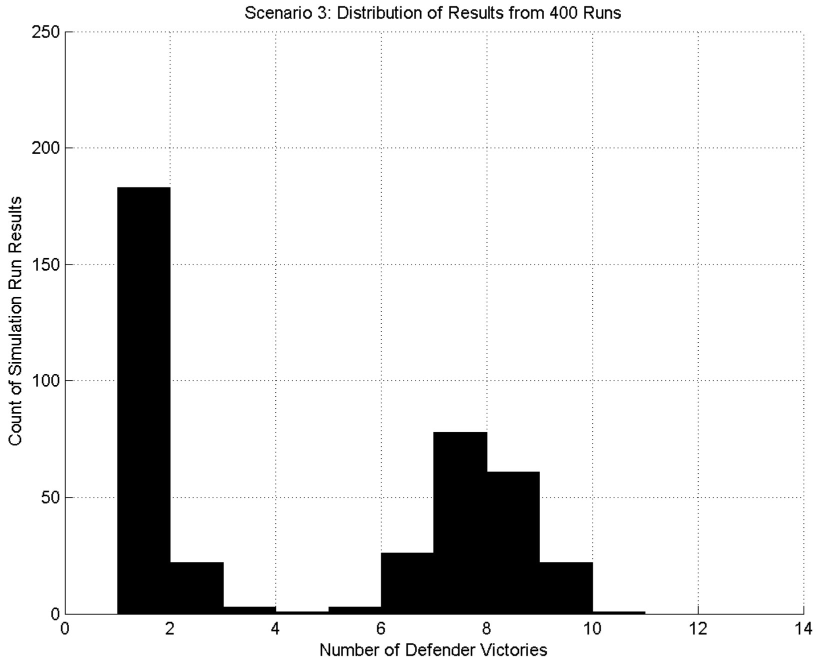

Figure 8 Small Network Scenario 3 Histogram

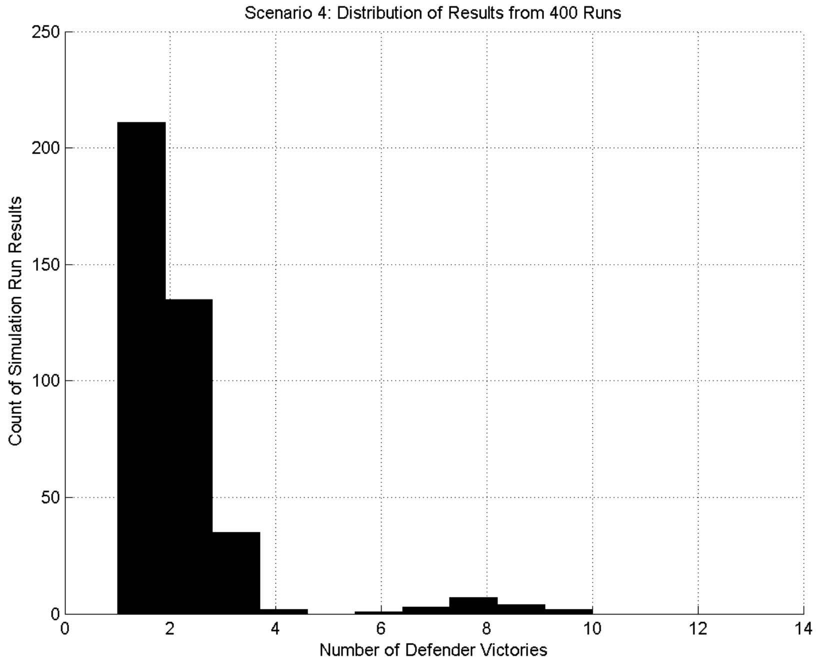

Figure 9 Small Network Scenario 4 Histogram

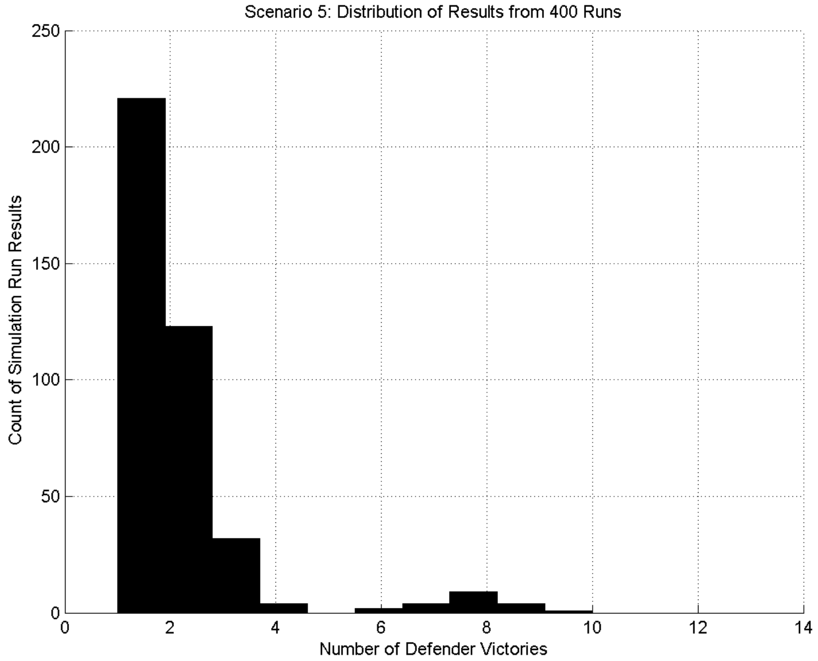

Figure 10 Small Network Scenario 5 Histogram

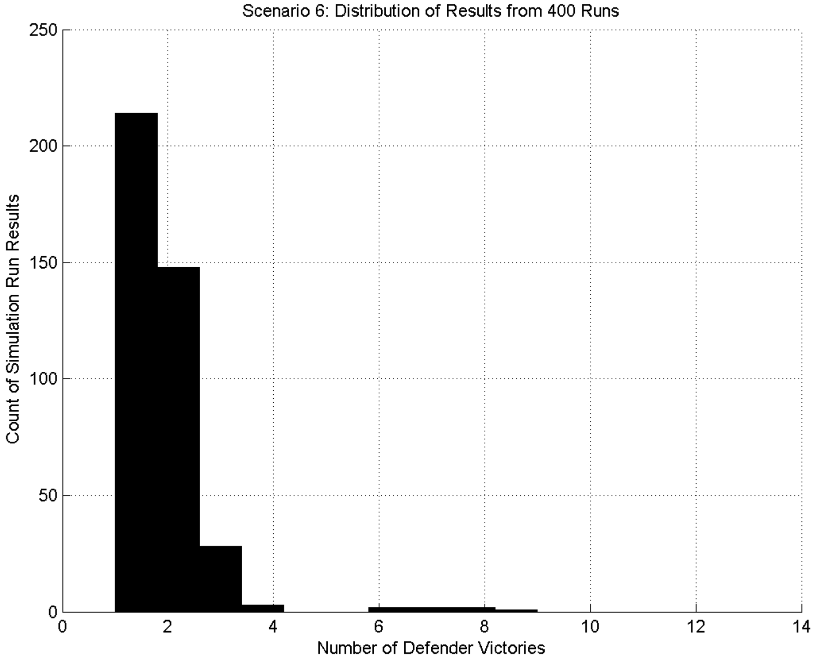

Figure 11 Small Network Scenario 6 Histogram

Table 3 Small Network Resource Allocations

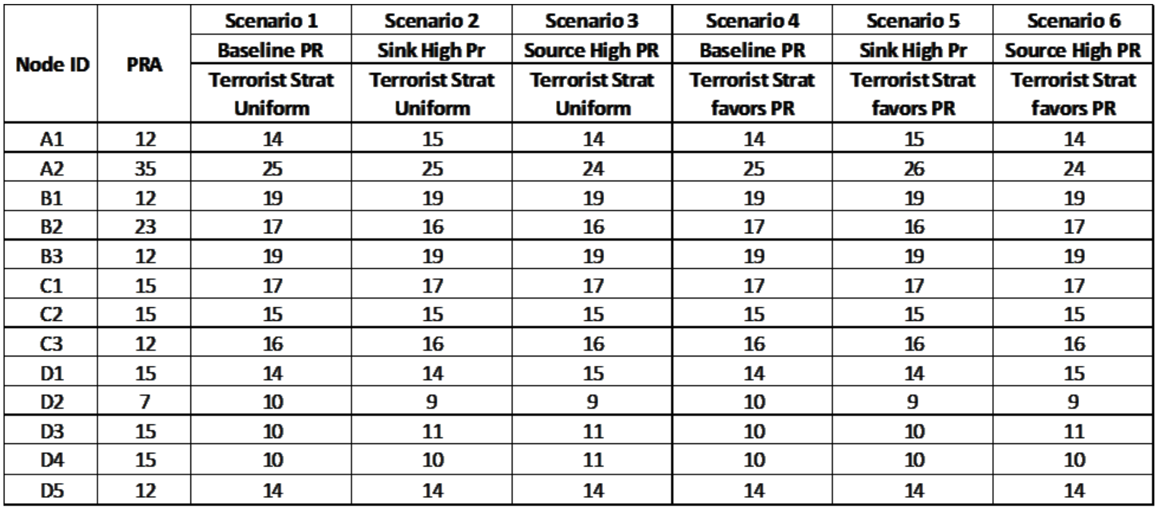

Node ID	PRA	Scenario 1	Scenario 2	Scenario 3	Scenario 4	Scenario 5	Scenario 6
		Baseline PR	Sink High Pr	Source High PR	Baseline PR	Sink High Pr	Source High PR
		Terrorist Strat Uniform	Terrorist Strat Uniform	Terrorist Strat Uniform	Terrorist Strat favors PR	Terrorist Strat favors PR	Terrorist Strat favors PR
A1	12	14	15	14	14	15	14
A2	35	25	25	24	25	26	24
B1	12	19	19	19	19	19	19
B2	23	17	16	16	17	16	17
B3	12	19	19	19	19	19	19
C1	15	17	17	17	17	17	17
C2	15	15	15	15	15	15	15
C3	12	16	16	16	16	16	16
D1	15	14	14	15	14	14	15
D2	7	10	9	9	10	9	9
D3	15	10	11	11	10	10	11
D4	15	10	10	11	10	10	10
D5	12	14	14	14	14	14	14

5.3 Comparison of Allocation Results

Using the same run parameters (2000 repetitions per run, results taken as the average of the last 500 repetitions, 400 runs for each allocation and scenario) the comparison model was then run for each of these scenarios. The comparison model examined the performance of the PRA allocation and the GT allocation against an evolving terrorist threat in all six of the scenarios. The results were examined for normality using QQ plots and tested using the Lilliefors test, as previously discussed. The QQ plots are shown in Figure 12 to Figure 17. The results of the Lillefors tests for normality are shown in Table 4. In this test, the null hypothesis was that the distribution is normal against a two-sided alternative that the distribution is not normal.

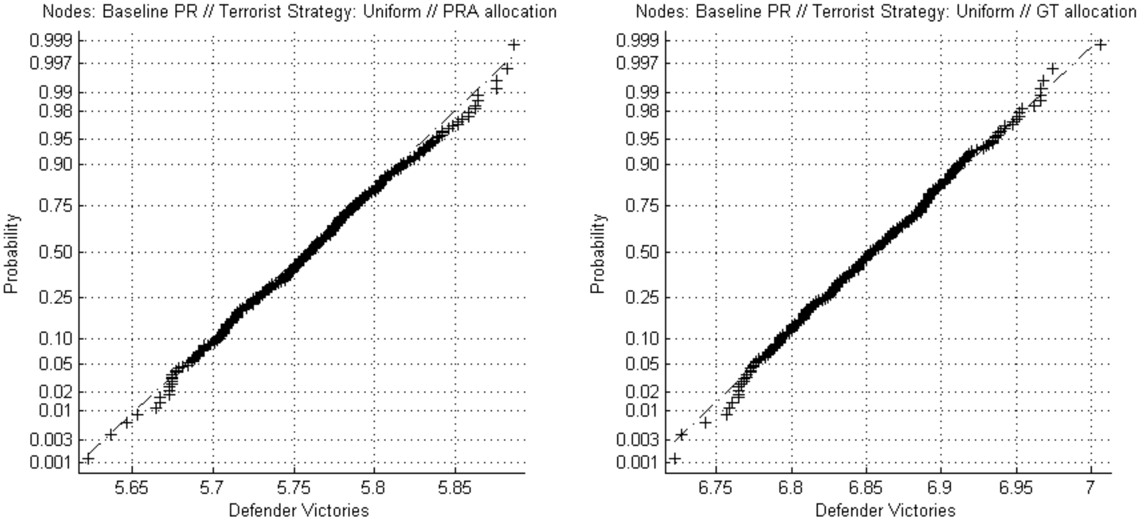

Figure 12 Small Network Scenario 1 QQ Plot

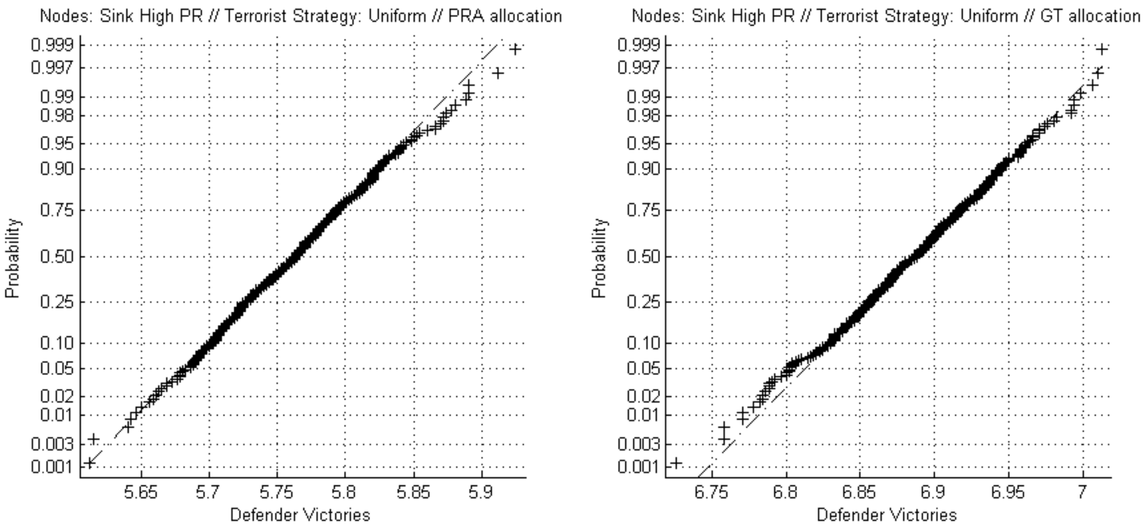
Nodes: Sink High PR // Terrorist Strategy: Uniform // PRA allocation

Nodes: Sink High PR // Terrorist Strategy: Uniform // GT allocation

Figure 13 Small Network Scenario 2 QQ Plot

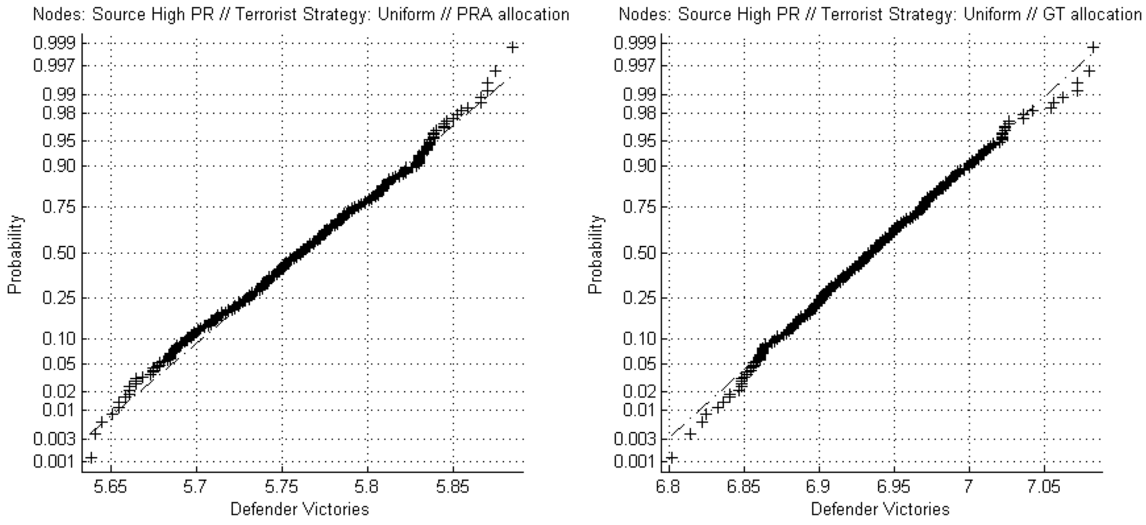
Nodes: Source High PR // Terrorist Strategy: Uniform // PRA allocation

