

UNIVERSITY OF CALIFORNIA

Santa Barbara

The Military Metaphor in Computer Network Defenses

A Thesis submitted in partial satisfaction of the
requirements for the degree Master of Arts
in Geography

by

Thomas James Pingel

Committee in charge:

Professor Sara Fabrikant, Chair

Professor Helen Couclelis

Professor Keith Clarke

Professor Giovanni Vigna

June 2004

The thesis of Thomas James Pingel is approved.

Helen Couclelis

Keith Clarke

Giovanni Vigna

Sara Fabrikant, Committee Chair

June 2004

The Military Metaphor in Computer Network Defenses

Copyright © 2004

by

Thomas James Pingel

ACKNOWLEDGEMENTS

It has been a long road to the completion of this thesis, and I couldn't have done it without the help of others. As a researcher and student, I've had the fortune of benefiting from the experience and ideas of those who have come before me. I'd like to thank my family- Greg, Denise, Adam, and Allison - for all of their love and support. My friends all across the country have consistently been a source of ideas, inspiration, and (not least) fun, which has really made the past few years happy and worthwhile. Thanks Kate, Meri, Mike, Chris, James, Yuki, Scott, Carlos, Tara, Tiger, Pickett, Erik, Jeff, Jeremy, Mark, and Dan.

The faculty and staff at the UCSB Geography Department are the best. My advisor, Sara Fabrikant has showed great patience, and has lent useful criticisms and comments throughout the project. Helen Couclelis, Keith Clarke, Giovanni Vigna, Dan Montello, and Jim Proctor have also given a great deal of time and help to my endeavors at UCSB.

ABSTRACT

The Military Metaphor in Computer Network Defenses

by

Thomas James Pingel

Abstract:

Currently, computer network defense borrows explicit language and concepts from physical security strategies. This thesis examines conceptual links between real-space and computer networks in order to provide justification for thinking of computer networks in traditional military terms. Many possible links are explored, and two factors – cover and concealment – are examined via two controlled experiments measuring firewall presence, number of services offered, and TCP port usage as independent variables and the number and intensity of computer network intrusion events as dependent variables. The empirical results are then compared to a classic study on the impacts of terrain on physical contests of force (Otterbein 1970) to determine whether cover and concealment behave in similar ways in real-space and in computer networks.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS.....	iv
ABSTRACT	v
TABLE OF CONTENTS	vi
LIST OF FIGURES	xi
LIST OF TABLES.....	xiii
1 INTRODUCTION	1
1.1 Problem Statement.....	1
1.2 Research Goals	3
1.3 Rationale of Approach.....	4
1.4 Research Questions.....	6
1.5 Relevance of Research.....	8
2 LITERATURE REVIEW	9
2.1 Introduction	9
2.2 Terrain	9
2.2.1 Terrain Defined.....	10
2.2.2 Terrain as a Metaphor.....	17
2.2.3 Elements of Terrain Analysis	22
2.3 The Defense of Geographic Space	27
2.3.1 Supplementation of Terrain Features	27
2.3.2 Size, Shape, and Material of Early Walls.....	28

2.3.3	Building Cities with Fortification in Mind	30
2.3.4	Evolution of Spatial Defense	31
2.3.5	Summary	33
2.4	Structure of Computer Networks	34
2.4.1	Relevant Communication Protocols	35
2.4.2	Topology of Computer Networks	39
2.5	Defense in Computer Networks	43
2.5.1	Attack in Computer Networks	43
2.5.2	Spatial vs. non-Spatial Elements of Defense	45
2.5.3	Firewalls	46
2.5.4	Bastion Hosts	48
2.5.5	Demilitarized Zones	49
2.5.6	Misleading Analogies	49
2.6	Metaphors of Spatiality Within Computer Networks	51
2.6.1	Distance	51
2.6.2	Density	52
2.6.3	Population Mobility	55
2.6.4	Visibility	57
2.6.5	Movement	59
2.6.6	Interaction	62
2.6.7	Summary	64

	2.7	Generalizing Spatial Elements Valuable to Computer Network	
		Defense	65
	2.7.1	Cover in Computer Networks.....	66
	2.7.2	Concealment in Computer Networks	69
	2.8	Summary.....	71
3		METHODS OF DATA COLLECTION AND ANALYSIS.....	73
	3.1	Introduction	73
	3.2	Testing Environment and Equipment.....	74
	3.2.1	Network Environment	74
	3.2.2	Hardware and Operating Systems.....	75
	3.2.3	Software.....	75
	3.3	Experiment One: Cover in Computer Networks	77
	3.3.1	Design.....	77
	3.3.2	Data Collection.....	83
	3.3.3	Method of Data Analysis	83
	3.4	Experiment Two: Concealment in Computer Networks	85
	3.4.1	Design.....	85
	3.4.2	Data Collection.....	91
	3.4.3	Method of Data Analysis	93
	3.5	Experiment Three: Cover and Fortification in Geographic Spaces.....	93
	3.5.1	Design.....	93
	3.5.2	Data Collection.....	95

	3.5.3	Method of Data Analysis	97
	3.6	Summary.....	98
4		RESULTS AND DISCUSSION.....	99
	4.1	Introduction	99
	4.2	Cover in Computer Networks.....	99
	4.2.1	General Results.....	99
	4.2.2	Attacks on Firewall by Phase	101
	4.2.3	Attacks on Clients by Phase	102
	4.2.4	Attacks on Firewall and Clients by Day	104
	4.2.5	Analysis and Summary of Cover Experiment	108
	4.3	Concealment in Computer Networks	109
	4.3.1	General Results.....	109
	4.3.2	Impact of Number of Services Offered.....	113
	4.3.3	Use of Concealment.....	115
	4.3.4	Quality of Concealment.....	119
	4.3.5	Searching Behavior.....	124
	4.3.6	Summary of Concealment Findings	126
	4.4	Cover and Fortification in Geographic Spaces.....	127
	4.5	Discussion.....	129
5		CONCLUSIONS AND OUTLOOK	132
	5.1	Conclusions	132
	5.2	Outlook.....	134

BIBLIOGRAPHY.....	136
APPENDIX	147

LIST OF FIGURES

Figure 2-A: TurboRisk Game Board	15
Figure 2-B: Topological Game Board of Risk	17
Figure 2-C: Physical Topology of a Network.....	42
Figure 2-D: Logical Topology of a Network.....	42
Figure 3-A: Logical Diagram of Cover Experiment at Outset	78
Figure 3-B: Logical Diagram of Cover Experiment at Conclusion.....	79
Figure 3-C: Testing for Cover - Expected Attacks on Firewall.....	81
Figure 3-D: Testing for Cover – Expected Attacks on Client	81
Figure 3-E: Expected Attacks on Server Running Services on Standard Ports (Server(S))	89
Figure 3-F: Expected Attacks on Server Running Services on Non-Standard Ports (Server(NS))	89
Figure 3-G: Logical Diagram of Concealment Experiment	92
Figure 4-A: Raw Attacks on Hosts by Day	105
Figure 4-B: Active Attacks on Hosts by Day	105
Figure 4-C: Linear Regression Plot of Raw Attacks by Day.....	106
Figure 4-D: Linear Regression Plot of Active Attacks by Day	106
Figure 4-E: Active Attacks by Day on Concealment Experiment Hosts.....	112
Figure 4-F: Packets Transmitted by Day on Concealment Experiment Hosts	112
Figure 4-G: Active Attacks by Phase on Concealment Experiment Hosts.....	117

Figure 4-H: Packets Transmitted by Phase on Concealment Experiment Hosts.....	117
Figure 4-I: Linear Regression Plot of Active Attacks by Phase on Concealment Experiment Hosts	118
Figure 4-J: Linear Regression Plot of Packets Transmitted by Phase on Concealment Experiment Hosts	119
Figure 4-K: Distribution of All Ports Contacted – Server(S).....	121
Figure 4-L: Distribution of All Ports Contacted - Server(NS)	121
Figure 4-M: Distribution of Ports Contacted (1-10,000) - Server(S).....	122
Figure 4-N: Distribution of Ports Contacted (1-10,000) - Server(NS).....	122
Figure 4-O: Distribution of Ports Contacted (1-2,000) - Server(S).....	123
Figure 4-P: Distribution of Ports Contacted (1-2,000) - Server(NS).....	123

LIST OF TABLES

Table 2-A: Metaphors of Spatiality Within Computer Networks	65
Table 3-A: Independent and Dependent Variables for Cover Experiment and the Expected Relationships Between Them	82
Table 3-B: Progression of Concealment Experiment	87
Table 3-C: Independent and Dependent Variables for Concealment Experiment and the Expected Relationship Between Them	90
Table 3-D: Independent and Dependent Variables Selected From Otterbein's (1970) Factors	94
Table 3-E: Recoding Of Otterbein's Variables	97
Table 4-A: Raw and Active Attacks on Firewall by Phase	102
Table 4-B: Raw and Active Attacks on Clients by Phase	104
Table 4-C: Regression and Correlation results for Cover Experiment Hosts.....	108
Table 4-D: Active Attack Rates and Correlations	114
Table 4-E: Attack and Data Traffic Data for Concealment Experiment	115
Table 4-F: Packets Transmitted To/From Ports Used in Concealment Experiment .	126
Table 4-G: Other Measures of Searching Behavior on Concealment Experiment Hosts	126
Table 4-H: Chi-squared Test P-Values for Cover and Fortification in Otterbein's (1970) Data.....	128
Table 4-I: Contingency Table for Cover vs Frequency of Attack	128

1 INTRODUCTION

1.1 Problem Statement

Computer networks are increasingly targets of malicious intent (CERT Coordination Center 2003). As the number of computer and computer network targeted attacks continues to rise, administrators and researchers are challenged to improve network security. Not only does the sheer number of attacks beg attention to the issue of computer security, but so also does the potential cost of a successful attack on the information infrastructure (President's Critical Infrastructure Protection Board 2002). The problem of network security is not new, but it becomes increasingly complex over time, as services offered by computers and networks increase.

The bulk of the work being done to solve computer security problems comes from the field of computer science as well as industry. This is evidenced by the only relatively recent entry of the United States Government into the cyber-security realm (Presidential Decision Directive 63 1998; Executive Order 13231 1998; Critical Infrastructure Assurance Office 2003). In turn, much of these solutions are engineering-related: new and better programs with fewer security bugs, patches to repair problems with existing software, and improved theoretical design with security firmly in mind are all quite common approaches to solving the computer security dilemma.

These technology and engineering approaches are invaluable toward easing the security problem. Yet it may be that these approaches alone neglect one of the most fundamental aspects of the security problem – namely, that for as long as computer networks are spatial entities, they will be subject to at least some of the same constraints as govern the “real” spaces like the ones humans inhabit in their everyday lives. Based on this assumption, it is appropriate to look to security problems in real-space, and see how they compare to security problems in cyber-space.

This is a study of a relation between methods and strategies to secure physical space (i.e., the fortification of territory) and similar methods and strategies to secure computers and computer networks. Specifically, I attempted to explain how two variables widely recognized within physical security – cover and concealment – manifest within the topological space of computer networks. Elements of manifestation include the characteristics of cyber-spatial features that could be classified as cover or concealment, but also an explanation of how these features compare to their physical counterparts in terms of how they affect deterrence of attack, and how well they protect a network or host from attack.

In order to provide a satisfactory explanation of the manifestation of cover and concealment within computer networks, it is necessary to first explore the broader concept of terrain in military security. The concept of terrain is presented in a highly abstracted way that enables a compression of fully three-dimensional space into a network space. While this approach does not allow for full application of a terrain metaphor, it does enable application of the military security concepts generally

discussed in connection with terrain. In effect, only network related topologies found in real space are used to describe the spatial defenses of computer networks.

Cover and concealment are presented as two metaphors with which to understand the spatial elements of computer network defenses. Instances of cover and concealment are then tested on real networks to determine how well functional characteristics apply to both military security in real space and to computer networks.

1.2 Research Goals

The objective of this thesis is to provide groundwork for a comparison of spatial elements common between the two very separate worlds of computer network security and military geography. Although computer network security techniques, technologies, and language all to some extent borrow from concepts of physical security, a rigorous exposition of their shared spatial commonalities may be of some use in each of the communities. As governments increasingly attempt to project their power into cyberspace and continue to maintain an overwhelming military presence in real spaces, a synthesized concept of security that overlaps both elements is crucial.

This thesis is about researching the justification for applying the metaphor of military security to computer networks. This is done by carefully examining the literature to find possible reasons for links between spatial strategies in geographic spaces and spatial strategies used to help defend computer networks. The goal is to

expose reasons why thinking of computer networks in traditional military terms is appropriate.

1.3 Rationale of Approach

In real space, various elements of landscapes, whether created by humans or not, contribute to or detract from the application of force. Examples of these include land cover (plains, deserts, mountains), barriers (rivers, mountains, constructed walls), and transportation routes (rivers, roads, railways). Landforms that reduce or negate an offensive force are defensively significant landforms.

There are many cases in which large-scale objects are created by humans to aid in defense (e.g., moats, walls, ditches). Additionally, the skill of a commander rests largely on his or her ability to master the spatial configurations of landforms and armies. The choice of an advantageous battlefield can have great repercussions on the outcome of any contest of force. All of this is well known in real space, and has been documented by generals, geographers, and historians for thousands of years (Peltier and Percy 1966; Clausewitz 1968; Keegan 1999).

Geographers O’Sullivan and Miller (1983, p. 7) have asserted that, “the fundamental strategic and tactical problems are geographical in nature.” Certainly this seems true in real space, but what about artificially created spaces like chessboards? Those with the same view as O’Sullivan and Miller have not made a mistake in their assertion, because strategy in games such as chess and *go* still depends on depth of understanding of spatial configuration (Atherton, Zhuang et al. 2003; Chen, Zhang et al. 2003). Still, games like chess and *go* are explicitly spatial.

Do geographic concerns also affect tactical and strategic problems that exist in more abstract spaces?

The abstraction of space in games like chess and *go* suggest that they do. Although these games are often performed in real space, the real game playing takes place in the mind. The manifestations of the games are usually physical, but they need not be. If the fundamental strategic and tactical problems (in general) are geographical in nature, then making the analytical jump from physical contests of force (e.g., traditional warfare) to the contests of force in the pseudo-space of computer networks should be possible.

Computer networks are spatial simply if only because they exist in our physical world. Computer workstations, servers, and the physical connections between them (wire, fiber optic, or wireless) exist, and exist in specific locations. More significant than just their locations, though, is that they are connected by lines, switches, and protocols to form networks. The connection and arrangement within a network forms a topology.

The topology of computer networks creates a “space” that is artificial in much the same way that a “space” is created during the play of a game such as chess or *go*. The rules that govern network data traffic are not unlike those that govern game play. Not every type of movement is possible. Movement is limited by the rule structure. Strategy and tactics in these artificial spaces (and real spaces for that matter) begins with an understanding of the rules, limitations, and capabilities of pieces within the sphere of the space created by the game.

Computer networks are now interlinked to the point that together they form the Internet. Although very large and quite dynamic, the Internet itself has its own topology. As people and data interact on the Internet, a virtual space is created. Since computer networks have their own topology and rules, contests of force within them could be thought of in much the same way as a game. Given that people, groups, and states have interests that are manifested in cyberspace, it is no wonder that conflicting interests have manifested themselves as contests of force. These contests of interests with respect to computers and networks are studied as Information Security. Information Security addresses the issues of information (data) confidentiality, integrity, availability, and privacy (Kemmerer 2002).