Nodes: Source High PR // Terrorist Strategy: Uniform // GT allocation

Figure 14 Small Network Scenario 3 QQ Plot

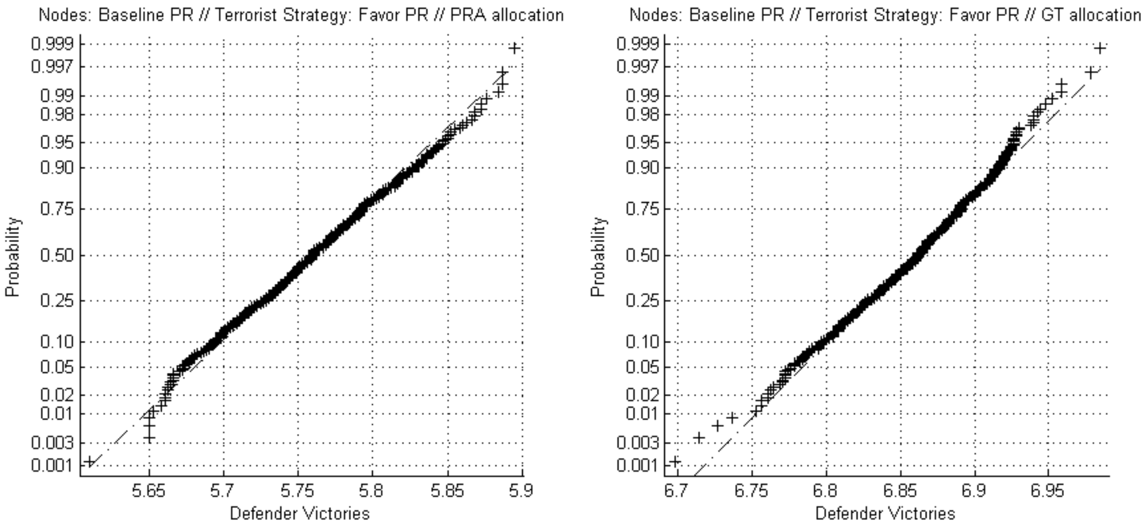

Figure 15 Small Network Scenario 4 QQ Plot

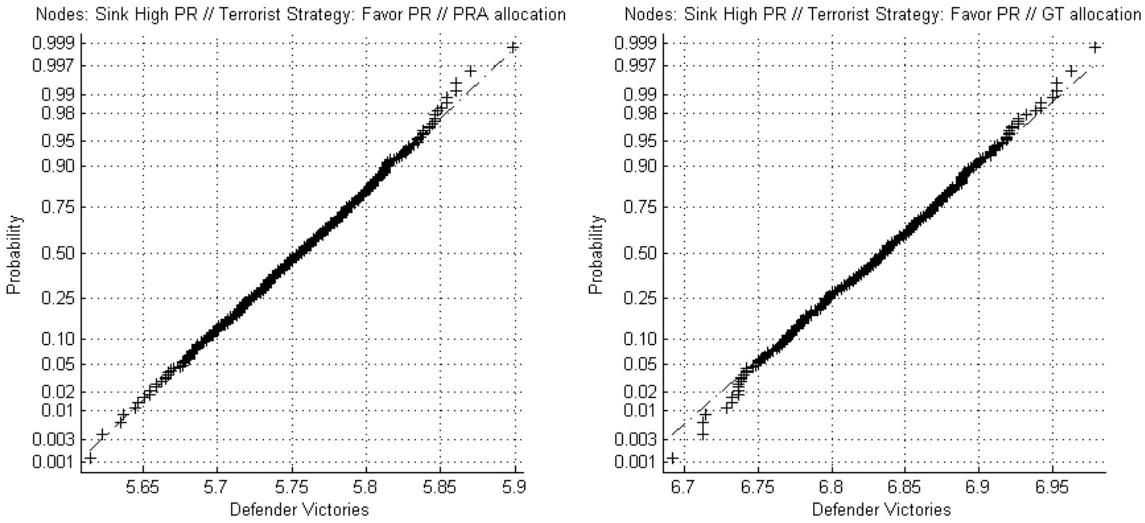

Figure 16 Small Network Scenario 5 QQ Plot

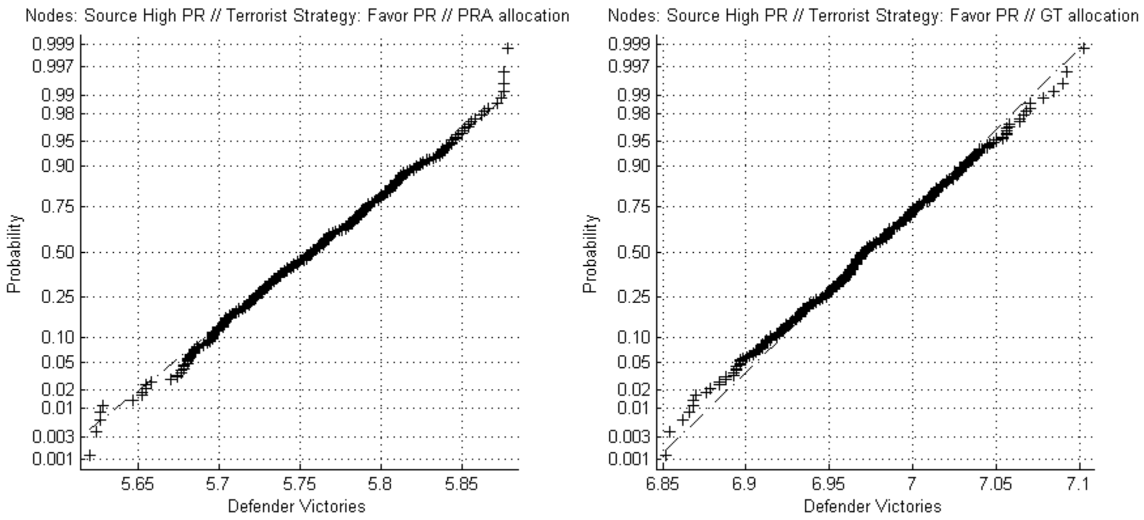

Figure 17 Small Network Scenario 6 QQ Plot

Table 4 Small Network Lilliefors Test Results

Lilliefors Test Results			
Scenario	**Allocation**	**p value**	**Result**
1	PRA	0.4122	Do Not Reject H_0
2	PRA	> .5	Do Not Reject H_0
3	PRA	> .5	Do Not Reject H_0
4	PRA	> .5	Do Not Reject H_0
5	PRA	> .5	Do Not Reject H_0
6	PRA	> .5	Do Not Reject H_0
1	GT	> .5	Do Not Reject H_0
2	GT	> .5	Do Not Reject H_0
3	GT	> .5	Do Not Reject H_0
4	GT	0.0161	Reject H_0
5	GT	0.4607	Do Not Reject H_0
6	GT	0.1122	Do Not Reject H_0

For each of the scenarios, the CDF of the number of defender victories observed with the PRA-based allocation and the CDF of the number of defender victories with the GT-based allocation was compared. These results are shown in Figure 18 to Figure 23.

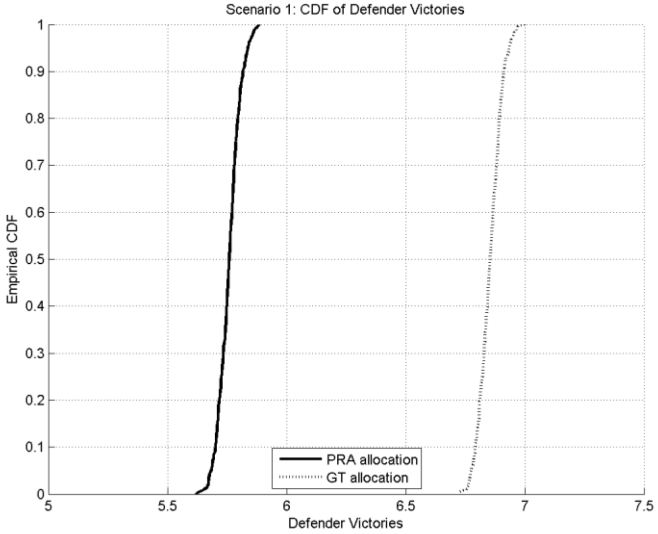

Figure 18 Small Network Scenario 1 CDF

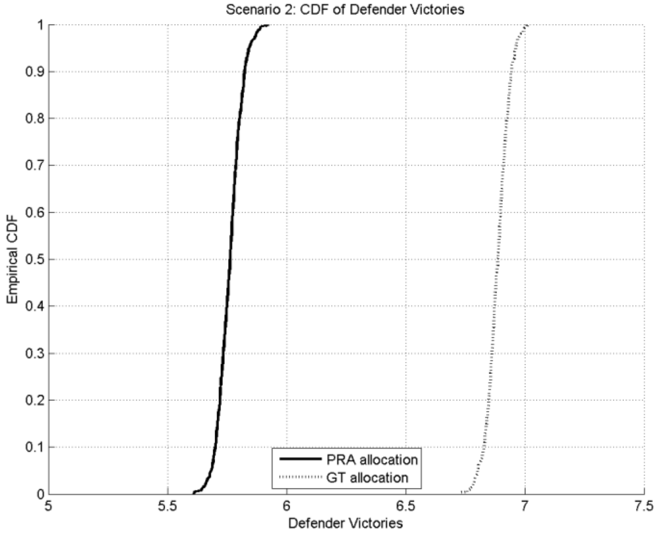

Figure 19 Small Network Scenario 2 CDF

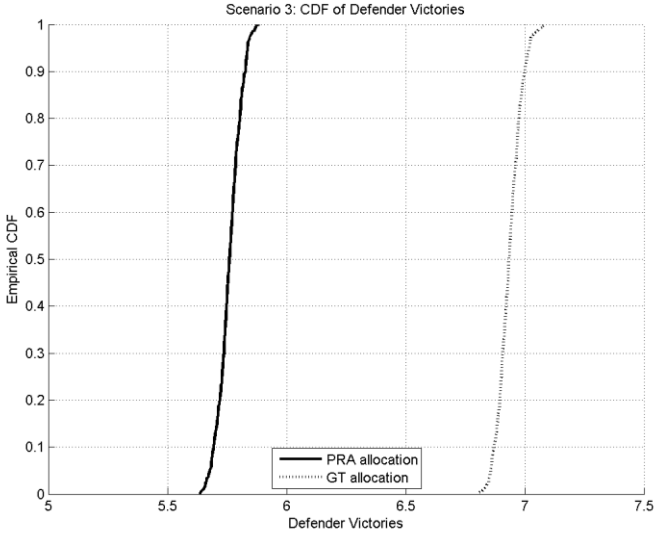

Figure 20 Small Network Scenario 3 CDF

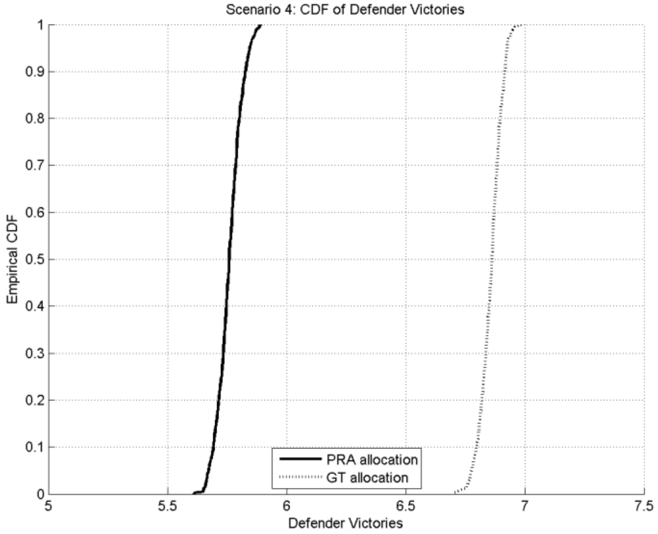

Figure 21 Small Network Scenario 4 CDF

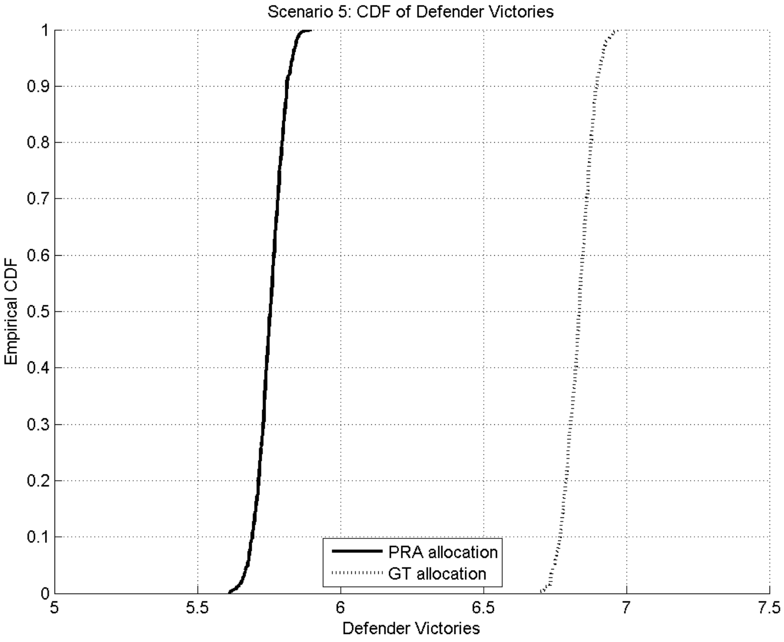

Figure 22 Small Network Scenario 5 CDF

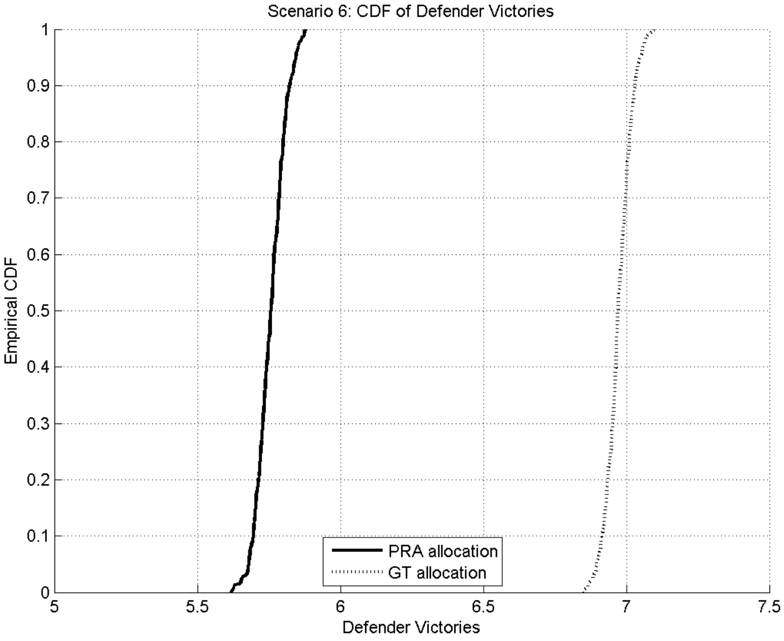

Figure 23 Small Network Scenario 6 CDF

Examination of the CDF plots in Figure 18 to Figure 23 shows that, in every case, the PRA-based allocation is dominated by the GT-based allocation in the sense that, for any specified value of the CDF, the GT-based approach always produces a higher number of defender victories than does the PRA-based approach. As shown in Table 5, the differences between the means of the GT- and PRA-based solutions are statistically significant.

Table 5 Small-Scale Network Means Test Results

Mean Test Result		
Scenario	**T statistic**	**Result**
1	-332.1778	Reject H_0
2	-326.0212	Reject H_0
3	-341.3158	Reject H_0
4	-320.6159	Reject H_0
5	-310.9925	Reject H_0
6	-359.9671	Reject H_0

Chapter 6: Mid-Scale Example

6.1 Evaluation Network

The mid-scale network example, which consists of 31 nodes, was used to see how the technique would scale to larger problems. The network was derived from publically available data on the form of the Irish power grid ("Transmission System," 2007). These data were then simplified to a smaller network of the desired size by consolidating nodes. Values for generating capacity, maximum flow, and demand were constructive and chosen for illustrative purposes. Actual data were not available, but the actual capacities were not critical to the research. The key was to determine if the behavior observed in the small-scale example discussed in the preceding chapter would continue in the larger example. The network representation used is shown in Figure 24. Details as to the function of the various nodes are shown in Table 6.

To examine the effect that the method of determining the value of each node might have on the overall problem, only the maximum-flow algorithm was used to determine the value of each node. Four different scenarios were created to facilitate the sensitivity analysis of different parameters. In each of these scenarios, there were 93 defenders who were fixed by assigning them to specific nodes. An additional 310 defenders were allowed to move from node to node depending upon their strategies.