If we have in computer networks a spatial configuration (topology) and a contest of interest (leading to a contest of force) we may wonder whether, in such a complex structure such as the Internet, local variability in spatial arrangement may yield offensive/defensive advantages in the same way as does similar variability in real landscapes.

1.4 Research Questions

To determine how the military metaphor might apply to computer network defense, the general research questions must be refined into specific research questions. These are:

What are the network-based topologies that humans have used as part of their defenses in military affairs?

How have computer network defenses been organized spatially?

How do specific uses of cover and concealment compare between real space and computer networks?

The problem is large, but it is one that can be attacked in stages. The first step is to investigate the current literature on terrain and fortification and their effect on physical security. Next, it is important to establish how computer networks actually function. With these two pieces, it is then possible to attempt to formulate a metaphorical mapping between physical space and computer networks by using network-related elements inherent in the concepts of terrain and fortification.

Apart from the theoretical framework developed from the relevant literature, an experiment was conducted to test whether two aspects of spatial defense, cover and concealment, behave the same way in computer networks as they do in real space. This is accomplished by first looking for the most likely candidates for examples of cover and concealment, and then subjecting each to a controlled experiment that measures changes in rates of attack based on the presence or absence of either cover or concealment. The results of these two experiments, when compared to data collected on the impact of terrain on traditional military operations, provide a benchmark for comparison between cyber-security and physical security within the context of cover and concealment.

1.5 Relevance of Research

Internet mapping projects are widespread, ranging from work done commercially to that done for academic and research interest (Cheswick and Burch 2001; Dodge and Kitchin 2001). Many of these mapping projects are content or semantically oriented (e.g., showing how web pages refer to other web pages), but some tools exist already to display useful, small-scale information about particular networks. Of the most useful of these are produced by network component manufacturers such as Cisco and 3Com. None of these applications are specifically designed with spatial security in mind, although clever systems administrators could possibly use them in that way. The lack of such tools was cited by the Advanced Network Defense Research workshop in 2000 as a critical piece of the network security puzzle that could significantly improve defensibility of computer networks (Anderson, Brackney et al. 2000).

Although development of such tools could be done independently and without consultation to cartographic representation techniques and widespread research in the disciplines of geography and military science, to do so would be unwise. The reason for this is that a spatial defense strategy of a particular place, time, and medium is closely related to other spatial defense strategies by virtue of their similar spatial configuration.

2 LITERATURE REVIEW

2.1 *Introduction*

The first step in unraveling the spatial links between physical and cyber-security is an integrative and interdisciplinary analysis of the current literature. Because of the relatively segregated nature of the literature, the review and analysis are initially treated separately here. The more traditional physical security is treated first, with computer networks and network security reviewed second. With both of these bodies of literature summarized, the theoretical development of the thesis proceeds with an examination of the spatial metaphors that are applicable to a discussion of network security. Finally, these spatial metaphors are used to provide a conceptual framework of cover and concealment within computer network security. Ultimately, these two concepts provide an opportunity for an empirical evaluation of the appropriateness of the military metaphor in computer network security in the following chapter.

2.2 *Terrain*

Before proceeding with an in-depth analysis of the physical security literature, some groundwork should be provided in the form of an explanation of the concept of “terrain.” The concept of terrain is important as a starting point to evaluating the military metaphor, since terrain (in its fully metric sense) is the basis for military

spatial analysis. As such, it is worth some time to clarify both the meaning of the term and the reasons for its inclusion as an element of analysis here.

2.2.1 Terrain Defined

Terrain is defined as, “A tract of country considered with regard to its natural features, configuration, etc.; in military use especially as affecting its tactical advantages, fitness for maneuvering, etc”(Unabridged Oxford English Dictionary 2002). The use of “country” implies land to the exclusion of deep-water ocean (Collins 1998). Country is then not only the land surface, but also rivers, lakes, coastline, and the like. Further, the land *surface* is not the only relevant component of terrain. If this were the case, terrain would only be relevant in as much as it expressed relief – the relative heights of the surfaces within a certain area.

It is interesting to note that there are no references to “the Earth’s surface” in this definition. “Surface” may well have been excluded for reasons mentioned above, while references to Earth may have been rejected since we now have knowledge of terrains on other planets and objects. In this sense, the historical ties to the earth connoted by the root *terra* are abandoned in favor of the looser term “land”. Indeed in other versions of the Oxford English Dictionary as well as other dictionaries, “land” is used in place of “country” (Oxford University Press. 2002).

Collins (1998) argues that terrain is “all physical and cultural geographical features within any given area.” However, this may be casting the net of definition too widely. While all physical and cultural features of an area are doubtlessly of some relevance to anyone engaging in conflict in that area, not all of these are

necessarily of relevance to terrain. To illustrate this small point, one may review the scope of United States military terrain analysts (United States. Dept. of the Army 1990). The scope of the terrain analyst's work is restricted to the gathering and reporting of "terrain intelligence", a subset of "operational intelligence". Terrain intelligence is thus the product of the analysis of terrain, which for the U.S. Army is restricted to "a portion of the earth's surface that includes man-made and natural features... and the influence of weather and climate on them" (*ibid.*, p. 531). Further, one engaging in terrain intelligence gathering (at least for the U.S. Department of Defense) is not interested in the terrain *per se*, but rather the military significance of that terrain (United States. Department of Defense 2003). Collins's definition corresponds more closely to the U.S. military's more general view of intelligence, the function of which is to gather information concerning foreign countries or areas (*ibid.*).

So, following the U.S. Army operational definition of terrain, one may reasonably restrict the scope of "terrain" somewhat further than the Collins definition. However the U.S. Department of Defense (DOD) definition is also slightly too narrow. It focuses primarily on the military significance of natural and human-made characteristics of an area (United States. Dept. of the Army 1990). While this is perfectly reasonable for a practical land-based army, a conceptual definition of terrain, even when restricted to military discourse, needs to be widened. The reason that this is necessary is that the usage of terrain in a military context could conceivably refer to terrains that are not surface based, as the usage of "area"

would imply. For instance, while naval combat is often done in what would seem to be the relatively smooth, uniform surface of the sea, it is still reasonable to refer to the under water surface irregularity. The combat of naval ships occurs within two surfaces, both of which are important – the surface of the water and the surface of the ocean floor. Further, underground variability may matter – for instance cave structures where opposition may find cover. Finally, as noted earlier, terrain analysis is more than just the study of surface variability (i.e., relief). Rather, military terrain analysts give concern to vegetation, water distribution, and soil types in an area of interest (United States. Dept. of the Army 1990). The result of this concern, especially when one also includes weather and climate impacts, changes one's concern from an *area* of interest to a *volume* of interest.

The important realization is that terrain is not merely the configuration of a *tract of country* or an *area*, but the configuration of whatever the medium of conflict happens to be. For historically land-bound humans on Earth, the components of terrain remain generally two-dimensional. In military geography, terrain has been expanded to include a vertical component as well, and thus a third dimension. As operations in the air and underwater became more widespread, the vertical component became increasingly important.

What is at issue here is not whether two or three dimensions are required to discuss terrain in a military context, but that the importance of the third dimension has largely been due to *the location of the conflict*. Similarly whether humans are moving on or through land, air, water, or some other substance is largely irrelevant.

As humans have changed how they move, the descriptions and inclusions of terrain have changed, and thus, so has the definition of the concept. As such, the analysis of terrain can be applied to various zones of conflict, including urban areas, underground tunnels, and outer space.

Before the question of the application of the term “terrain” to computer networks can be addressed, there remains one more fundamental spatial element of terrain other than its dimensionality that needs to be addressed. The importance of topological properties of a space and how these influence the movement of combatants is of vital importance. This is hinted at by the term “configuration” used in the first definition provided above. Although the metric properties of terrain are of course important (demonstrated by the precision to which military organizations have tried to document terrain), their relative positions to each other are just as important. The unpleasant terrain between a commander’s position and his destination is important precisely because it is *between* – not just because there are some given number of miles of it.

This topological concern with terrain comes in large part from its function – both in the context of warfare and in other analyses of it. The definitions provided within the academic discourse on military geography include this functional concern with terrain automatically, but it is explicitly noted in the definition found in the Oxford English Dictionary. Penguin’s “Dictionary of Geography” defines terrain as, “an area of land in respect to its physical characteristics or condition, *especially if*

considered for its fitness or use for a special purpose, e.g. for laying a railway track, or for a military operation,” (Clark 1998).

It is this functional aspect of terrain that provides clues as to why a more abstract, non-geocentric definition of terrain (restricted to military discourse) should be *the configuration of the medium of conflict*. Although in geographic spaces, this typically occurs on land, it may not – as in the case of submarine warfare. Although typically thought of as an area, its application is not restricted to two dimensions. Finally, although it is usually thought of as applying to metric spaces, it includes topological elements as well, even if the topological properties arise out of Euclidean metric underpinnings. All of these matter because terrain is fundamentally about the effects of one’s medium on one’s movement in it.

With this conceptual definition of terrain in hand we can sketch an argument for what elements of *terrain* rightly describe computer networks, and thus why traditional techniques of terrain analysis may help provide valuable insights to how defenses of computer networks might be better organized. Computer networks, as shall be described, have both metric and topological properties. Although representations of computer networks (both in the mind and on paper) frequently abandon or reduce the importance of these metric properties, this is not to say that they are unimportant – they are simply inessential to certain types of understanding. It is possible to reduce the importance of metric properties to consider only how the topological properties affect military operations. An example may help to provide illustration of the concept.

Hasbro's Risk™ is a popular parlor game where the goal of the game is to conquer the world by occupying geographically-related entities (usually referred to informally as “countries” or “territories”) with “armies”. Territories are connected to each other across borders, and movement only occurs between adjacent countries. The game board is essentially a two dimensional map of the real earth's surface, excluding Antarctica and generally displayed in a Mercator-like map projection. The territories are grouped into “continents” such that rewards go to players who control all of the constituent territories of a continent. These continents correspond exactly to the six major human-bearing continents on the real Earth. There are numerous copies of the game in electronic format, one of which is TurboRisk (Ferrari 2003). Figure 2-A (below) shows the opening game board.

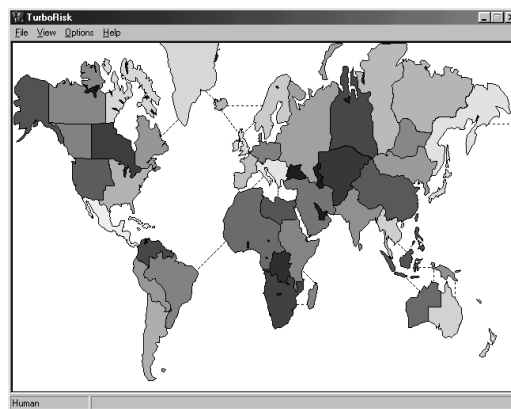


Figure 2-A: TurboRisk Game Board

One may notice that the continuities and discontinuities in the land surface are all roughly preserved such that – as a whole – both the physical geography of the land and the human geography of nations and territories are recognizable. Further, rules of movement for land-based armies are generally preserved as well – armies may

only move into territories directly adjacent to them. There are, however, only two distinctions between surfaces – land and sea. Movement over sea is only permitted between certain territories (e.g., between Brazil and Western Africa). Land surfaces are treated as homogeneous – there are no mountains, rivers, etc.

The game surface itself is two-dimensional. It is both metric and explicitly network-topological, and there is a strong relationship between the two – namely, that the nearer the territories, the more likely they are to be “connected”. But if the metric space is discarded – as it is for computer players – the result is a topology that retains all of the strategic elements of the fully metric space. The same geography that produces the choke points and long indefensible borders in the real-world ultimately produces similar spaces within the game of Risk. Figure 4-B shows the abstracted topological space generated from the Risk game board with some geographic names provided for visual reference. The circles and links that represent the territories and connections remain in their original position and orientation for ease of identification between Figures 2-A and 2-B. Orientation and metric distance could be discarded, as long as the relative positions remain the same.




Figure 2-B: Topological Game Board of Risk

If the effects of terrain could be seen in geographically based generation of territories, the contrast of land and sea spaces, and the aggregation of territories into continents then the same should be visible within the abstracted topological space. At least some of the strategic elements arising from the space of the game board are probably more visible. It is certainly true that the game board of Risk does not encompass all of the aspects of terrain – that would be impossible in almost any context, let alone in a parlor game. Yet if any remain at all then the argument for at least *looking for* some elements of terrain within computer networks seems justified since many of the features of this game are also represented within computer networks (i.e., movement, segmentation of space, connectedness of entities, path/node structure).

2.2.2 Terrain as a Metaphor

The previous arguments for application of some elements of the concept of terrain to computer network defense are suggestive, but not conclusive. In order to

make the case more strongly, it is important to return to the question of how *thinking of* computer networks in terms of traditional military spatial defense might be beneficial to the process of improving network security.

As Lakoff and Johnson (1980) have pointed out, geographic experiences underlie much of our thinking. People often use geographic terms like “landscape” and “terrain” to apply to things that do not exist in space. “Landscape,” for example, is sometimes used to describe purely mental phenomena, as in the term *conceptual landscape*. In this case a field of possible concepts or ideas is organized in some way. The evocation of landscape implies a cohesive spatial layout. The ideas do not literally have spatial relationships with one another, but such relationships can be imagined.

A spatial relationship between terms can even be formalized to develop a “space” on which conceptual elements may be placed, as in the case with multi-dimensional scaling techniques. Much of this kind of quantification and its subsequent visualization involves spatializing non-spatial information. Couclelis (1998) clarifies “true” or geographic spatializations as those which “reproduce aspects of the kinds of spaces that are familiar to people from everyday experience: those of working areas (e.g., desktops), of rooms, buildings or of larger geographic-scale spaces.” The key to these kinds of spatializations comes from interaction with them: users must be able to interact with representations in the same way as they would interact with the real thing. It is in this way that spatializations are considered metaphors – they establish *prima facie* commonalities from which other, less

obvious relationships also hold true. The power of the metaphor comes directly from the innumerable, subtle, but conceptually important ways in which one thing resembles another. It is in this way that metaphors help us understand one thing by virtue of another, and also lead us into making fallacious conceptual leaps.

The spatial metaphor has been applied to computer networks to such an extent that it is near impossible to discuss them without appealing to spatial concepts. The term “network” itself acquires much of its cognitive content from a physical net. One of the more culturally prevalent terms, “cyberspace,” explicitly invokes space as a means toward understanding the digital interlinked world of information technologies. If what Lakoff and Johnson (1980) say is correct, the spatial metaphor so pervades everyday speech that untangling the morass of figurative language to get back to some literal (i.e., non-metaphorical) truth about how computer networks function is completely impossible. The pervasiveness of metaphor in language and in thought would undermine such an endeavor before it could even begin. Whether or not computer networks are spatial, people understand them in spatial terms by virtue of the metaphors they have employed to name, describe, and build them.

Part of the justification for viewing information as terrain comes from this linguistic and conceptual borrowing of spatial terminology. Similarly, a computer network could be thought of as a metaphorical terrain, simply by virtue of the many spatial terms. Unfortunately, as is described in detail below, the borrowed terms often do not well-describe their counterparts in the real world. As a result, a simple explicit mapping of the source and target domains would not result in a spatialization

that would satisfy the admonition that interaction is the same in the spatialization as it would be in the real versions. Although important to note in its particularities, this should not be surprising. Metaphors always abandon some differences to emphasize the commonalities; they do not claim that *all* qualities of one thing apply to another.