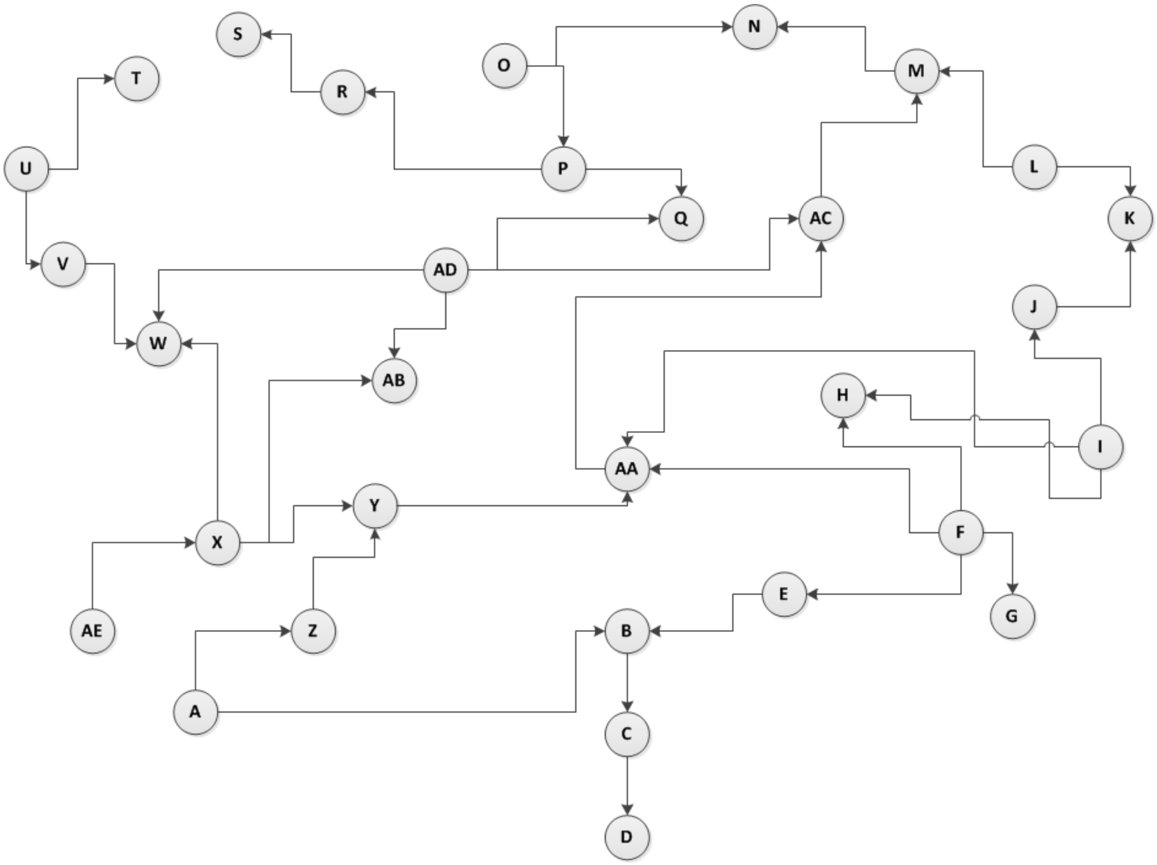

Figure 24 Mid-Scale Network

Table 6 Mid-Scale Network Node Details

Node	Function	Capacity	Demand	Inflow	Node	Function	Capacity	Demand	Inflow
A	SOURCE	500		300	Q	SINK	500	150	
B	NODE	500	50		R	NODE	500	50	
C	NODE	500	50		S	SINK	500	100	
D	SINK	500	100		T	SINK	500	100	
E	NODE	500	50		U	SOURCE	500		300
F	SOURCE	500		300	V	NODE	500	50	
G	SINK	500	100		W	SINK	500	100	
H	SINK	500	100		X	NODE	500	50	
I	SOURCE	500		300	Y	NODE	500	50	
J	NODE	500	50		Z	NODE	500	50	
K	SINK	500	200		AA	NODE	500	50	
L	SOURCE	500		300	AB	SINK	500	100	
M	NODE	500	50		AC	NODE	500	50	
N	SINK	500	100		AD	SOURCE	500		500
O	SOURCE	500		300	AE	SOURCE	500		300
P	NODE	500	50						

There were 403 terrorists in the analysis, all of whom were free to evolve their strategies. Scenarios were created by varying assumptions. The two assumptions that varied were the allocation of fixed defenders and the distribution of the terrorists. In some scenarios, the fixed, or inherent, defenders were allocated in a non-uniform manner; i.e., some nodes had more inherent defenders than others. In other scenarios, these were allocated uniformly with three inherent defenders to each node. The other assumption that varied was that in some of the scenarios the terrorists had a uniform probability of being any one of the six tactics, and in the other case the probability that they would favor PR was 0.5, with the remaining probability uniformly distributed among the other tactics, in a similar manner previously described for the small-scale network example. The PR values were arbitrarily assigned to nodes and did not vary from one scenario to another. Table 7 lists a summary of the varying values for the scenarios, and Table 8 provides a summary of the actual values used.

Table 7 Mid-Scale Scenarios

Scenario	**Inherent defense allocation**	**Terrorist Strategies**
1	targeted	Constant
2	targeted	Favor PR
3	uniform	Constant
4	uniform	Favor PR

Table 8 Mid-Scale Network Parameters

			Inherent Defense	
Node	Value	PRValue	Targeted	Uniform
A	300	10	5	3
B	200	10	1	3
C	150	10	1	3
D	100	10	1	3
E	50	10	1	3
F	300	10	9	3
G	100	10	1	3
H	100	10	1	3
I	300	10	9	3
J	50	10	1	3
K	340	20	1	3
L	300	10	8	3
M	650	10	1	3
N	690	30	1	3
O	300	10	9	3
P	200	10	1	3
Q	150	10	1	3
R	150	10	4	3
S	100	10	1	3
T	100	10	1	3
U	300	10	9	3
V	190	10	1	3
W	140	10	1	3
X	290	30	1	3
Y	280	10	1	3
Z	90	30	1	3
AA	510	10	1	3
AB	100	40	1	3
AC	700	10	1	3
AD	300	10	9	3
AE	300	10	9	3

The allocation model was run for the four scenarios (combinations of inherent defense allocations and terrorist strategies) and the best allocation determined from those model runs. The allocation is called the Game Theoretic (GT) allocation. The comparison model was then run on the PRA allocation and on the GT allocation for all four scenarios and the results compared.

6.2 Determination of Resource Allocations

The resource allocation model was run for 2000 repetitions for each of these four scenarios. Examination of the results showed that the method of analysis used on the small network example was applicable. The allocation of resources stabilized after 1000 repetitions, and the average number of defenders at each node over the last 500 repetitions was used as the basis for a distribution of results. The results for these cases can be seen in Figure 25 to Figure 28. The number of defender victories refers to the number of nodes that were successfully defended. The GT and PRA allocations were calculated as discussed in the small network example. Table 9 lists the final allocations that were calculated.

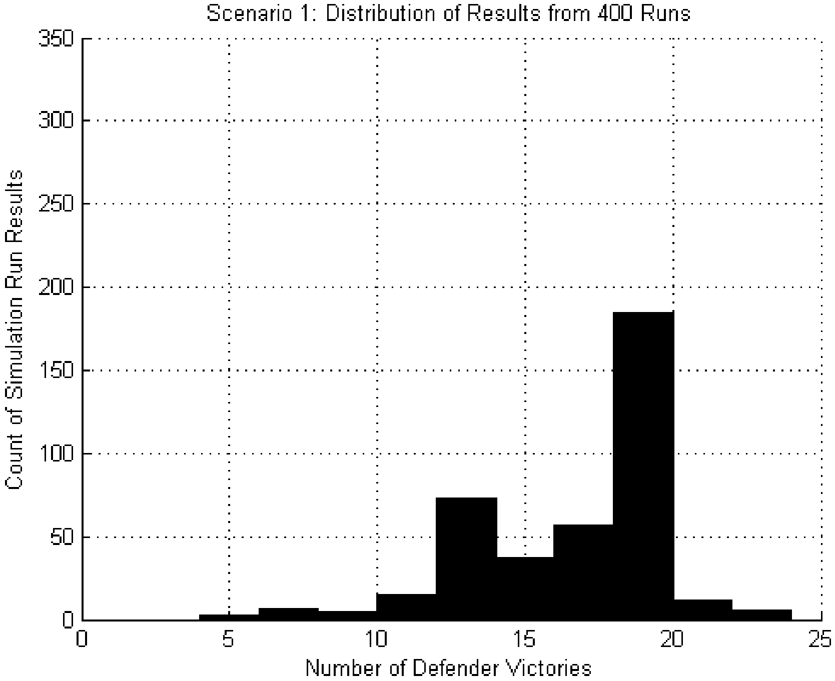

Figure 25 Mid-Scale Network Scenario 1 Histogram

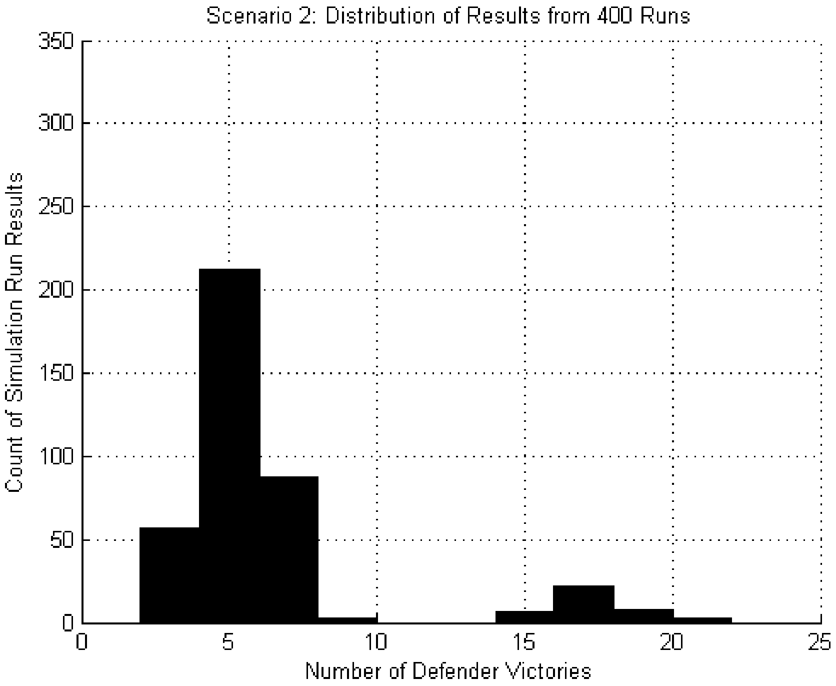

Figure 26 Mid-Scale Network Scenario 2 Histogram

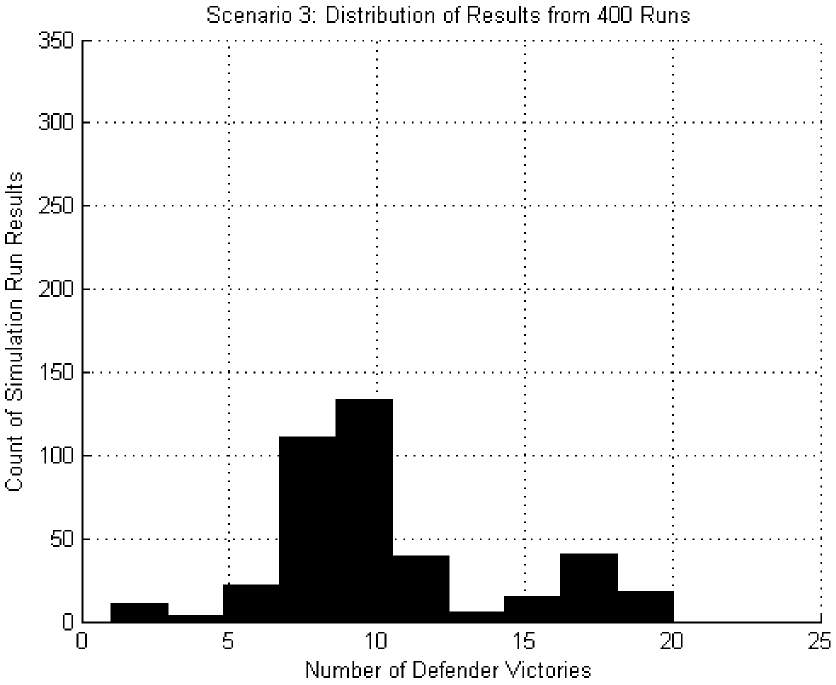

Figure 27 Mid-Scale Network Scenario 3 Histogram

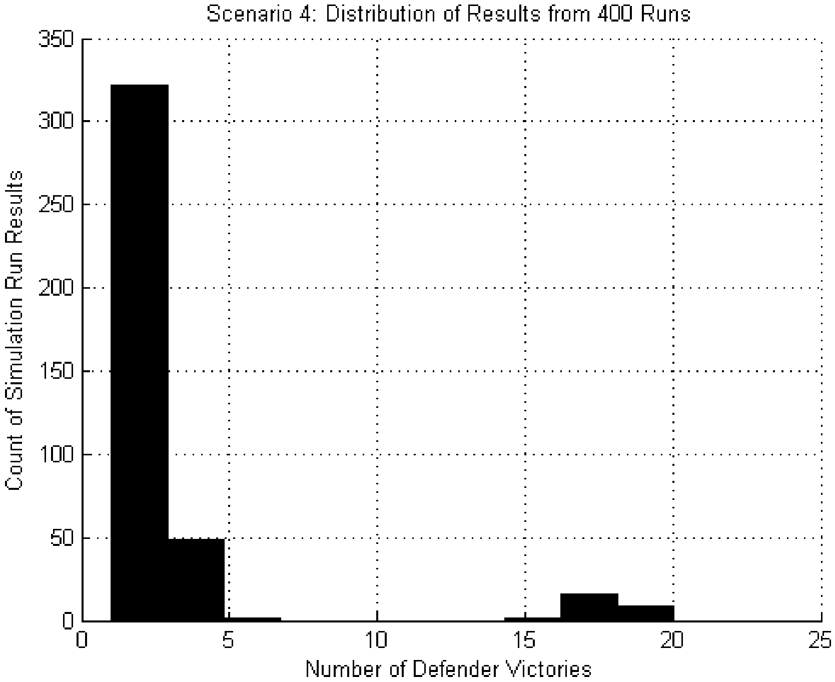

Figure 28 Mid-Scale Network Scenario 4 Histogram

Table 9 Mid-Scale Network Resource Allocations

			Scenario 1	Scenario 2	Scenario 3	Scenario 4
	PRA		Targeted Defense	Targeted Defense	Uniform Defense	Uniform Defense
Node ID	Targeted	Uniform	Terrorist Strategy Uniform	Terrorist Strategy favors PR	Terrorist Strategy Uniform	Terrorist Strategy Favors PR
A	17	15	**16**	**16**	**14**	**14**
B	9	11	9	9	11	11
C	7	9	9	9	11	11
D	5	7	9	9	11	11
E	3	5	9	9	11	11
F	21	15	19	19	12	13
G	5	7	11	11	14	14
H	5	7	9	9	11	11
I	21	15	18	18	12	12
J	3	5	9	9	11	11
K	14	16	14	14	16	16
L	20	15	16	16	11	11
M	27	29	11	11	13	13
N	28	30	16	17	18	18
O	21	15	18	18	12	12
P	9	11	9	8	11	10
Q	7	9	9	8	10	11
R	10	9	12	12	11	11
S	5	7	9	9	11	11
T	5	7	9	9	11	11
U	21	15	18	18	12	12
V	8	10	9	9	11	11
W	7	9	9	9	11	11
X	12	14	10	10	12	12
Y	12	14	9	8	11	10
Z	4	6	10	10	13	13
AA	21	23	14	15	16	16
AB	5	7	13	13	15	15
AC	29	31	33	34	35	35
AD	21	15	17	17	11	11
AE	21	15	20	20	14	14

6.3 Comparison of Allocations

Using the same run parameters (2000 repetitions per run, results taken as the average of the last 500 repetitions, 400 runs for each allocation and scenario) the comparison model was then run for each of these scenarios. The comparison model examined the performance of the PRA allocation and the GT allocation against an evolving terrorist threat in all four of the scenarios. The results were examined for normality using QQ plots and tested using the Lilliefors test. The QQ plots are shown in Figure 29 to Figure 32. The Lilliefors tests for normality are shown in Table 10.