The problem with using such a simple mapping of computer network security elements to their geographic, real-world counterparts is that one can neither use the mapping at face value to think about computer security (as might be done by a relatively naïve user community) nor completely abandon it and study computer security completely independently of wider geographic and spatial security problems. The solution, as I see it, is to a) explore the relevant differences in terminology and techniques as they exist, and b) try to outline an informed correspondence between physical and network security that includes elements of traditional terrain analysis. What makes this possible, in part, is the rich metaphorical adaptation of basic spatial concepts into abstract categories.

Applying metaphors of military spatial defenses to computer network defense is different in some important ways to many spatializations. First, there are explicit spatial relationships already built-in to the computer network. As is described later, a great deal of security problems arise from these spatial relationships. This is quite different from many information visualizations (Fabrikant and Buttenfield 2001) that rely on coding similarity-as-distance. This kind of spatialization is attempted on the rationale that since (following Tobler's First Law of Geography) nearer things are more related than distant things, people will similarly *think of* nearer things as more

similar. Attempts to reveal this “First Law of Cognitive Geography” have been successful in showing that people do tend to think of closer objects as more similar by default (Montello, Fabrikant et al. 2003).

Fisk, Smith, *et al.* (2003) have developed a spatialization of computer network security based on the metaphor of territory. This spatialization was developed with Lakoff and Johnson’s theory of metaphor firmly in mind. A virtual territory is created based on the arrangement of the internal network, and attacks on the network are visualized based on information about the attack. The achievement of the visualization is that it graphically shows the difference between the self (the internal network) and the other (external networks) and projects attack data onto this schema.

This visualization represents a giant leap forward in the visualization of network security data by allowing for rapid discernment of different attack patterns. One drawback is that the system uses a minimum of geographic concepts. The visualization, as a result, does not produce a fully “real” environment that can be interacted with. To some extent this is understandable – any metaphorical mapping has similarities and differences between the source and target domain. Failure to produce a “fully geographic” terrain is not a failure of the visualization. One of the aims of any visualization is to compress a great deal of data into an intuitive form. Use of fully-fleshed geographic terrain may allow for better compression of information, and better performance of the system as a whole, provided that the mapping between the representation of geographic-object and the source phenomenon is based on some principle that can be intuitively grasped by the user.

One such principle may be the functional similarities between some technologies and techniques used in computer network defense and corresponding categories of terrain used in military geography.

The next section highlights some of these traditional elements of terrain analysis, with the goal of providing conceptual and metaphorical categories into which computer network security elements can later be grouped. Two of these three will ultimately serve as a basis for empirical evaluation of the applicability of the military spatial defense metaphor in computer networks.

2.2.3 Elements of Terrain Analysis

The specific methods of utilizing different terrain to one's benefit in conflict arise from the simple idea that terrain influences the outcome of conflict. In order to explain the specific examples (the *how* of terrain influence) one must first have (in varying degrees of explicitness) a knowledge of *why* terrain matters. Fundamentally, terrain seems to matter in as much as it limits or enhances (a) movement (of combatants, weapons, and ammunition) and (b) observable knowledge about combatants.

Evidence of this claim can be found in the five terrain factors identified by the DOD in their literature. The five basic categories of interest to the terrain analyst are: (a) Avenues of Approach, (b) Points of Observation and Fields of Fire, (c) Key Terrain, (d) Cover and Concealment, and (e) Obstacles (United States. Dept. of the Army 1986). *Avenues of Approach* and *Observation* are clearly related to movement (in the case of the former) and acquisition of knowledge (in the case of the latter).

Fields of Fire are areas or points that offer strong offensive positions. They are places where fire can be directed in a way such that the enemy cannot avoid it, and cannot return fire (O'Sullivan 1991). This has the effect of limiting movement – either by the actual destruction of adversaries that attempt to cross a field of fire, or simply the lack of movement due to the anticipated result of such a crossing (i.e., the deterrent effect).

Key Terrain (also known as *Critical Terrain*) generally refers to any feature that offers particular advantage to one who controls it (United States. Dept. of the Army 1986; Collins 1998). In some cases, battlefields offer no such manifestly extraordinary features. Further, what may be considered key terrain for one set of goals may not be considered crucial for another set of goals, even within the same broad area under consideration. In some cases, one feature may be key terrain for one side of a conflict but highly disadvantageous for another side, due to training differences, equipment, and so on (O'Sullivan 1991). Key terrain, then, may be any of the other four types of features outlined above, but rather than being merely one consideration among many, is of such paramount importance that it thereby greatly reduces the relative importance of other factors.

The remaining two categories of terrain identified by the DOD actually contain three separate elements: Cover, Obstacles, and Concealment. Clausewitz, in his treatment of terrain in the classic *On War*, reduced the impact of terrain on military operations to these since “all other properties can be traced back to these three” (Clausewitz 1984). As a result, it is important to give these categories a somewhat

more in-depth examination. At the end of this chapter, we return to two of these three elements to discuss how they apply to computer network security.

2.2.3.1 Cover

The U.S. Army Operations Manual defines cover as protection from observation and fire (United States. Dept. of the Army 1986). The size of the covering object is related to the size of the object that needs protecting. For example, a large tree or automobile may provide cover to a single human, but not to a tank. The effectiveness of cover is related to the type of offensive weapon being employed. The same automobile in the previous example may provide cover for small weapons fire, such as a pistol, but not against small munitions such as a grenade. Cover generally inhibits the ability of an attacking weapon or munition to reach a defender. Its purpose is to interdict the munition or weapon before it reaches its target, not to interdict the attacker's motion.

One feature of cover is that so long as it is measured during threatening circumstances one can observe a displacement of fire from the target object to the covering object. To use a military example, if a soldier hides behind a car, and someone is shooting at him, the car will be shot and not the soldier – at least to the extent that the car is successful as cover. Based on this principle of displaced attack, we can test whether an object behaves as cover if we observe this displacement phenomenon.

Objects behaving as cover, in conventional military terms, are spatial since they are real physical objects that occupy locations, and have sizes and shapes. Their behavior as cover is also spatial from a more abstract point of view. The behavior of an object as cover can be represented topologically as *in-between* two other objects. For instance, consider opponents A and B engaged in some kind of struggle. If player A is using some object C as cover, that object is useful only in as much as it exists *in-between* players A and B. Its in-between-ness can be represented in a spatial, but not metric way (e.g., topologically).

Cover can be used to describe non-spatial phenomena (e.g., a reporter may cover a story), but these usages are metaphorical adaptations of the more familiar spatial use of the concept. The more everyday usages of cover (e.g., to cover a table with a tablecloth) not only utilize the notion of *between* (e.g., the tablecloth is between the table and the diner) but also typically utilize a notion of vertical directionality. To cover something is typically to place something on top of another object. It is interesting that cover, when used in the military context, sheds the link to vertical directionality, and thus is a wider sense of the term.

Cover-fire is a counter-offensive use of weapons fire to provide protection from attack. In this case, weapons fire is trained on the attacker, who must then mind his or her own defenses, thus reducing his or her offensive threat. The rear-guards of armies, whose purpose is to fight and delay a pursuing army, perform the same function of providing cover.

In this sense it can be seen that the emphasis of cover is really on protecting from fire rather than observation. The close link in the literature is probably due to both a) the primary usage of cover as an object of passive defense, and b) the natural opacity of any object typically used for cover in a military setting.

2.2.3.2 Obstacles

Terrain features that inhibit movement are called obstacles (United States. Dept. of the Army 1986; United States. Dept. of the Army 1990). As with cover, the constitution of an obstacle is relative to type of offense employed and the object or objects to be defended. Impediments to movement are highly sensitive to the abilities of the moving object. Terrain obstacles may be natural or human-made. Swamps, rivers, and cliffs are examples of natural terrain features that may slow or inhibit movement. Human-made examples of obstacles include walls, trenches, and barricades.

2.2.3.3 Concealment

Concealment, like cover, protects objects from observation. Unlike cover, however, concealment offers no protection (Collins 1998). A classic example is the camouflage of a soldier's jacket or the green paint on a tank. Neither offers physical protection from attack, but both offer some limitation to an opponent's perception. In terms of terrain, an example of concealment might be foliage that prevents visual identification but will hardly stop bullets. Concealment is highly sensitive to the perceptual aids and abilities, sizes, and number of combatants. Increasingly, tactical

concealment has become more difficult as remote sensing and alternative detection technologies have become more widespread (United States. Dept. of the Army 1986).

2.3 The Defense of Geographic Space

The previous section described some of the important conceptual aspects of terrain of interest to military geographers. Many of these focus on natural features rather than on human augmentation to the existing environment. This is due, in part, to the abandonment of fixed fortification in favor of a defense that relies on mobility. Despite this relatively recent change in strategy, it is important to devote some attention to the history of fortification, because this affords both a greater depth of knowledge of possible defensive techniques while at the same time building on the concept of terrain, as it incorporates both natural and human-made spatial structures in defense.

The study of fortification is also important because, as will be shown in later sections, fortification is the pre-eminent model for network defense today. The details provided in this section are intended to provide a fuller understanding of the nature of the commonalities present between defenses of physical and network spaces.

2.3.1 Supplementation of Terrain Features

In many cases, the natural terrain offers insufficient protection to settlements. As a result, humans have tended to fortify existing settlements. The earliest recorded

defensive walls for settlements can be found in Jericho, although it is more than probable that earlier villages were also fortified in some manner (Yadin 1963; Keeley 1996). Although the adding of walls and ditches for defense has been a dominant theme throughout the history of fortification, the best designs were those that merged terrain features with human-made structures. Examples of such fortification plans include the Greek city of Knidos, which used reinforcement of nearby ridges in its defense, as well as the sites chosen by the Etruscans, who pragmatically fortified plateaus and hilltops (De la Croix 1972). Indeed the very idea of the *acropolis* (directly translated as *high-city*) is one that merges defensive augmentation with the natural terrain advantage of higher ground.

2.3.2 Size, Shape, and Material of Early Walls

Walls occupy a pre-eminent place in fortification strategy, and it is difficult (though possible) to imagine fortifications without walls. Topologically, walls exist in between two sets of objects. In the case of walls of fortification, they define by their existence an inside and an outside. In most cases, walls also have doors or portals, connections between the inside and outside. Through their existence, they transform a relatively open area of space, where many paths are available, to a restricted space of movement, where many fewer paths are available. With this in mind, we can briefly trace some of the history of walls in order to more fully understand how walls have evolved over time, and the reasons for their general abandonment in military affairs. This more specific knowledge should help provide

for a more detailed comparison to defensive practices used in computer networks later in the thesis.

Most walls from pre-history were likely made of earth and wood, and as such few survive today (Keeley 1996). The earliest stone fortifications are found at Jericho, date to roughly 7500 B.C., and are preserved to the impressive height of almost four meters (De la Croix 1972). The composition of walls varied with resources and technology available, but typically included earth, wood, brick, stone (Collins 1998). Early walls had to resist attacks via sapping, mining, battering ram, and escalade. By the fourth century B.C., range attacks were brought to bear by means of catapults and, later, ballistas (Winter 1971; Payne-Gallwey 1995). These offensive methods could be overcome by building walls thicker and higher, but without paying a great deal of attention to shape. As a result, outer walls were often either irregular (following terrain) or else regular, but somewhat simple in shape – generally rectangular or ovaloid (De la Croix 1972).

Some specific examples may help to illustrate the enormous proportions of early walls. Evidence suggests that the walls of Babylon and Ur in ancient Mesopotamia were twenty-five meters high, and over twenty-five meters thick, respectively (*ibid.*). Winter (1971) estimates typical early Greek walls during the Iron Age at a height of 3.5-4.5 meters, with a width of 1.75 meters, while walls after the fifth century were typically over 6 meters high, with those from the Hellenistic period an impressive 10-12 meters high.

Early frontier Roman fortification differed from Greek fortification in that materials were more often timber and earth than brick and stone, with similar widths and heights (Johnson 1983). Later, the Romans increasingly used symmetrical stone-based walls, again with similar widths and heights (*ibid.*). The Aurelian wall protecting Rome itself was an exception to the general rule of relatively weaker but highly planned fortification, being an irregular 19 kilometers long and reinforced to over 12 meters in height (De la Croix 1972; Leigheb 1998).

2.3.3 Building Cities with Fortification in Mind

Many times the establishment of a fortified wall or palisade around a settlement was added after its establishment. The likelihood and extent of fortifications varied based on perceived need of security, economic considerations, and political organization (Keeley 1996). In some cases, however, especially in frontier regions where risk of attack was high, settlements were designed with fortification in mind at the outset. Examples of this can be found throughout human history, from prehistoric settlements to Roman outposts in antiquity, to cities built during the Middle Ages and Renaissance (De la Croix 1972).

As warfare became increasingly advanced, building cities with a security plan involved much more than just building walls. As previously mentioned, the choice of the city's site was made with an attention to the advantages that the terrain afforded. City plans incorporated an improved road structure that allowed for the easy repositioning and deployment of troops during a siege (*ibid.*). Towers were built to allow for more advanced warning of an impending attack (Keeley 1996).

Thus, planners increasingly realized the importance of cover (in the form of walls), but also of mobility and observation, even in the defense of a fixed position.

2.3.4 Evolution of Spatial Defense

From antiquity until the Renaissance, fortification brooked no significant advances. With the development of the cannon in the fifteenth century, however, fortification patterns changed significantly (Howard 1976). The cannon was capable of reducing even the strongest of stone walls to rubble, so fortification methods were adapted by modifying the material and arrangement of defenses. Earthworks were resurrected as a means to bolster walls by absorbing the impact of the cannon fire. At the same time, walls were arranged in more complex shapes. Increasingly angular (and sometimes circular) walls were adapted so that the force of impact was reduced (De la Croix 1972).

It was also during this period that *bastions* developed. Bastions were fortified projections from walls that served to cover avenues of approach with interlocking fields of fire, such that any point could be engaged by at least two bastions (De la Croix 1972; Oxford University Press. 2002). While the term *bastion* has come to mean any particularly well-fortified structure, or even something akin to “a final holdout” as in the expression “the last bastion”, the term originally referred to a counter-offensive structure within the defense. The development occurred as cities increasingly began to realize that the day of the passive defense had ended, and a successful resistance depended on a way to demolish an attacker, not simply to hold out long enough for reinforcements.

As fixed fortifications became costly in terms of constriction of commercial growth and less effective as a defense, a mobile counter-offensive strategy of defense began to take hold. Rather than relying on passive defense, states increasingly used artillery itself as a defense rather than walls or other obstacles. In this case the function of obstacles as barriers to movement was superseded by offensive weapons, particularly artillery. Rather than relying on a wall or ditch (or natural terrain feature) an army would limit enemy troop movement either by the deterrent factor of an artillery storm, or if that failed, by preventing enemy troops from moving anywhere – by killing them. Use of offensive power in this way was not *new*. Rather, it simply became a much more powerful and reliable method than it had in the past while at the same time passive defense became less and less reliable. A good offense had become the best defense.

The technological superiority of defense – despite the uselessness of walls – comes from a change in how defense was organized. Again, rather than relying on passive defense solely, defenses included components that exacted a cost (of lives or at least serious injury) for movement. An example of this (other than artillery as defense) is employment of barbed-wire (and later concertina, or razor wire) in the defense. Rather than trying to completely inhibit movement, it is counter-offensive as it seeks to cause injury to the transgressor. Even more prescient now are the land and sea mines that limit movement by injury, to be sure, but also by the tremendous fear of injury that they *threaten*.

With the advent of atomic and nuclear weapons, the offense has finally gained what seems at present to be an insurmountable advantage. However, even discounting nuclear weapons given their highly restricted role in combat, offensive capability still appears to have the upper hand, at least on the battlefield. Defending fixed positions over time is not the way of the new doctrine of maneuver that modern armies have followed since World War II (Addington 1994). From that period the development of rapid deployment troops and weaponry (e.g., tanks and aircraft) has favored a more fluid battlefield (Levy 1984). Tactically, armor (rather than fortification) and effective use of terrain (required by increased mobility) have once again become the principal defensive practices on a tactical level (United States. Dept. of the Army 1986).