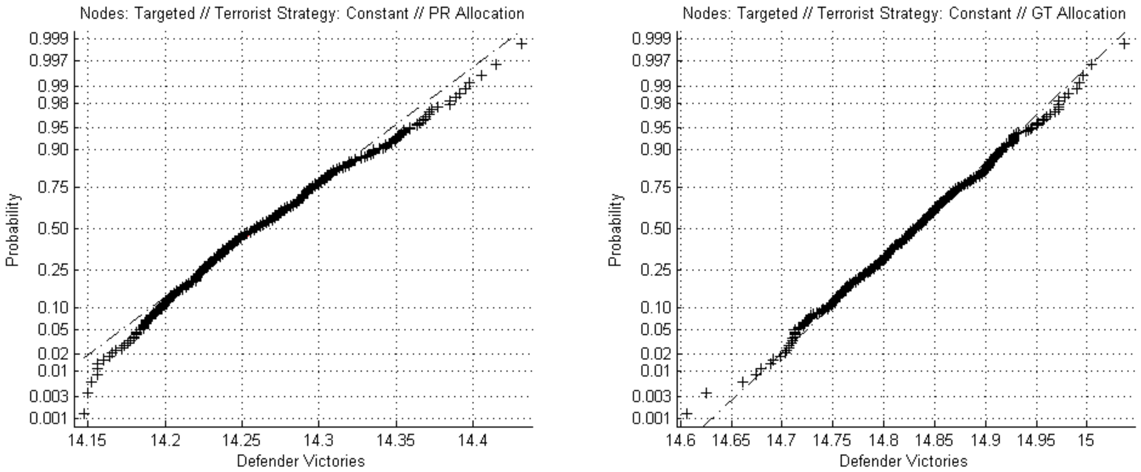

Figure 29 Mid-Scale Network Scenario 1 QQ Plot

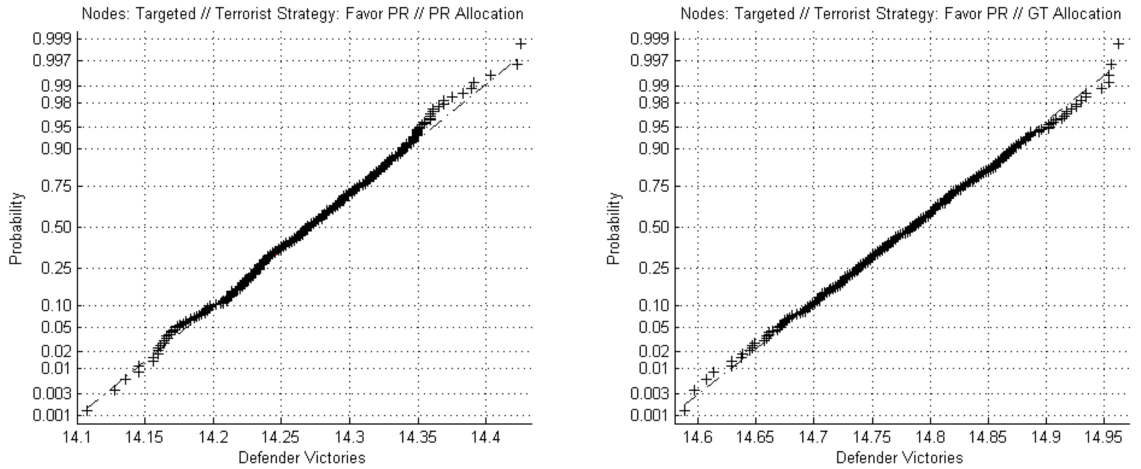

Figure 30 Mid-Scale Network Scenario 2 QQ Plot

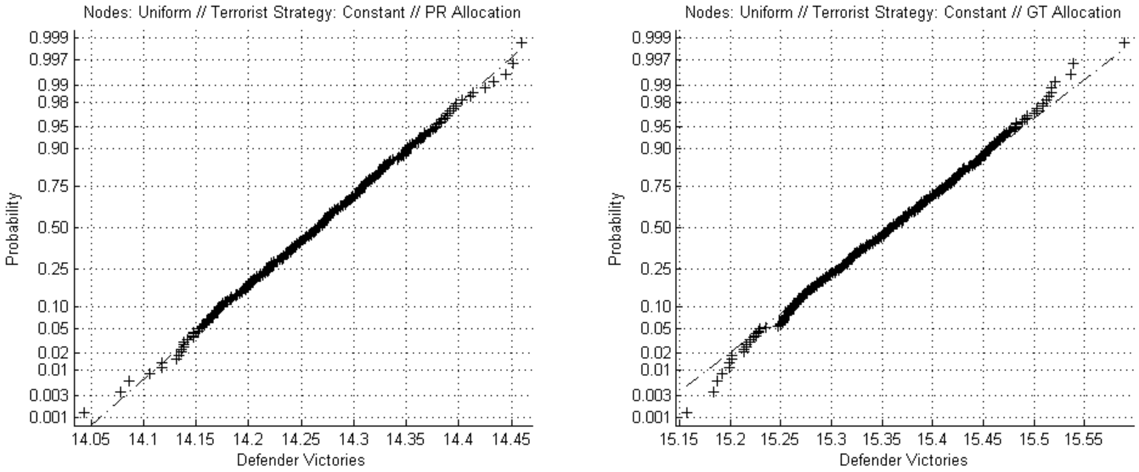

Figure 31 Mid-Scale Network Scenario 3 QQ Plot

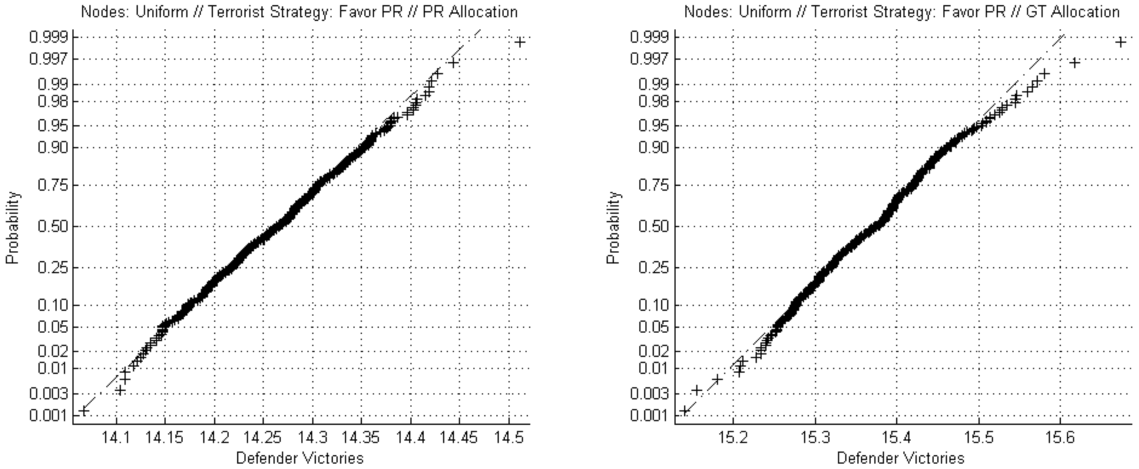

Figure 32 Mid-Scale Network Scenario 4 QQ Plot

Table 10 Mid-Scale Network Lilliefors Test Results

Lilliefors Test Results			
Scenario	Allocation	p value	Result
1	PRA	> .5	Do Not Reject H_0
2	PRA	> .5	Do Not Reject H_0
3	PRA	> .5	Do Not Reject H_0
4	PRA	> .5	Do Not Reject H_0
1	GT	> .5	Do Not Reject H_0
2	GT	> .5	Do Not Reject H_0
3	GT	> .5	Do Not Reject H_0
4	GT	0.2404	Do Not Reject H_0

For each of the scenarios, the CDF of the number of defender victories observed with the PRA-based allocation and the CDF of the number of defender victories observed with the GT-based allocation were compared. These results are shown in Figure 33 to Figure 36.

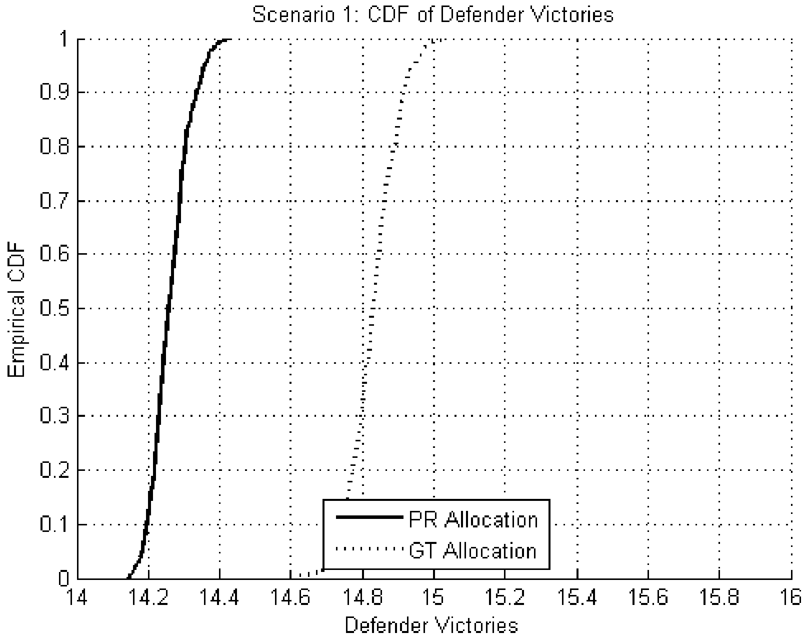

Figure 33 Mid-Scale Network Scenario 1 CDF

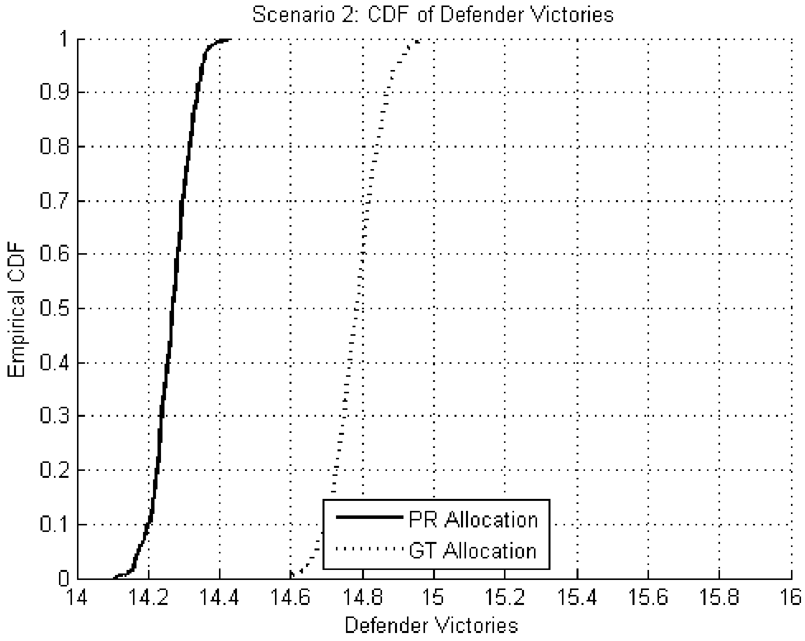

Figure 34 Mid-Scale Network Scenario 2 CDF

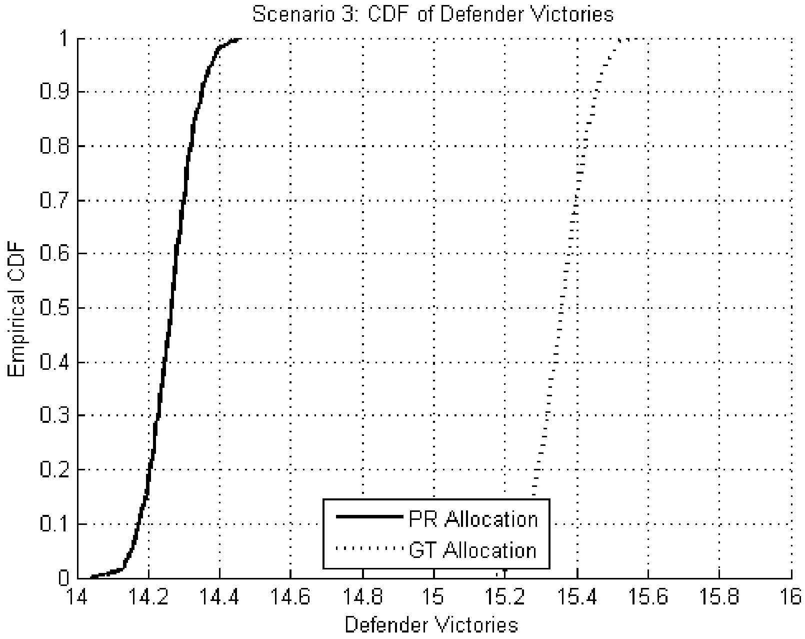

Figure 35 Mid-Scale Network Scenario 3 CDF

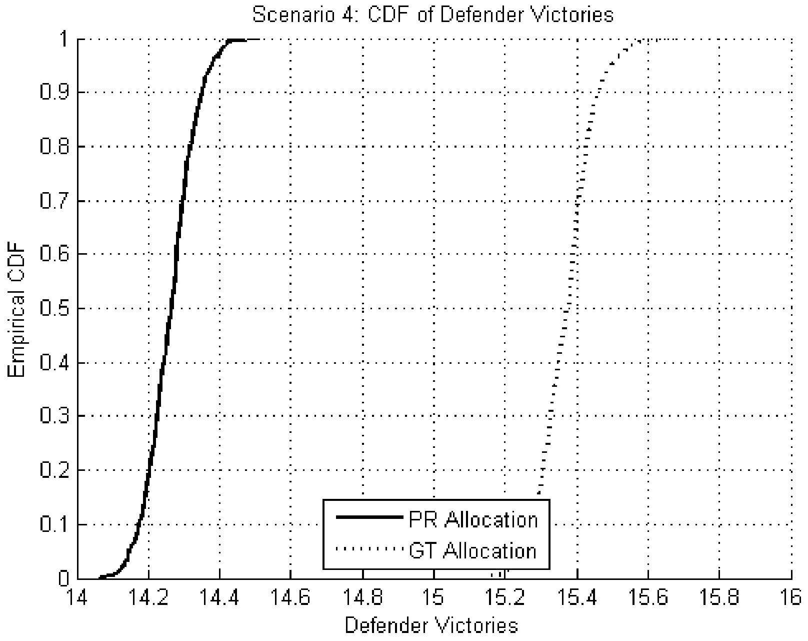

Figure 36 Mid-Scale Network Scenario 4 CDF

Examination of the CDF plots in Figure 33 to Figure 36 shows that in every case the PRA-based allocation is dominated by the GT-based allocation in the sense that, for any specified value of the CDF, the GT-based approach always produces a higher number of defender victories than does the PRA-based approach. These results are consistent with the results for the small-scale network presented in the previous chapter. The differences in means are statistically significant. The means test results are shown in Table 11.

Table 11 Mid-Scale Network Means Test Results

Mean Test Result		
Scenario	**T statistic**	**Result**
1	-131.3408	Reject H_0
2	-118.095	Reject H_0
3	-216.9933	Reject H_0
4	-217.7804	Reject H_0

Chapter 7: Large-Scale Example

7.1 Evaluation Network

The large-scale network, which consists of 85 nodes, was used to continue to examine the ability of the technique to work on problems of increasing scale. This network was also derived from publically available data on Kinder-Morgan Corporation's Plantation Pipeline ("Plantation Pipeline System Map," 2012). The capacities and demands were chosen for illustrative purposes only. The key elements of this example that make it different from the previous two examples are both the scale of the network and the very linear arrangement of the network. The pipeline consists of a main branch running up and down the East Coast of the United States, with a number of smaller branches running off and leading to locations where the product is demanded. As in the Irish power grid example, actual data were not available, but the actual capacities were not critical to the research. The goal was to see if the results continued to scale up with larger networks and with a different network topology. The network representation is shown in Figure 37. Note that the physical arrangement of this figure has been modified to allow for easy publication. In actuality, the lower section (from Anderson, SC to Washington, DC) would be a continuation of the "main trunk," which runs from Baton Rouge, LA to Hartwell, GA. This Northeastern portion of the network was "moved" for illustration to allow it to fit comfortably on a single page. Details as to the function of the various nodes are shown in Table 12.