2.3.5 Summary

In the progression of fortification structures, it is important to note both the relatively static elements as well as those that have changed over time. Static elements include a constant devotion to the control of space around important sites and use (when possible) of advantageous natural terrain. Elements that have changed over time include increased use of counter-offensive methods to provide cover, employment of more complex shapes and configurations to deflect offensive fire, and more centrally planned defensive features, when such defensive features are present at all.

Most of computer network defense more closely resembles the early years of fortification than the modern counter-offensive systems. Increasingly however, more complex spatial patterns are emerging. In order to discuss some of these spatial arrangements, it is necessary to first discuss the underlying structure of computer networks.

2.4 Structure of Computer Networks

With a brief exposition of the history of physical security methods involving terrain and fortification completed, one can proceed to examine similar structures in computer networks. Before detailing the security aspects of networks, however, it is worthwhile to briefly describe the technologies involved in the global Internet, and how the Internet behaves as a spatial entity.

The Internet commonly refers to the interconnected computers and computer networks that communicate via the Transmission Control Protocol / Internet Protocol (TCP/IP) suite (Krol and Hoffman 1993). The multitude of computers and the ease with which they can be accessed has increased information transfer, but it has also increased the incidence of attacks on computers. Between 1988 and 2002, the number of security incidents reported to the first collector of public Internet security data, CERT/CC, has increased every year (save one), with 6 incidents reported in 1988 and over 82,000 reported in 2002 (CERT Coordination Center 2003). The open structure of the Internet, where any computer on the network can route data flows, serve information, or simply be a client plays a large role in the high number of security incidents.

2.4.1 Relevant Communication Protocols

TCP/IP is a simple set of protocols that controls the flow of information across the Internet. A detailed examination of TCP/IP is not strictly necessary in order to understand the spatial aspect of computer network security. However, a brief look at how information moves through the Internet should provide the requisite knowledge.

First, each network interface (e.g., a computer) on the global TCP/IP Internet has an IP address. In the most widely used version 4 of the Internet Protocol (IPv4) the IP address is a unique identifier, consisting of four numbers separated by dots (e.g., 192.168.132.107). Each of the four numbers may range from 0-255 (2^8), with a total number of unique IP addresses totaling 4,294,967,296. In actuality, the amount of available space is much more limited. Large chunks of the available IP address space have been allocated to governments, universities, and industries (Internet Assigned Numbers Authority 2003). Users obtain a unique IP address from their Internet Service Provider (ISP) which in turn procures their pool of available IP addresses from a local, regional, or national Internet registry (Internet Assigned Numbers Authority 2003).

Modern TCP/IP network design can be complex, but the conceptual bulk of it rests on two relatively simple devices: the hub, and the router. The hub acts as a repeating device – any information it receives is retransmitted to all other devices connected to the hub. Devices connect to hubs (and routers) via ports. Ports (like doors) in this case are physical connection points - as where a cable plugs into the hardware. Hubs vary in size (i.e., number of ports) that they have. Many devices

sold to small private networks only accommodate four to eight network interfaces, but commercial hubs are typically larger, often with sixteen, twenty-four, or in some cases more ports available. Hubs can be chained together to form even larger repeater units when necessary.

In larger networks, switches are generally used in place of hubs. Switches perform a similar role as hubs, however data flows are directed to specific ports rather than simply repeated across all ports. As a result, a network implementing switches instead of hubs (generally simply referred to as a *switched network*) is more secure (as well as more efficient) because information (typically) only reaches the intended host, not all hosts connected to the switch (as would be the case with the hub).

The router (sometimes called a gateway) is the network device responsible for transferring information between networks. Its job is to move (or route) data through the network based on a table of available routers that it maintains. When a web page or other information is requested of some computer on the Internet by a host, small parcels of information, called packets, are transmitted from router to router until they are ultimately delivered to and reassembled by the intended recipient.

Many of the computers, or hosts, on the Internet, have static IP addresses. This means that the IP address that they use as their identification to send and receive traffic is fixed. Any change in IP configuration must be manually entered by the user. Static IP addressing generally implies stronger central management of the network, because the network administrator must allocate IP addresses manually.

As network size increased and more users were added to individual networks, more and more hosts allocated IP addresses dynamically. Rather than issue IP addresses on a case-by-case basis, the Bootstrap Protocol (BOOTP) and later the Dynamic Host Configuration Protocol (DHCP) were developed to automate the IP address allocation process (Croft and Gilmore 1985; Droms 1993). Under the model, a host sends out a request for an IP address which a server then responds to. Software then configures the requesting host with an IP address, thus enabling it to connect to the TCP/IP network, and the Internet as a whole.

The most commonly used protocols for transmission (i.e., Transport Protocols) over the Internet are the Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) (Socolofsky and Kale 1991). The most relevant difference between the two is that TCP guarantees delivery, while UDP does not – so TCP is considered more reliable at the protocol level (Socolofsky and Kale 1991).

Communication between IP addresses occurs over ports (different than the sense used above with routers and hubs), similar in many ways to television channels. Just as a television can receive different stations over different channels, computers on the Internet receive data over these different ports. Both TCP and UDP have 2^{16} (65,536) available ports.

Following the above analogy, information *services* are like the various stations available. Services that the reader may be familiar with include world wide web (WWW) service, telnet, and file transfer protocol (FTP). Each has a separate function within the scope of the Internet, and there are hundreds of such services

available. Any given *server* on the network may offer any or all of these services, although in most cases, they only offer a few.

To complete the analogy between television and computer networks, just as a station is assigned a channel within a television community, a service is run on a port. For example, the Cable News Network (CNN) may be broadcast on channel 22. This enables a viewer to rapidly locate CNN if they wish; they do not have to search the range of channels once they have learned the assignment of channels in their locale. Similarly, many services are typically assigned standard (i.e., generally accepted and used) ports. For instance, World Wide Web servers typically wait for connections on TCP port 80. This simplifies the connection process because to access a given service the user already knows what “channel” to tune into. In most cases, the user doesn’t even need to know what port to use. Because of the high degree of standardization, most of these well-used services and ports are preprogrammed into software.

The Internet Control Message Protocol (ICMP) is commonly used in Internet communication to relay basic messages about the status of hosts and transmissions (Postel 1981). It is the protocol used in the network tool “Ping”, a small program designed to see if a given host on the network is responsive or not. A small packet is sent to the host, to which it replies if it is connected to the network. ICMP could be thought of as another “channel” in the previous analogy, useful for diagnostic information about the network.

2.4.2 Topology of Computer Networks

The previous section detailed the important technological and software components of the Internet, but the arrangement of these pieces can vary considerably between networks. The arrangement of these various pieces is called the topology of the network. There are two principal (and strongly related) topologies within a given network. The first of these is the *physical* arrangement of nodes. The second is *logical* arrangement of nodes.

The physical arrangement of nodes includes both wired and wireless networks. The typical means of connecting computer networks still involve a “hard” physical connection between hardware components (e.g., computers, routers, hubs), because physical lines as a whole still offer a far greater rates of information transfer than do wireless methods. The maximum rate of information transfer is called bandwidth, and it varies, in part, according to the type of material that connects the different components of a network. Among the fastest types of material in use is fiber optic cable, and among the slowest is communication via modem and modular phone cable. Bandwidth is often thought of as a “pipe” where having a connection with more bandwidth is like having a pipe with a larger diameter. Just as more water can flow through a larger pipe, more data can flow through a connection with higher bandwidth.

The physical differences between wired and wireless networks tend to impact topology of the network to some degree. Wireless networks present new challenges since there are an infinite number of access points to the network. Wired networks

can manage access by relying on physical control of the hardware, since there is something physical available to control (e.g., locking doors to rooms that house routers). Wireless networks have no such option.