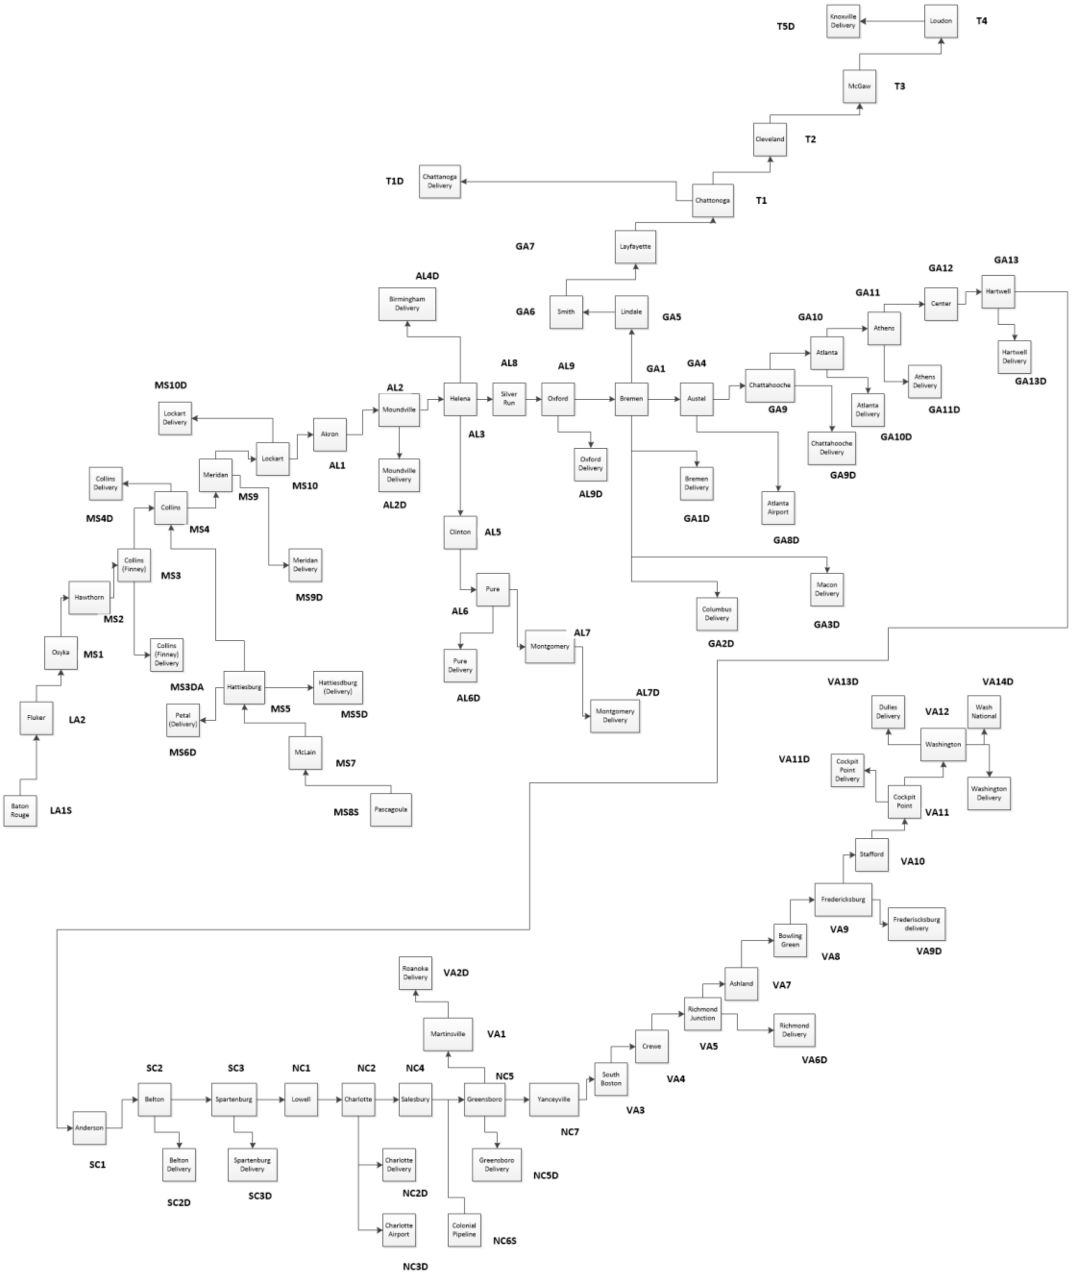

Figure 37 Large-Scale Network

Table 12 Large-Scale Network Node Details

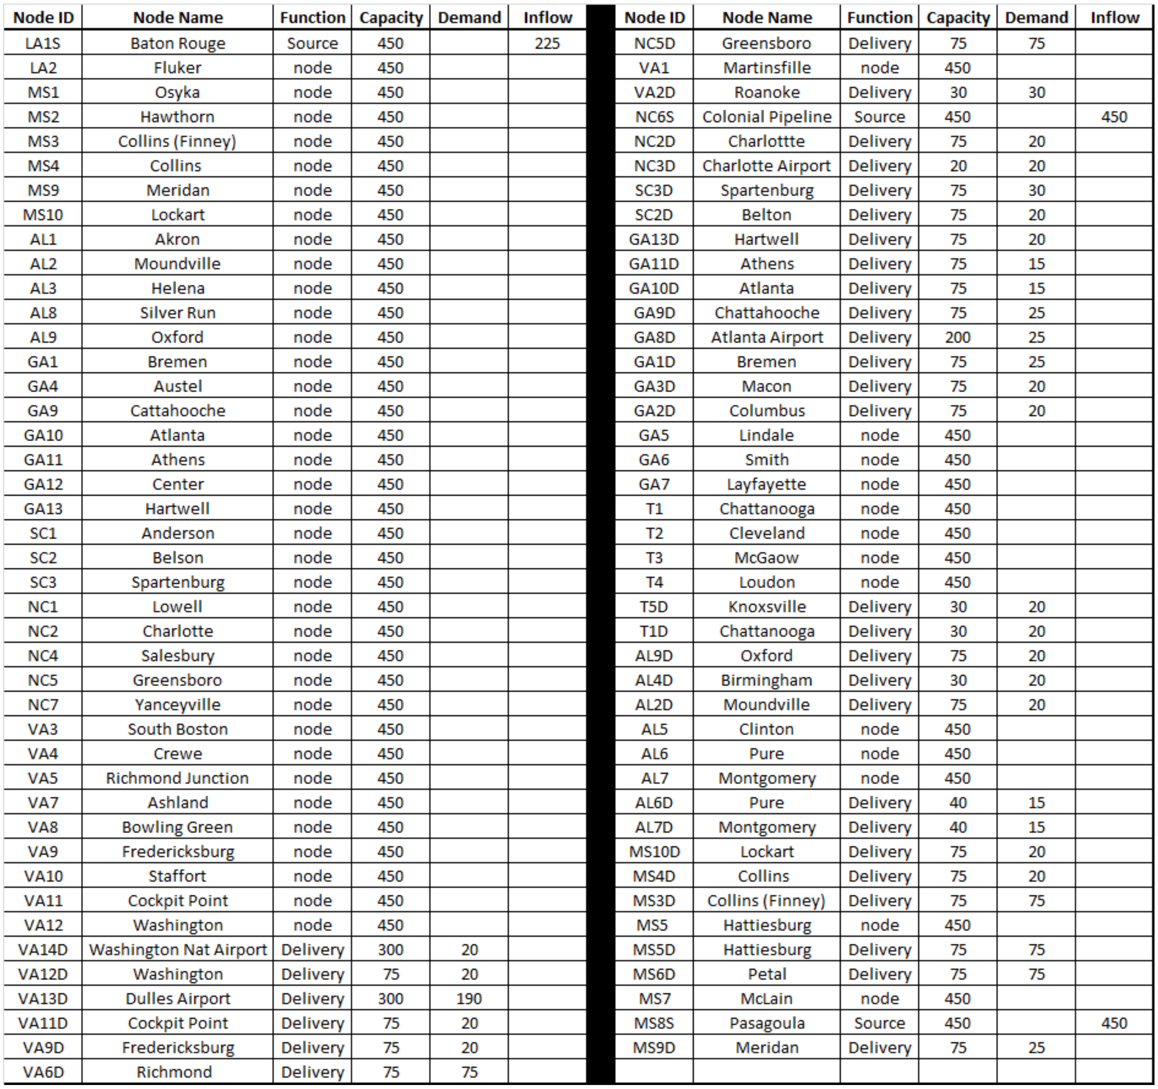

Node ID	Node Name	Function	Capacity	Demand	Inflow
LA1S	Baton Rouge	Source	450		225
LA2	Fluker	node	450		
MS1	Osyka	node	450		
MS2	Hawthorn	node	450		
MS3	Collins (Finney)	node	450		
MS4	Collins	node	450		
MS9	Meridan	node	450		
MS10	Lockart	node	450		
AL1	Akron	node	450		
AL2	Moundville	node	450		
AL3	Helena	node	450		
AL8	Silver Run	node	450		
AL9	Oxford	node	450		
GA1	Bremen	node	450		
GA4	Austel	node	450		
GA9	Cattahooche	node	450		
GA10	Atlanta	node	450		
GA11	Athens	node	450		
GA12	Center	node	450		
GA13	Hartwell	node	450		
SC1	Anderson	node	450		
SC2	Belson	node	450		
SC3	Spartenburg	node	450		
NC1	Lowell	node	450		
NC2	Charlotte	node	450		
NC4	Salesbury	node	450		
NC5	Greensboro	node	450		
NC7	Yanceyville	node	450		
VA3	South Boston	node	450		
VA4	Crewe	node	450		
VA5	Richmond Junction	node	450		
VA7	Ashland	node	450		
VA8	Bowling Green	node	450		
VA9	Fredericksburg	node	450		
VA10	Stafford	node	450		
VA11	Cockpit Point	node	450		
VA12	Washington	node	450		
VA14D	Washington Nat Airport	Delivery	300	20	
VA12D	Washington	Delivery	75	20	
VA13D	Dulles Airport	Delivery	300	190	
VA11D	Cockpit Point	Delivery	75	20	
VA9D	Fredericksburg	Delivery	75	20	
VA6D	Richmond	Delivery	75	75	
NC5D	Greensboro	Delivery	75	75	
VA1	Martinsfille	node	450		
VA2D	Roanoke	Delivery	30	30	
NC6S	Colonial Pipeline	Source	450		450
NC2D	Charlottte	Delivery	75	20	
NC3D	Charlotte Airport	Delivery	20	20	
SC3D	Spartenburg	Delivery	75	30	
SC2D	Belton	Delivery	75	20	
GA13D	Hartwell	Delivery	75	20	
GA11D	Athens	Delivery	75	15	
GA10D	Atlanta	Delivery	75	15	
GA9D	Chattahooche	Delivery	75	25	
GA8D	Atlanta Airport	Delivery	200	25	
GA1D	Bremen	Delivery	75	25	
GA3D	Macon	Delivery	75	20	
GA2D	Columbus	Delivery	75	20	
GA5	Lindale	node	450		
GA6	Smith	node	450		
GA7	Layfayette	node	450		
T1	Chattanooga	node	450		
T2	Cleveland	node	450		
T3	McGaow	node	450		
T4	Loudon	node	450		
T5D	Knoxsville	Delivery	30	20	
T1D	Chattanooga	Delivery	30	20	
AL9D	Oxford	Delivery	75	20	
AL4D	Birmingham	Delivery	30	20	
AL2D	Moundville	Delivery	75	20	
AL5	Clinton	node	450		
AL6	Pure	node	450		
AL7	Montgomery	node	450		
AL6D	Pure	Delivery	40	15	
AL7D	Montgomery	Delivery	40	15	
MS10D	Lockart	Delivery	75	20	
MS4D	Collins	Delivery	75	20	
MS3D	Collins (Finney)	Delivery	75	75	
MS5	Hattiesburg	node	450		
MS5D	Hattiesburg	Delivery	75	75	
MS6D	Petal	Delivery	75	75	
MS7	McLain	node	450		
MS8S	Pasagoula	Source	450		450
MS9D	Meridan	Delivery	75	25	

Because of the topology of this network, only the maximum-flow algorithm was used to determine the value of the nodes. Given the serial arrangement of the network nodes, an interruption on the main branch would affect any flow upstream with little ability to find new sources. Therefore, the interdiction algorithm was considered of limited utility in this example. Four different scenarios were created for sensitivity analysis. The parameters for defining the scenarios were the same as was done for the mid-scale example, with fixed defenders either targeted to specific nodes or uniformly distributed among all the nodes, and the attackers either with uniformly distributed strategies or favoring PR targets. In each of these scenarios, there were 170 defenders who were fixed by assigning them to specific nodes. An additional 400 defenders were allowed to move from node to node, depending upon their strategies. There were 500 terrorists in this analysis. This was the first model analyzed that had an imbalance between terrorist assets and defender assets. All the terrorists were free to evolve their strategies. In some scenarios, the fixed, or inherent, defenders were allocated in a non-varying but non-uniform manner; i.e., some nodes had more inherent defenders than others. In other scenarios, they were allocated uniformly with two inherent defenders per node. The other assumption that varied was that in some of the scenarios the terrorists had a uniform probability of using any of the tactics, and in the other case the probability that they would favor PR was 0.5, with the remaining probability uniformly distributed among the other tactics, in a similar manner previously described for the mid-scale network example. The PR values were arbitrarily assigned and did not vary from one scenario to another. Table 13 provides a summary of the varying values for the scenarios, and Table 14 provides a summary of the actual values used.

Table 13 Large-Scale Scenarios

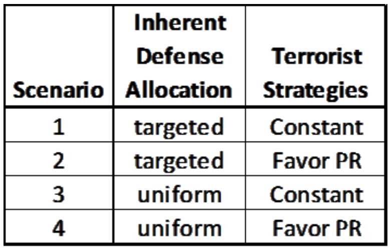

Scenario	Inherent Defense Allocation	Terrorist Strategies
1	targeted	Constant
2	targeted	Favor PR
3	uniform	Constant
4	uniform	Favor PR

Table 14 Large-Scale Network Parameters

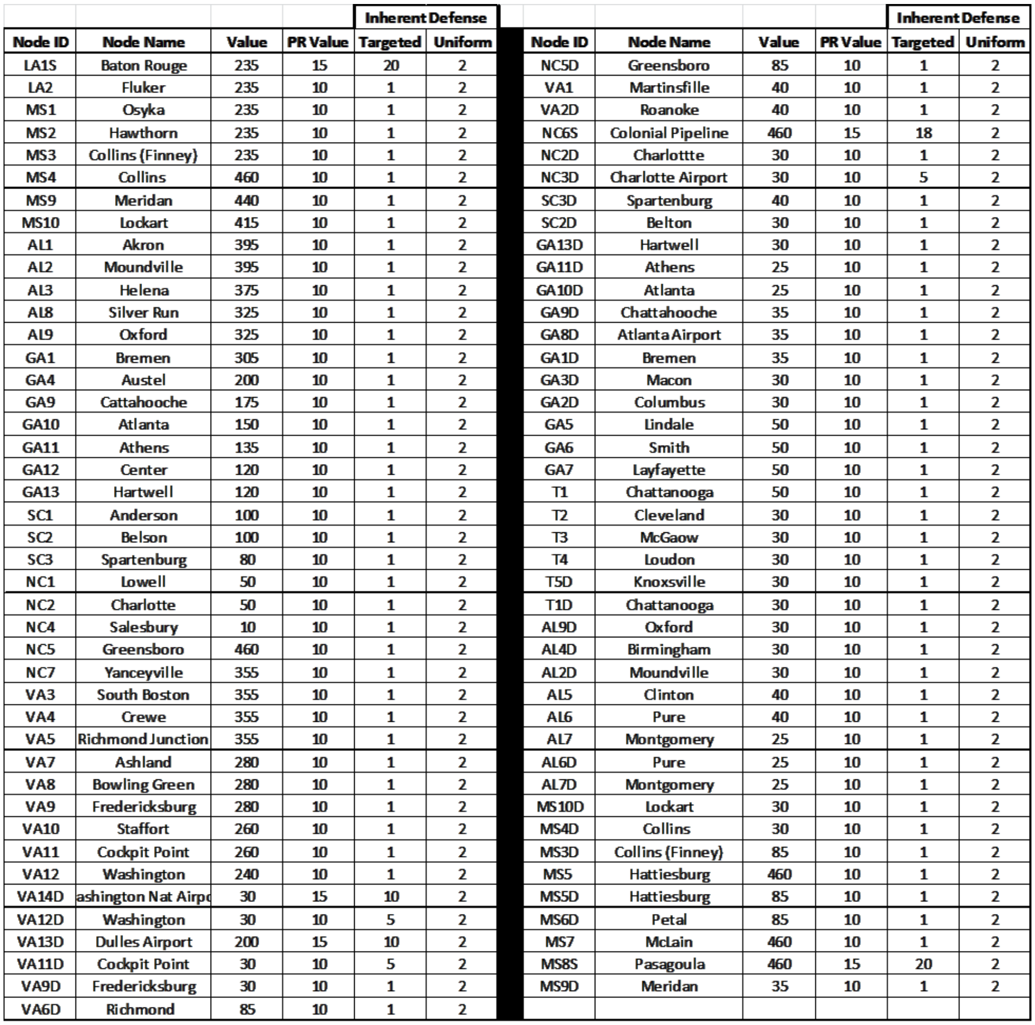

				Inherent Defense	
Node ID	Node Name	Value	PR Value	Targeted	Uniform
LA1S	Baton Rouge	235	15	20	2
LA2	Fluker	235	10	1	2
MS1	Osyka	235	10	1	2
MS2	Hawthorn	235	10	1	2
MS3	Collins (Finney)	235	10	1	2
MS4	Collins	460	10	1	2
MS9	Meridan	440	10	1	2
MS10	Lockart	415	10	1	2
AL1	Akron	395	10	1	2
AL2	Moundville	395	10	1	2
AL3	Helena	375	10	1	2
AL8	Silver Run	325	10	1	2
AL9	Oxford	325	10	1	2
GA1	Bremen	305	10	1	2
GA4	Austel	200	10	1	2
GA9	Cattahooche	175	10	1	2
GA10	Atlanta	150	10	1	2
GA11	Athens	135	10	1	2
GA12	Center	120	10	1	2
GA13	Hartwell	120	10	1	2
SC1	Anderson	100	10	1	2
SC2	Belson	100	10	1	2
SC3	Spartenburg	80	10	1	2
NC1	Lowell	50	10	1	2
NC2	Charlotte	50	10	1	2
NC4	Salesbury	10	10	1	2
NC5	Greensboro	460	10	1	2
NC7	Yanceyville	355	10	1	2
VA3	South Boston	355	10	1	2
VA4	Crewe	355	10	1	2
VA5	Richmond Junction	355	10	1	2
VA7	Ashland	280	10	1	2
VA8	Bowling Green	280	10	1	2
VA9	Fredericksburg	280	10	1	2
VA10	Staffort	260	10	1	2
VA11	Cockpit Point	260	10	1	2
VA12	Washington	240	10	1	2
VA14D	ashington Nat Airpo	30	15	10	2
VA12D	Washington	30	10	5	2
VA13D	Dulles Airport	200	15	10	2
VA11D	Cockpit Point	30	10	5	2
VA9D	Fredericksburg	30	10	1	2
VA6D	Richmond	85	10	1	2
NC5D	Greensboro	85	10	1	2
VA1	Martinsfille	40	10	1	2
VA2D	Roanoke	40	10	1	2
NC6S	Colonial Pipeline	460	15	18	2
NC2D	Charlottte	30	10	1	2
NC3D	Charlotte Airport	30	10	5	2
SC3D	Spartenburg	40	10	1	2
SC2D	Belton	30	10	1	2
GA13D	Hartwell	30	10	1	2
GA11D	Athens	25	10	1	2
GA10D	Atlanta	25	10	1	2
GA9D	Chattahooche	35	10	1	2
GA8D	Atlanta Airport	35	10	1	2
GA1D	Bremen	35	10	1	2
GA3D	Macon	30	10	1	2
GA2D	Columbus	30	10	1	2
GA5	Lindale	50	10	1	2
GA6	Smith	50	10	1	2
GA7	Layfayette	50	10	1	2
T1	Chattanooga	50	10	1	2
T2	Cleveland	30	10	1	2
T3	McGaow	30	10	1	2
T4	Loudon	30	10	1	2
T5D	Knoxsville	30	10	1	2
T1D	Chattanooga	30	10	1	2
AL9D	Oxford	30	10	1	2
AL4D	Birmingham	30	10	1	2
AL2D	Moundville	30	10	1	2
AL5	Clinton	40	10	1	2
AL6	Pure	40	10	1	2
AL7	Montgomery	25	10	1	2
AL6D	Pure	25	10	1	2
AL7D	Montgomery	25	10	1	2
MS10D	Lockart	30	10	1	2
MS4D	Collins	30	10	1	2
MS3D	Collins (Finney)	85	10	1	2
MS5	Hattiesburg	460	10	1	2
MS5D	Hattiesburg	85	10	1	2
MS6D	Petal	85	10	1	2
MS7	McLain	460	10	1	2
MS8S	Pasagoula	460	15	20	2
MS9D	Meridan	35	10	1	2

The allocation model was run for the four scenarios (combinations of inherent defense allocations and terrorist strategies) and the best allocation determined from those model runs. The allocation is called the Game Theoretic (GT) allocation. The comparison model was then run on the PRA allocation and on the GT allocation for all four scenarios and the results compared.

7.2 Determination of Resource Allocations

The resource allocation model was run for 2000 repetitions for each of these four scenarios. Examination of the results showed that the method of analysis used on the small- and mid-scale network examples was applicable. The allocation of resources stabilized after 1000 repetitions, and the average number of defenders at each node over the last 500 repetitions was used as the basis for a distribution of results. The results for these cases can be seen in Figure 38 to Figure 41. The number of defender victories refers to the number of nodes that were successfully defended. The GT and PRA allocations were calculated as discussed in the small network example. Table 15 lists the final allocations that were calculated.

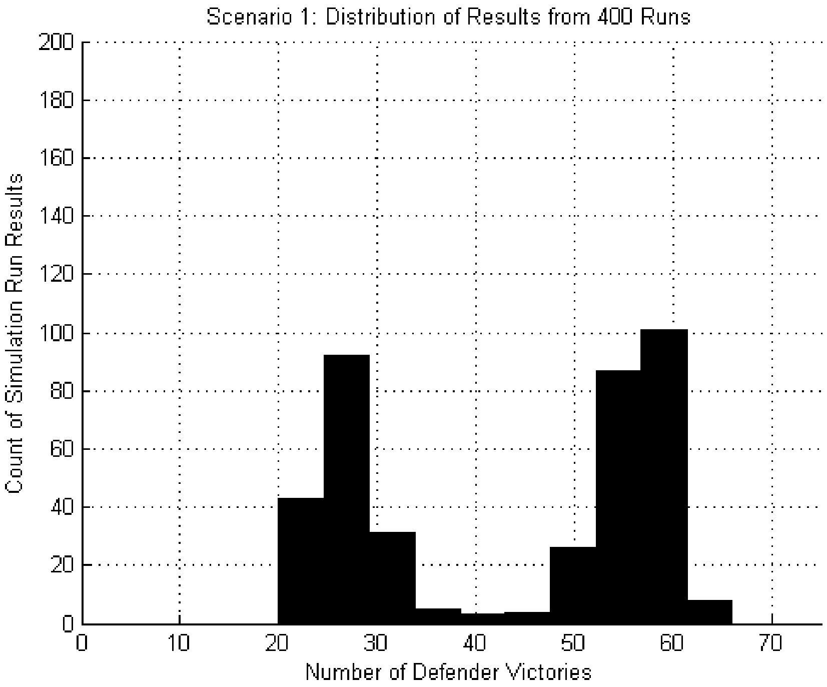

Figure 38 Large-Scale Example Scenario 1 Histogram

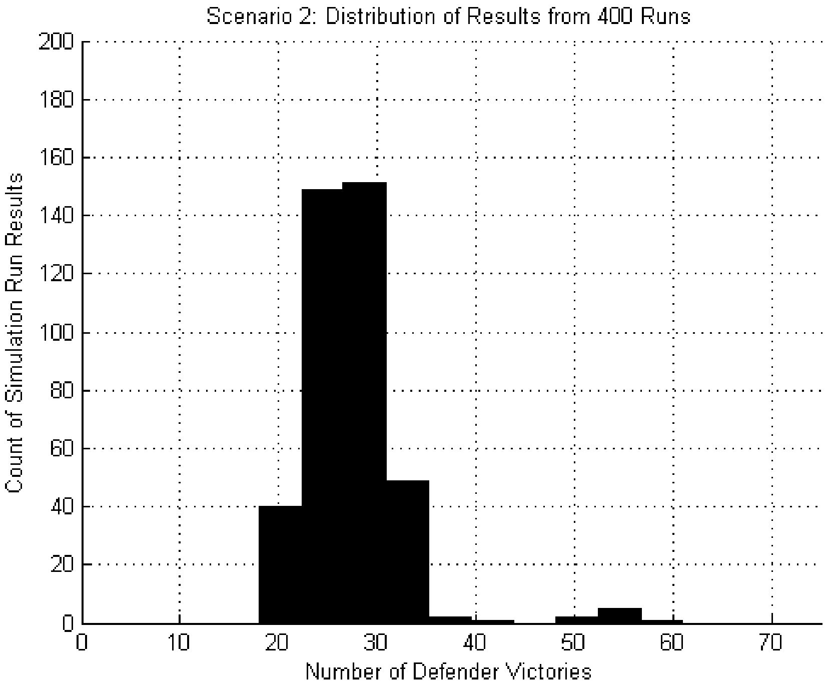

Figure 39 Large-Scale Network Scenario 2 Histogram

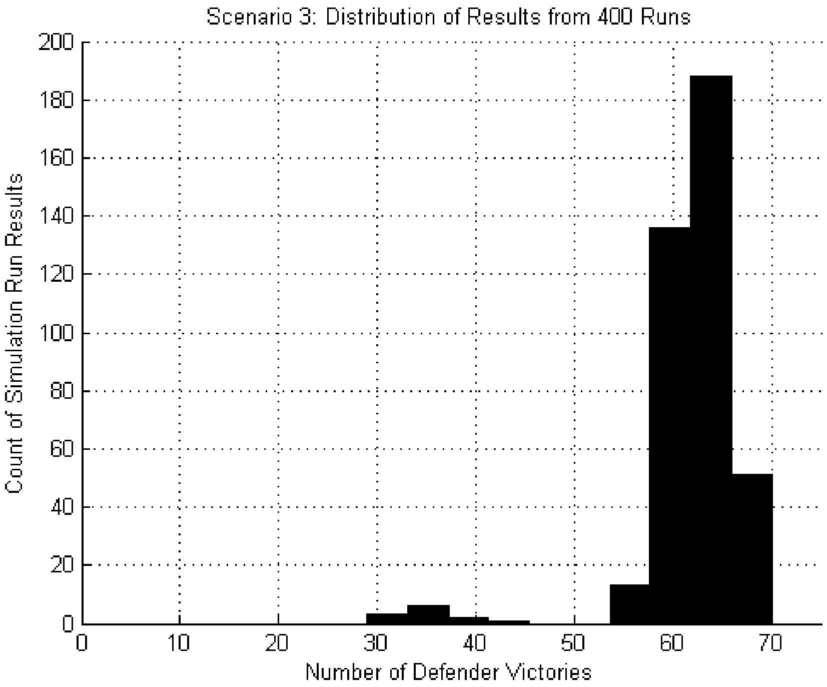

Figure 40 Large-Scale Network Scenario 3 Histogram

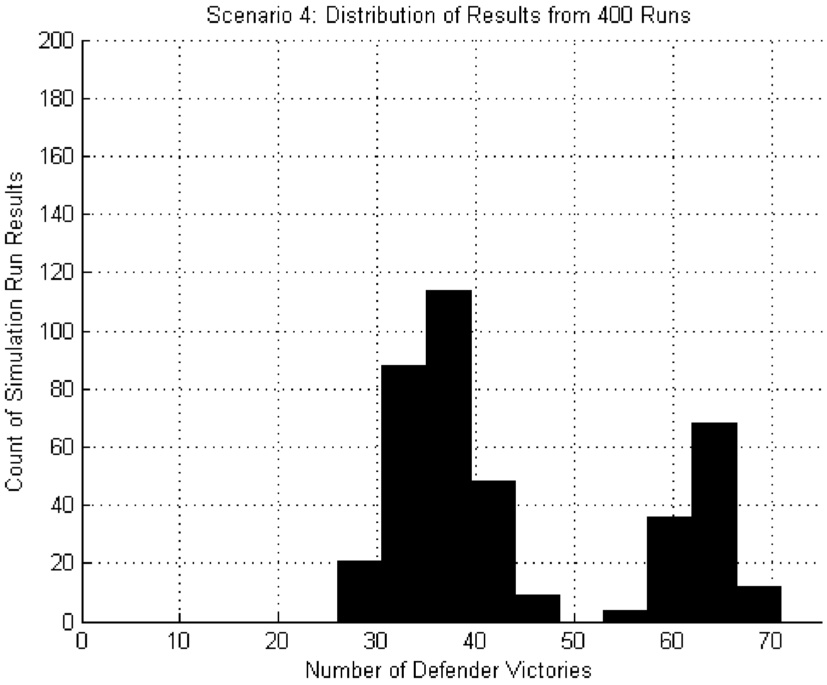

Figure 41 Large-Scale Network Scenario 4 Histogram

Table 15 Large-Scale Network Resource Allocations

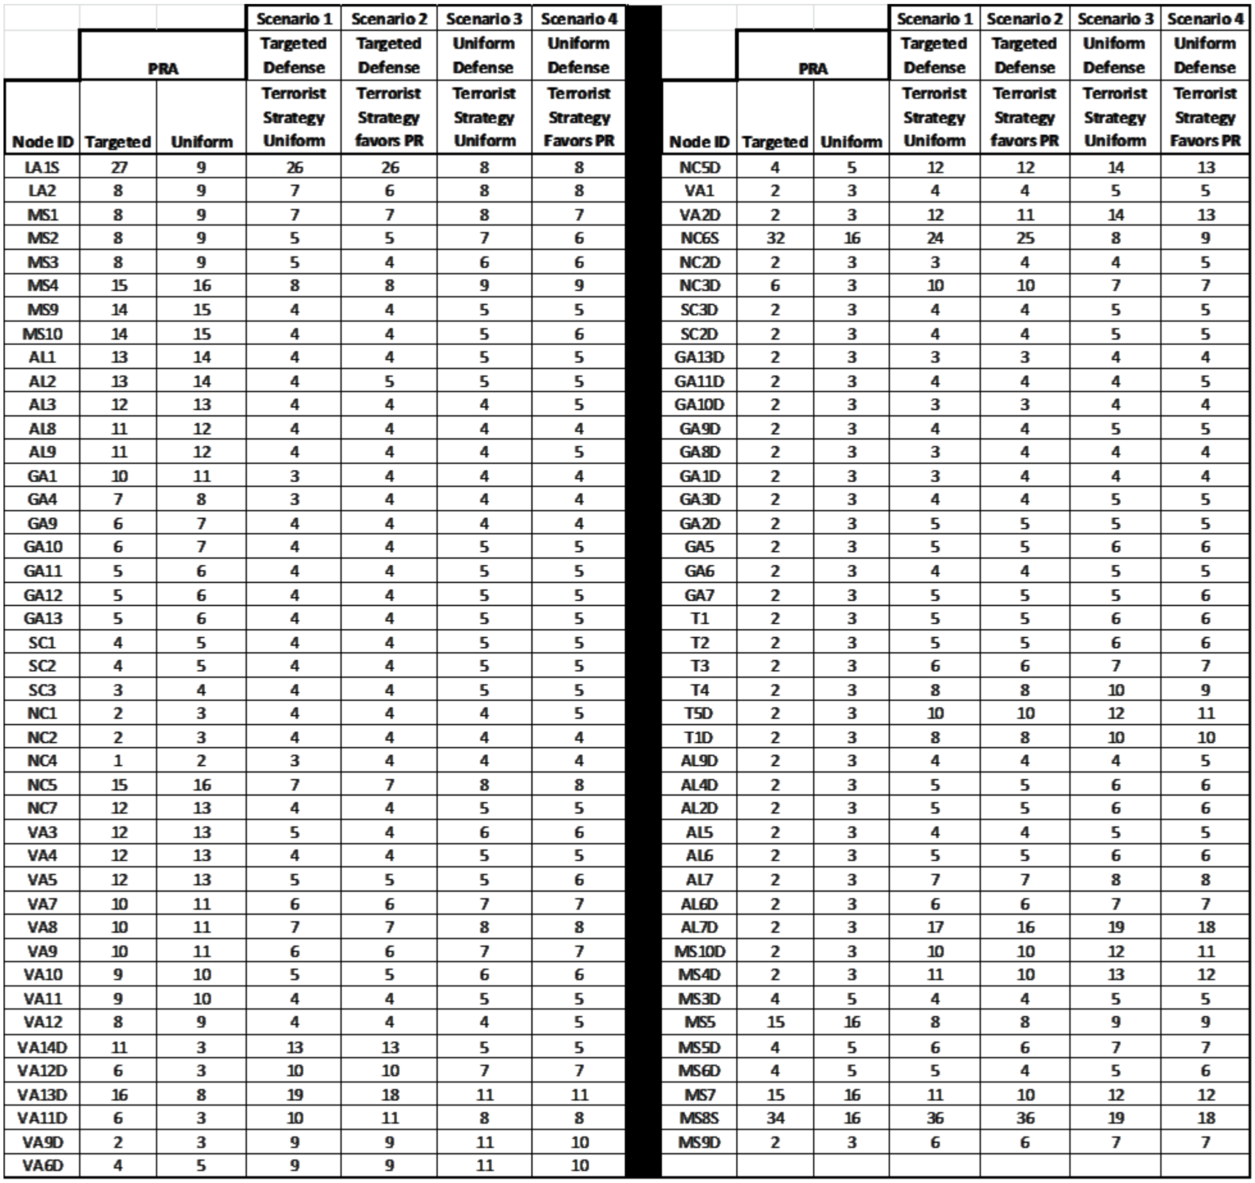

			Scenario 1	Scenario 2	Scenario 3	Scenario 4
	PRA		Targeted Defense	Targeted Defense	Uniform Defense	Uniform Defense
Node ID	Targeted	Uniform	Terrorist Strategy Uniform	Terrorist Strategy favors PR	Terrorist Strategy Uniform	Terrorist Strategy Favors PR
LA1S	27	9	26	26	8	8
LA2	8	9	7	6	8	8
MS1	8	9	7	7	8	7
MS2	8	9	5	5	7	6
MS3	8	9	5	4	6	6
MS4	15	16	8	8	9	9
MS9	14	15	4	4	5	5
MS10	14	15	4	4	5	6
AL1	13	14	4	4	5	5
AL2	13	14	4	5	5	5
AL3	12	13	4	4	4	5
AL8	11	12	4	4	4	4
AL9	11	12	4	4	4	5
GA1	10	11	3	4	4	4
GA4	7	8	3	4	4	4
GA9	6	7	4	4	4	4
GA10	6	7	4	4	5	5
GA11	5	6	4	4	5	5
GA12	5	6	4	4	5	5
GA13	5	6	4	4	5	5
SC1	4	5	4	4	5	5
SC2	4	5	4	4	5	5
SC3	3	4	4	4	5	5
NC1	2	3	4	4	4	5
NC2	2	3	4	4	4	4
NC4	1	2	3	4	4	4
NC5	15	16	7	7	8	8
NC7	12	13	4	4	5	5
VA3	12	13	5	4	6	6
VA4	12	13	4	4	5	5
VA5	12	13	5	5	5	6
VA7	10	11	6	6	7	7
VA8	10	11	7	7	8	8
VA9	10	11	6	6	7	7
VA10	9	10	5	5	6	6
VA11	9	10	4	4	5	5
VA12	8	9	4	4	4	5
VA14D	11	3	13	13	5	5
VA12D	6	3	10	10	7	7
VA13D	16	8	19	18	11	11
VA11D	6	3	10	11	8	8
VA9D	2	3	9	9	11	10
VA6D	4	5	9	9	11	10
NC5D	4	5	12	12	14	13
VA1	2	3	4	4	5	5
VA2D	2	3	12	11	14	13
NC6S	32	16	24	25	8	9
NC2D	2	3	3	4	4	5
NC3D	6	3	10	10	7	7
SC3D	2	3	4	4	5	5
SC2D	2	3	4	4	5	5
GA13D	2	3	3	3	4	4
GA11D	2	3	4	4	4	5
GA10D	2	3	3	3	4	4
GA9D	2	3	4	4	5	5
GA8D	2	3	3	4	4	4
GA1D	2	3	3	4	4	4
GA3D	2	3	4	4	5	5
GA2D	2	3	5	5	5	5
GA5	2	3	5	5	6	6
GA6	2	3	4	4	5	5
GA7	2	3	5	5	5	6
T1	2	3	5	5	6	6
T2	2	3	5	5	6	6
T3	2	3	6	6	7	7
T4	2	3	8	8	10	9
T5D	2	3	10	10	12	11
T1D	2	3	8	8	10	10
AL9D	2	3	4	4	4	5
AL4D	2	3	5	5	6	6
AL2D	2	3	5	5	6	6
AL5	2	3	4	4	5	5
AL6	2	3	5	5	6	6
AL7	2	3	7	7	8	8
AL6D	2	3	6	6	7	7
AL7D	2	3	17	16	19	18
MS10D	2	3	10	10	12	11
MS4D	2	3	11	10	13	12
MS3D	4	5	4	4	5	5
MS5	15	16	8	8	9	9
MS5D	4	5	6	6	7	7
MS6D	4	5	5	4	5	6
MS7	15	16	11	10	12	12
MS8S	34	16	36	36	19	18
MS9D	2	3	6	6	7	7

7.3 Comparison of Allocations

Using the same run parameters (2000 repetitions per run, results taken as the average of the last 500 repetitions, 400 runs for each allocation and scenario) the comparison model was then run for each of these scenarios. The comparison model examined the performance of the PRA allocation and the GT allocation against an evolving terrorist threat in all four of the scenarios. The results were examined for normality using QQ plots and tested using the Lilliefors test. The QQ plots are shown in Figure 42 to Figure 45. The results of the Lilliefors tests for normality are shown in Table 16.

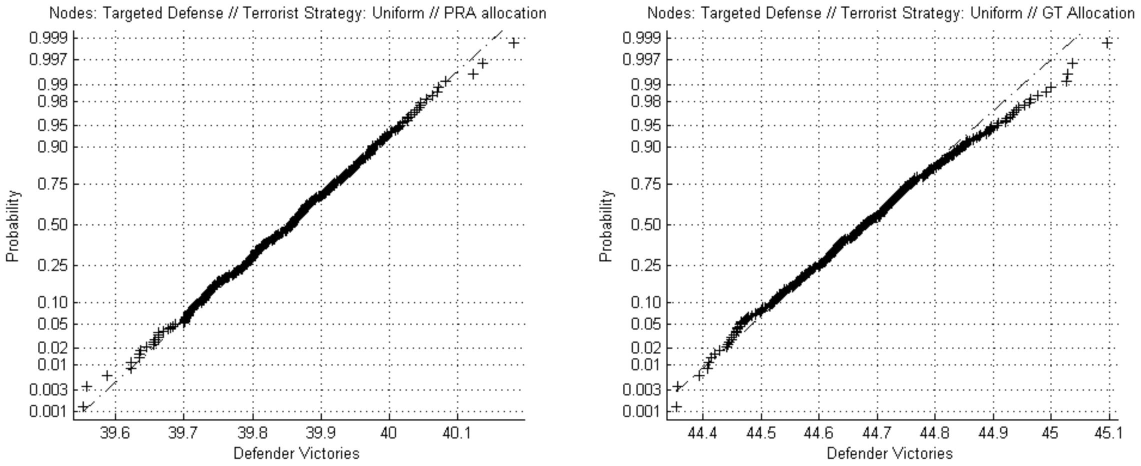

Figure 42 Large-Scale Network Scenario 1 QQ Plot

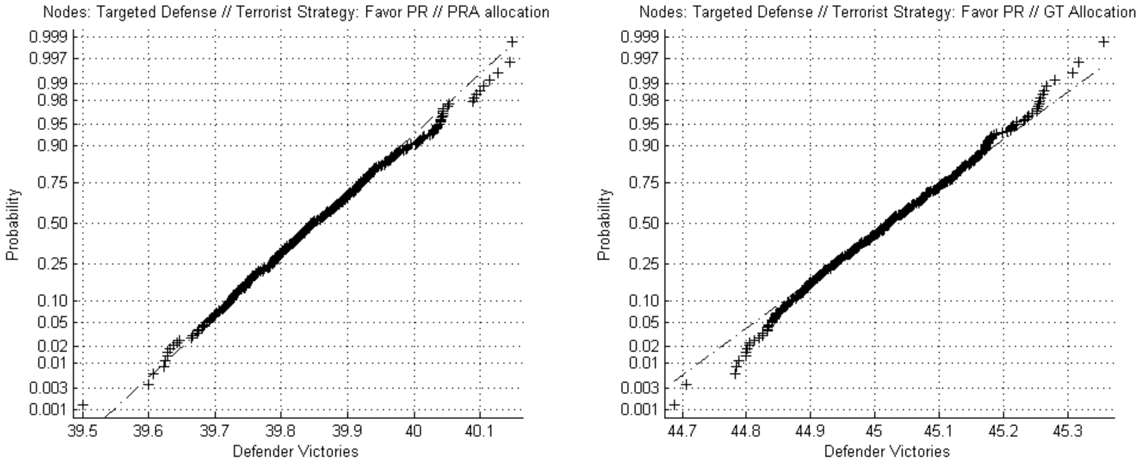

Figure 43 Large-Scale Network Scenario 2 QQ Plot

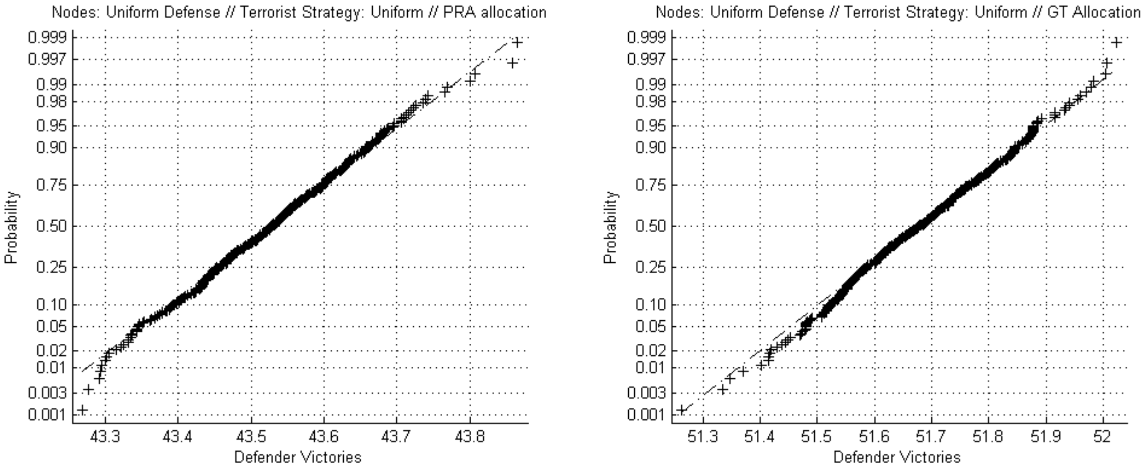

Figure 44 Large-Scale Network Scenario 3 QQ Plot

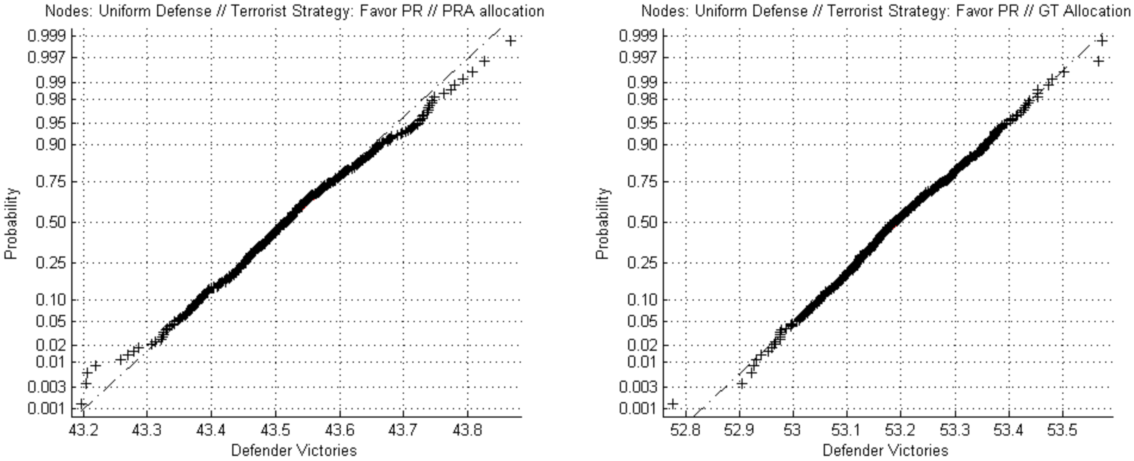

Figure 45 Large-Scale Network Scenario 4 QQ Plot

Table 16 Large-Scale Network Lilliefors Test Result

Lilliefors Test Results			
Scenario	Allocation	p value	Result
1	PRA	0.4521	Do Not Reject H_0
2	PRA	0.3750	Do Not Reject H_0
3	PRA	>.5	Do Not Reject H_0
4	PRA	0.0457	Reject H_0
1	GT	0.1753	Do Not Reject H_0
2	GT	>.5	Do Not Reject H_0
3	GT	>.5	Do Not Reject H_0
4	GT	0.1166	Do Not Reject H_0

For each of the scenarios, the CDF of the number of defender victories observed with the PRA-based allocation and the CDF of the number of defender victories observed with the GT-based allocation was compared. These results are shown in Figure 46 to Figure 49.

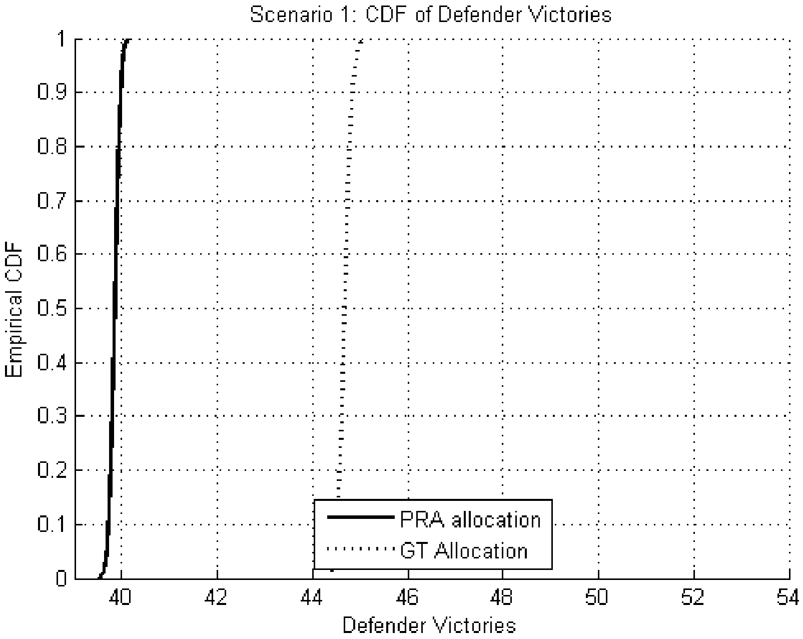

Figure 46 Large-Scale Network Scenario 1 CDF

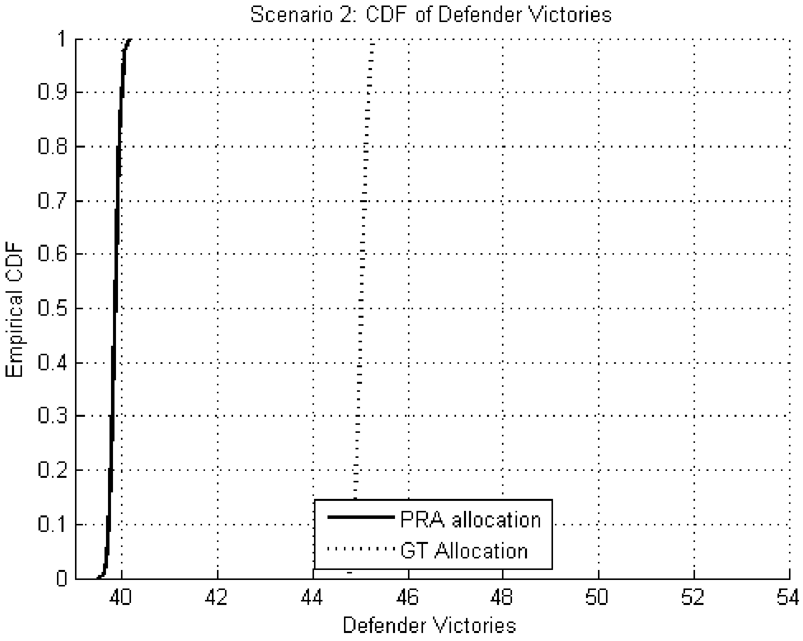

Figure 47 Large-Scale Network Scenario 2 CDF

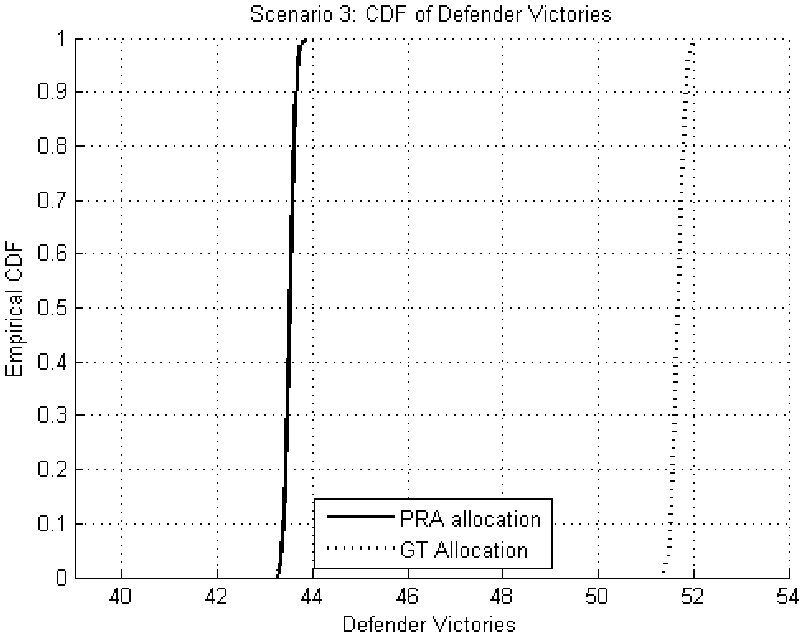

Figure 48 Large-Scale Network Scenario 3 CDF

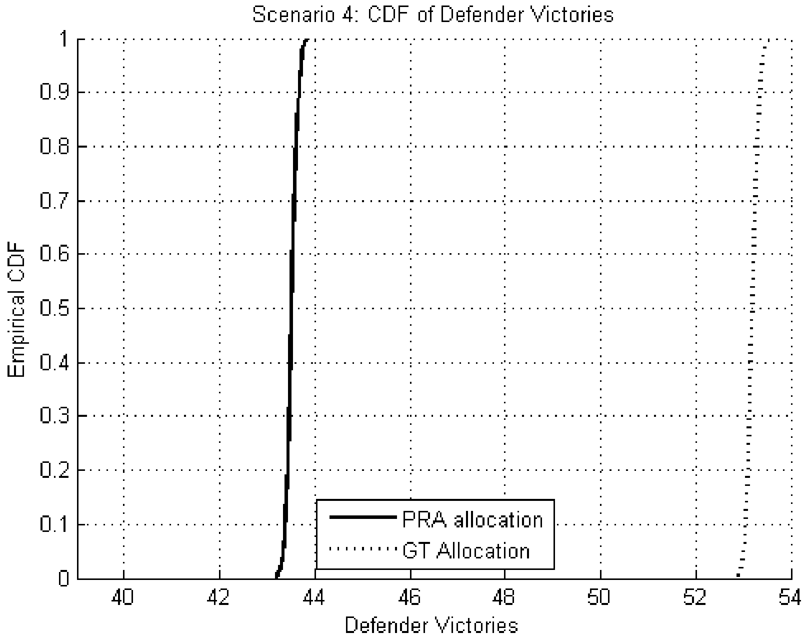

Figure 49 Large-Scale Network Scenario 4 CDF

Examination of the CDF plots in Figure 46 to Figure 49 shows that, in every case the PRA-based allocation is dominated by the GT-based allocation in the sense that, for any specified value of the CDF, the GT-based approach always produces a higher number of defender victories than does the PRA-based approach. These results are consistent with the results for the small and mid-scale networks presented in the previous two chapters. The differences in mean are statistically significant. The means test result is shown in Table 17.

Table 17 Large-Scale Network Means Test Results

Mean Test Result		
Scenario	**T statistic**	**Result**
1	-594.3224	Reject H_0
2	-658.4476	Reject H_0
3	-992.697	Reject H_0
4	-1188.4	Reject H_0

Chapter 8: Summary and Conclusions

8.1 Summary

The goal of this research was to determine if there was a useful way to formally bring the concept of strategic uncertainty (ARA-based approach) into the allocation of resources to protect spatially distributed networks against a planning adversary that would allow flexibility for multiple adversary goals and multiple methods of network node valuation. A hypothesis was developed that the insights gained through the use of an evolutionary agent-based solution to the “leader–follower” game would provide an allocation that was at least as good as, and hopefully better than, an allocation based strictly on a more traditional probabilistic risk analysis approach (PRA-based approach) using the value of the network nodes, determined in some logical manner related to their importance to the network structure.

After finding a dearth of literature on the application of an evolutionary agent-based model to this problem, such a model was developed. Currently, this model assumes that there is a single value for the node’s contribution to the network’s mission, and it has another value relating to the publicity that attacking such a node may generate. This model assumes that both the attacker and defender of the network know the values of the nodes, and that, for each node, the values are the same for both the attacker and defender. In all cases, the model followed the “leader–follower” or “attacker–defender” model in which the defender made allocations without the knowledge of the attacker, but the attacker made their deployment of resources with the full knowledge of the defender’s dispositions.

Due to the difficulty of obtaining actual data on attacks—and the fact that any single attack would only represent a single realization of attacks against that particular network—simulation was used to test the allocations. For each scenario, a GT-based allocation was developed using this allocation model. A PRA-based allocation was also developed using the values of the nodes. These two allocations were then each run in an evaluation model that compared their results against an adversary who could adapt to the network's defenses. Three differently sized networks were used: a small network of 13 nodes, a mid-sized network of 31 nodes, and a larger network of 85 nodes. The small network was chosen to represent a network topology that would allow the ready use of both a network flow model and a network interdiction model to determine node values. The mid-size network was based on publically available data on the Irish power grid, and the large network was based on the Plantation Pipeline, which runs along the East Coast of the United States. The large network was chosen to be a network based on a main "trunk" with several branches, rather than a ring or network rich with paths and interconnections. Both the mid-sized and large networks had their node values taken from a network flow model. The defenders of the mid-sized and large network consisted of some "inherent" defenders that were always assigned to the node, and some other defenders who could move from node to node on each repetition as their strategy evolved.

In the small network, three different allocations of PR value to the nodes were tested. One was considered a baseline, one favored source nodes, and one favored sink nodes. In the medium and large networks, the PR values were held fixed, but the inherent defense was varied: one that was uniformly distributed among all the nodes, and a second where the same number of resources were allocated, but some nodes received a higher defense value than others. Finally, two different distributions of terrorist strategies were examined. In the first case, the starting strategies for all the terrorists were uniformly distributed among all the possible strategies. In the second case, the terrorists greatly favored attacks against targets with PR values.

For all cases, the resulting distributions of the percentage of defender victories (expressed as the number of nodes where the defender won) was examined and compared. A Lilliefors test for normality of results was run, and the CDFs were compared. Finally, a test of means for the distribution of the proportion of defender victories was run to see if the differences observed were statistically significant at the 5% level. Rejection of the null hypothesis that $\mu_{PRA} = \mu_{GT}$ indicates that the difference in average number of defender victories was statistically significant. The results are summarized in Table 18.

Table 18 Summary of All Run Test Results

Network	Scenario	Inherenet Defense	Publicity Values	Terrorist Strategy	Allocation Model	Lilliefors P value	Lilliefors Result α-.05	Mean Defender Victories	Median Defender Victories	Means Test t Statistic	Means Test Result α-.05
Small 13 Nodes	1	None	Baseline	Even	PRA	0.4122	Do Not Reject H_0	5.76	5.76	-332.1778	Reject H_0
					GT	>.5	Do Not Reject H_0	6.85	6.85		
	2	None	Sinks High	Even	PRA	>.5	Do Not Reject H_0	5.76	5.76	-326.0212	Reject H_0
					GT	>.5	Do Not Reject H_0	6.89	6.89		
	3	None	Sources High	Even	PRA	>.5	Do Not Reject H_0	5.76	5.76	-341.3158	Reject H_0
					GT	>.5	Do Not Reject H_0	6.94	6.93		
	4	None	Baseline	Favor PR	PRA	>.5	Do Not Reject H_0	5.76	5.76	-320.6159	Reject H_0
					GT	0.0161	Reject H_0	6.86	6.86		
	5	None	Sinks High	Favor PR	PRA	>.5	Do Not Reject H_0	5.76	5.76	-310.9925	Reject H_0
					GT	0.4607	Do Not Reject H_0	6.83	6.84		
	6	None	Sources High	Favor PR	PRA	>.5	Do Not Reject H_0	5.76	5.76	-359.9671	Reject H_0
					GT	0.1122	Do Not Reject H_0	6.97	6.97		
Mid 31 Nodes	1	Targeted	Baseline	Even	PRA	>.5	Do Not Reject H_0	14.26	14.26	-131.3408	Reject H_0
					GT	>.5	Do Not Reject H_0	14.83	14.83		
	2	Targeted	Baseline	Favor PR	PRA	>.5	Do Not Reject H_0	14.27	14.27	-118.095	Reject H_0
					GT	>.5	Do Not Reject H_0	14.78	14.78		
	3	Uniform	Baseline	Even	PRA	>.5	Do Not Reject H_0	14.26	14.27	-216.9933	Reject H_0
					GT	>.5	Do Not Reject H_0	15.36	15.36		
	4	Uniform	Baseline	Favor PR	PRA	>.5	Do Not Reject H_0	14.26	14.27	-217.7804	Reject H_0
					GT	0.2404	Do Not Reject H_0	15.37	15.37		
Large 85 Nodes	1	Targeted	Baseline	Even	PRA	0.4521	Do Not Reject H_0	39.85	39.86	-594.3224	Reject H_0
					GT	0.1753	Do Not Reject H_0	44.68	44.68		
	2	Targeted	Baseline	Favor PR	PRA	0.375	Do Not Reject H_0	39.85	39.85	-658.4476	Reject H_0
					GT	>.5	Do Not Reject H_0	45.02	45.02		
	3	Uniform	Baseline	Even	PRA	>.5	Do Not Reject H_0	43.53	43.53	-992.697	Reject H_0
					GT	>.5	Do Not Reject H_0	51.68	51.68		
	4	Uniform	Baseline	Favor PR	PRA	0.0457	Reject H_0	43.52	43.52	-1188.4	Reject H_0
					GT	0.1166	Do Not Reject H_0	53.20	53.19		

8.2 Conclusions

In all the cases examined, the allocation of resources determined by the GT approach dominated the allocation of resources determined by the PRA approach. This improvement grew larger as the number of nodes in the network grew. The assumption of normality of results held at the 5% level for all but two of the 28 cases run. While no simulation can “prove” what will happen in reality, this research has shown that using a GT approach to resource allocation based on an evolutionary agent-based model shows distinct promise. Under simulated conditions, the agent-based evolutionary allocation method clearly performed better that a PRA-based allocation in all scenarios tested, varying the network size, complexity, topology, terrorist tactics mix, node valuation method, and fixed defensive resource allocations.

8.3 Recommendations for Further Research

There are several natural outgrowths of this research that should be subject to further study. Broadly, they fall into four areas: (1) refinements on the terrorist tactics, (2) refinements on the network valuations, (3) continued research into methods that will scale effectively to even larger networks, and (4) evaluation of proposed networks.

The terrorists' strategies examined in the effort should be considered as starting points for further research. The model assumed that the terrorist would either go to a random node and attack or attack a node based on its PR value, its value to the network, its physical proximity to the terrorist, or some combination of node value and proximity. There have been several highly interesting works published in the last few years about using agents to determine how terrorists organize themselves. A fusion of these efforts with the resource allocation capability developed in this dissertation would be most interesting and would provide a means of examining the resource allocation problem in a more "end-to-end" approach.

Examples run in this research considered only two possible valuations of network: one for "flow" or value of the node in delivering product, and one for publicity generated by an attack. Nothing inherent in the model or method limits the valuations to only two. Other dimensions of value to terrorists (and defenders) need to be explored. In this work, the values of the nodes (both flow and PR) were known to both sides in the conflict. Variations on the scope of each opponent's knowledge, and their valuation functions, should be further explored. Other methods of evaluating each node's contribution to the network should be examined. Because a physical network is also vulnerable to natural disaster, a comparison of the PRA- and GT-based allocations under a purely random act of nature is another obvious area to study. Another obvious area for further study would be to compare the PRA and GT allocations when the PRA allocation is based on some combination of node values that combined the PR value and flow value of the node in some logical and realistic manner. The research presented in this dissertation is also based on a single time step. Extension of these techniques to a multi-time period "phased" attack is also a natural area for further exploration. The use of different metrics (such as total flow through the network after attack) as opposed to the number of nodes successfully defended would also be of interest.

In terms of scaling to ever larger networks, more efficient simulation models could be developed. As the network size grows, comparisons to the "attacker–defender" linear programming solutions would be very interesting and informative. Finally, this work could prove to be a basis for the evaluation of proposed network architectures, evaluating the resilience of different architectures to attack by comparing them under identical scenarios of attacker and defender resources.

References

Agah, A., Basu, K., & Das, S. K. (2005, 16-20 May 2005). *Preventing DoS attack in sensor networks: a game theoretic approach.* Paper presented at the Communications, 2005. ICC 2005. 2005 IEEE International Conference on.

An, B., Shieh, E., Yang, R., Tambe, M., Baldwin, C., DiRenzo, J., . . . Meyer, G. (2012). PROTECT - A Deployed Game-Theoretic System for Strategic Security Allocation for the United States Coast Guard. *AI Magazine, 33*(4), 96-110.

Anderton, C. H., & Carter, J. R. (2006). Applying Intermediate Microeconomics to Terrorism. *Journal of Economic Education, 37*(4), 442-458.

Arce, D., & Sandler, T. (2005). Counterterrorism: A Game-Theoretic Analysis. *The Journal of Conflict Resolution, 49*(2), 183.

Axelrod, R. (1984). *The Evolution of Cooperation*. Cambridge: Basic Books.

Ayyub, B. M., McGill, W. L., & Kaminskiy, M. (2007). Critical asset and portfolio risk analysis: An all-hazards framework. *Risk Analysis, 27*(4), 789-801. doi: 10.1111/j.1539-6924.2007.00911.x

Banks, D., Petralia, F., & Wang, S. (2011). Adversarial risk analysis: Borel games. *Applied Stochastic Models in Business and Industry, 27*(2), 72-86. doi: 10.1002/asmb.890

Bazaraa, M. S., Jarvis, J. J., & Sherali, H. D. (1990). *Linear Programming and Netork Flows* (2 ed.). New York: John Wiley & Sons.

Beitel, G. A., Gertman, D. I., & Plum, M. M. (2004). *Balanced scorecard method for predicting the probability of a terrorist attack.* Paper presented at the Fourth International Conference on Computer Simulation in Risk Analysis and Hazard Mitigation, RISK ANALYSIS IV, September 27, 2004 - September 29, 2004, Rhodes, Greece.

Bier, V. (2007). Choosing What to Protect. *Risk Analysis, 27*(3), 607.

Bier, V., Haphuriwat, N., Menoyo, J., Zimmerman, R., & Culpen, A. (2008). Optimal Resource Allocation for Defense of Targets Based on Differing Measures of Attractiveness. *Risk Analysis, 28*(3), 763.

Bier, V., Nagaraj, A., & Abhichandani, V. (2005). Optimal Allocation of Resources for Defense of Simple Series and Parallel Systems from Determined Adversaries. *Reliability Engineering & System Safety, 87*, 313-323.

Brown, G., Carlyle, M., Diehl, D., Kline, J., & Wood, K. (2005). A Two-Sided Optimization for Theater Ballistic Missile Defense. *OPERATIONS RESEARCH, 53*(5), 745-763. doi: 10.1287/opre.1050.0231

Brown, G., Carlyle, M., Salmerón, J., & Wood, K. (2006). Defending Critical Infrastructure. *INTERFACES, 36*(6), 530-544.

Brown, G., & Cox, J. L. (2011). How Probabilistic Risk Assessment Can Mislead Terrorism Risk Analysts. *Risk Analysis, 31*(2), 196.

Chen, Y., Hobbs, B., Leyffer, S., & Munson, T. (2006). Leader-Follower Equilibria for Electric Power and NO x Allowances Markets. *Computational Management Science, 3*(4), 307-330. doi: 10.1007/s10287-006-0020-1

Conitzer, V., & Sandholm, T. (2006). *Computing the Optimal Strategy to Commit to*.

Cox, J. L. (2009). Game Theory and Risk Analysis. *Risk Analysis, 29*(8), 1062.

Cox, L. A. T., Jr. (2008). Some limitations of "Risk = Threat x Vulnerability x Consequence" for risk analysis of terrorist attacks. *Risk Analysis: An Official Publication Of The Society For Risk Analysis, 28*(6), 1749-1761.

Dempe, S., Kalashnikov, V., & Rios-Mercado, R. Z. (2005). Discrete bilevel programming: Application to a natural gas cash-out problem. *European Journal of Operational Research, 166*(2), 469-488.

Dillon, R. L., Liebe, R. M., & Bestafka, T. (2009). Risk-based decision making for terrorism applications. *Risk Analysis, 29*(3), 321-335. doi: 10.1111/j.1539-6924.2008.01196.x

Ezell, B. C., Bennett, S. P., Von Winterfeldt, D., Sokolowski, J., & Collins, A. J. (2010). Probabilistic Risk Analysis and Terrorism Risk. *Risk Analysis: An International Journal, 30*(4), 575-589. doi: 10.1111/j.1539-6924.2010.01401.x

GAMS Development Corporation. (2009). General Algebraic Modeling System. Washington, DC: GAMS Development Corporation. Retrieved from www.gams.com

Gilbert, N. (2008). *Agent-Based Models*. Los Angeles: Sage Publications.

Gintis, H. (2009). *Game Theory Evolving* (2 ed.). Princeton: Princeton University Press.

Golany, B., Kaplan, E. H., Marmur, A., & Rothblum, U. G. (2009). Nature plays with dice - terrorists do not: allocating resources to counter strategic versus probabilistic risks. *European Journal of Operational Research, 192*(Copyright 2008, The Institution of Engineering and Technology), 198-208.

Goldstein, H. (2006). Modeling Terrorists. *Spectrum, IEEE, 43*(9), 26-34.

Gurobi Optimization. (2010). GUROBI: Gurobi Optimization, Inc. Retrieved from www.gurobi.com

Haimes, Y. Y., & Longstaff, T. (2002). The role of risk analysis in the protection of critical infrastructures against terrorism. *Risk Analysis, 22*(3), 439-444. doi: 10.1111/0272-4332.00055

Hall Jr, J. R. (2009). The Elephant in the Room Is Called Game Theory. *Risk Analysis, 29*, 1061-1061. doi: 10.1111/j.1539-6924.2009.01246.x

Hong, S. (2011). *Strategic Network Interdiction*. Fondazione Eni Enrico Mattei, Working Papers: 2011.43. Retrieved from http://proxygw.wrlc.org/login?url=http://search.ebscohost.com/login.aspx?direct=true&db=ecn&AN=1235371&site=ehost-live

http://www.feem.it/userfiles/attach/2011611213444NDL2011-043.pdf

Insua, I. R., Rios, J., & Banks, D. (2009). Adversarial Risk Analysis. *Journal of the American Statistical Association, 104*(486), 841.

Kantzavelou, I., & Katsikas, S. (2010). A game-based intrusion detection mechanism to confront internal attackers. *Computers & Security, 29*(8), 859-874. doi: DOI: 10.1016/j.cose.2010.06.002

Lapan, H. E., & Sandler, T. (1988). To Bargain or Not To Bargain: That Is The Question. *The American Economic Review, 78*(2), 16-21.

Lye, K.-w., & Wing, J., M. (2005). Game strategies in network security. *International Journal of Information Security, 4*(1-2), 71.

Macal, C. M., & North, M. J. (2010). Tutorial on agent-based modelling and simulation. *J of Sim, 4*(3), 151-162.

Machado, R., & Tekinay, S. (2008). A survey of game-theoretic approaches in wireless sensor networks. *Computer Networks, 52*(16), 3047-3061. doi: DOI: 10.1016/j.gaceta.2008.07.003

Moore, J. T., & Bard, J. F. (1990). THE MIXED INTEGER LINEAR BILEVEL PROGRAMMING PROBLEM. *OPERATIONS RESEARCH, 38*(5), 911.

Myerson, R. B. (1991). *Game Theory: Analysis of Conflict*. Cambridge, Mass: Harvard University Press.

Nisan, N., Roughgarden, T., Tardos, E., & Vazirant, V. (2007). *Algorithmic Game Theory*. New York: Cambridge University Press.

Overgaard, P. B. (1994). The scale of terrorist attacks as a signal of resources. *Journal of Conflict Resolution, 38*, 452-478.

Paruchuri, P., Pearce, J. P., Marecki, J., Tambe, M., Ordonez, F., & Kraus, S. (2008). *Playing games for security: an efficient exact algorithm for solving Bayesian Stackelberg games*.

Paruchuri, P., Tambe, M., Ordóñez, F., & Kraus, S. (2006). *Security in multiagent systems by policy randomization*.

Pita, J., Jain, M., Ord\, F., \#243, \#241, ez, . . . Magori-Cohen, R. (2009). *Effective solutions for real-world Stackelberg games: when agents must deal with human uncertainties*. Paper presented at the Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 1, Budapest, Hungary.

Pita, J., Jain, M., Ordóñez, F., Portway, C., Tambe, M., Western, C., . . . Kraus, S. (2009). Using Game Theory for Los Angeles Airport Security. *AI Magazine, 30*(1), 43.

Plantation Pipeline System Map. (2012). Retrieved from http://www.kindermorgan.com/business/products_pipelines/plantation.cfm

Rios, J., & Insua, D. R. (2012). Adversarial Risk Analysis for Counterterrorism Modeling. *Risk Analysis, 32*(5), 894-915. doi: 10.1111/j.1539-6924.2011.01713.x

Rose, A., Benavides, J., Chang, S., Szczesniak, P., & Lim, D. (1997). The regional economic impact of an earthquake: Direct and indirect effects of electricity. *Journal of Regional Science, 37*(3), 437.

Rose, A., & Liao, S.-Y. (2005). Modeling Regional Economic Resilience to Disasters: A Computable General Equilibrium Analysis of Water Service Disruptions. *Journal of Regional Science, 45*(1), 75-112. doi: 10.1111/j.0022-4146.2005.00365.x

Rose, A., Oladosu, G., & Liao, S.-Y. (2007). Business interruption impacts of a terrorist attack on the electric power system of Los Angeles: Customer resilience to a total blackout. *Risk Analysis, 27*(3), 513-531. doi: 10.1111/j.1539-6924.2007.00912.x

Rothschild, C., McLay, L., & Guikema, S. (2012). Adversarial risk analysis with incomplete information: a level-kapproach. *Risk analysis : an official publication of the Society for Risk Analysis, 32*(7), 1219-1231. doi: 10.1111/j.1539-6924.2011.01701.x

Sandler, T., & Arce, D. (2003). Terrorism & game theory. *Simulation & Gaming, 34*(3), 319.

Sandler, T., & Lapan, H. E. (1988). The Calculus of Dissent: An Analysis of Terrorists' Choice of Targets. *Synthese, 76*(2), 245-261.

Sandler, T., & Siqueira, K. (2009). Games and Terrorism: Recent Developments. *Simulation & Gaming, 40*(2), 164.

Scaparra, M., & Church, R. (2008). A bilevel mixed-integer program for critical infrastructure protection planning. *Computers & Operations Research, 35*(6), 1905.

Schaffer, M. E. (1988). Evolutionarily stable strategies for a finite population and a variable contest size. *Journal of theoretical biology, 132*(4), 469; 469-478; 478. doi: 10.1016/S0022-5193(88)80085-7; pmid:

Schaffer, M. E. (1989). Are profit-maximisers the best survivors?: A Darwinian model of economic natural selection. *Journal of economic behavior & organization, 12*(1), 29-45. doi: http://dx.doi.org/10.1016/0167-2681(89)90075-9

Scott, S., & Koehler, M. (2011). *A Field Guide to NetLogo*. Book draft. Department of Computational Social Science. George Mason University.

Sheskin, D. J. (2000). *Handbook of Parametric and Nonparametric Statistical Procedures* (2nd ed.). Boca Raton: Chapman & Hall/CRC.

Shoham, Y., & Leyton-Brown, K. (2009). *Multiagent Systems: Algorithimic, Game-Theoretic, and Logical Foundations*. Cambridge: Cambridge University Press.

Siddiqui, S., & Gabriel, S. A. (2013). An SOS1-Based Approach for Solving MPECs with a Natural Gas Market Application. *Networks and Spatial Economics, 13*(2), 205-227. doi: http://dx.doi.org/10.1007/s11067-012-9178-y

Smith, J. M. (1982). *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.

Smith, J. M., & Price, G. R. (1973). The Logic of Animal Conflict. *Nature, 246*(5427), 15-18.

Sorrentino, F., & Mecholsky, N. (2011). Stability of strategies in payoff-driven evolutionary games on networks. *Chaos (Woodbury, N.Y.), 21*(3), 033110-033110-033110. doi: 10.1063/1.3613924

Steffensen, S., & Ulbrich, M. (2010). A NEW RELAXATION SCHEME FOR MATHEMATICAL PROGRAMS WITH EQUILIBRIUM CONSTRAINTS*. *SIAM Journal on Optimization, 20*(5), 2504-2539.

Transmission System. (2007). Retrieved from http://www.geni.org/globalenergy/library/national_energy_grid/ireland/Eirgrid-SONI-Map-A3-Oct2007.pdf

Vega-Redondo, F. (1996). *Evolution, Games and Economic Behaviour*. Oxford: Oxford University Press.

Wilensky, U. (1999). NetLogo (Version 4.1.3). Evanston, Illinois: Center for Connected Learning and Computer-Based Modeling, Northwestern University. Retrieved from http://ccl.northwestern.edu/netlogo/

Wood, K. (1993). Deterministic Network Interdiction. *Mathematical and Computer modeling, 17*(2), 1-18.