In addition to the wide variety of physical connections possible is the variety of logical topologies that can be defined within a network. The logical structure refers to the complex rule structure that governs how information flows across a physical network. Networks are not formed by physical hardware components alone, but also with software that regulates the flow. Even on a simple network where all computers are physically connected to all others, information may not flow evenly. Sub-networks (or *subnets*) may be formed as groups of computers can define different configurations over which to exchange information.

There are a wide variety of software configurations that affect the logical structure of networks. Within a TCP/IP network, each host defines a set of configuration parameters that establish its gateway (the computer from which it gains its access to the network), as well as other “assisting” servers that may provide network storage space, domain names, or virtually any other type of network service. How each computer is configured within the network can establish hierarchies within it. These sets of dependencies and channels regulate the flow of information and data traffic in ways that are often much more complex than a superficial examination of the physical topology would show.

Controlling the logical topology of a network is often easier than controlling the physical topology. This is true largely because of the dependence on physical

cabling that undergirds most computer networks. It is often costly (and unnecessary) to redesign physical networks to establish rules of access. In many cases, it is quite impossible to establish things like selective filters on a physical basis alone, and logical variation via software rules is the most efficient choice to regulate information flows.

The topology of a computer network is generally represented as a number of shapes interconnected by lines. Figures 2-C and 2-D (below) represent the physical topology and logical topology of the same (imaginary) network. The strong similarity between the two is clear – the main difference between the two is that in the logical diagram the hub does not appear. The logical diagram shows only nodes on the network; since the hub's only function is to connect computers together, it is irrelevant to the logical connection. The other main difference is that the orientations between the diagrams are slightly different. Routers A and B in Figure 4-C are shown vertically, while in 4-D they are shown horizontally. Similar changes in placement can be observed for the computers. Since these are representations of the topology of the network, these changes in orientation or absolute placement on the page are irrelevant as far as being true to the topology is concerned. Such shifts could be used to mislead – for instance one computer could be shown with a longer link-line than others to make it appear somehow periphery. For this reason, care should be taken not to infer anything about the network other than what is clearly shown, and any metric or orientation information should be disregarded. Any representation that tries to encode additional information into a topological diagram

(e.g., using metric distance to represent some other measure of “closeness”) should be clear about its purposes and methods.

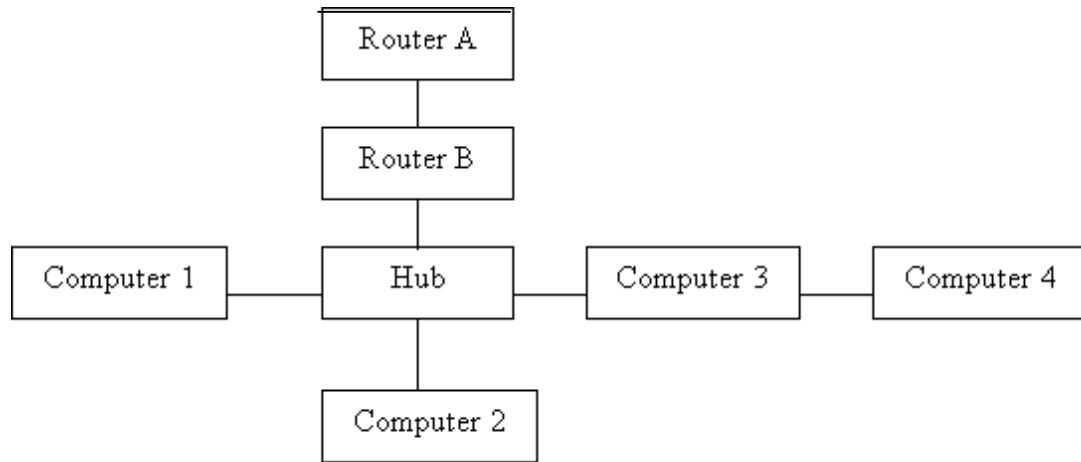


Figure 2-C: Physical Topology of a Network

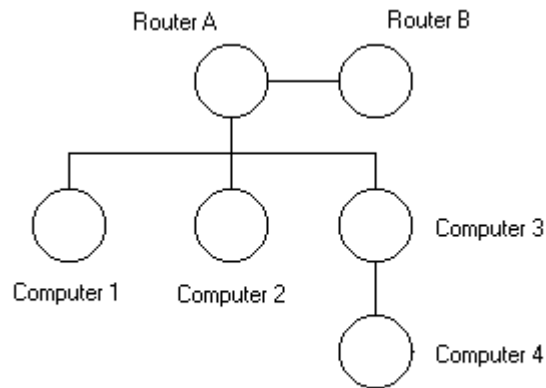


Figure 2-D: Logical Topology of a Network

The logical topology of computer networks is much like the topology of Risk, discussed previously. Each node represents a stopping or processing point of data as it passes through a network. At each node, the data can be examined and then forwarded on, modified, or dropped, according to the configuration of that node. The topology of the constituent networks represents the real underlying spatial structure of the Internet, and attempts to map it usually focus on exposing this path-node structure in some form or another. An attempt to describe the defenses of computer networks, therefore, begins with an examination of similarities between spatial structures in networks and similar spatial structures observable in geographic space. In the next section we discuss some of these kinds of shared spatial features.

2.5 Defense in Computer Networks

2.5.1 Attack in Computer Networks

Before proceeding with an analysis of defense, it is useful to quickly describe some of methods of attack within computer networks. Since there are a wide variety of attacks possible, it is not possible to fully treat the subject here. However, some general patterns or types of attack can be usefully identified.

First, there are differences between passive and active attacks. Passive attacks include software such as viruses and worms that automatically exploit and spread according to the designs of the authors, but not targeting any individuals or groups

specifically. These types of attacks may or may not have malicious intent behind them, and many authors (e.g., the creator of the first Internet worm) claim not to have had any desire to cause the harm that actually occurred (Schwartau 1996).

Active attacks, on the other hand, are perpetuated by thinking humans on intended, specific hosts. They may exploit the same or very similar misconfigurations in hardware or software as some passive attacks, but often they require exploitations of situational weaknesses (e.g., poor network configurations, bad passwords, outdated software) (Scambray, McClure et al. 2001). These attacks are much more interesting because they originate with people who find targets based on much more dynamic criteria and attack computers in a much more varied way than do the automated attacks like worms and viruses.

Analysis of these attack decisions may be very useful to solving network security problems. Why do individual hackers or crackers (i.e., those who attack computers or networks) choose to target one computer over another? Some of this may be accounted for by rational calculations of expected cost (e.g., perhaps time, or risk of being caught) versus expected gain (e.g., perhaps useful information, money, or a way to attack another target with less risk).

Developers of Intrusion Detection Systems (IDS) are faced with the challenge of defining rule sets or heuristics and building them into software that can analyze data flowing through networks for attack (Kemmerer and Vigna 2002). This technique relies on comparing information extracted from data packets traveling over the network and comparing them to the rule set to see if they match the given criteria. In

the case of Snort (one such IDS) new rule sets are released regularly, and users can also define their own rules to fine-tune the system (Richard 2001).

The general ideas of active versus passive attacks should suffice the purposes of this thesis. Attacks may take so many forms that defenders must be creative, thorough, and above all vigilant to defend against them.

2.5.2 Spatial vs. non-Spatial Elements of Defense

Defense within computer networks involves many cooperative technologies and methods. Many of these are for the most part non-spatial. Examples of these technologies and methods include authentication, encryption, and access auditing (Kemmerer 2002). Authentication attempts to make sure that a user or host is what it says it is. This is commonly accomplished by means of a password, although other methods (e.g., certificates) are possible. Encryption has been used for millennia to guard sensitive information so it is not surprising to find its use in cyberspace. Although most protocols were not originally encrypted, there is recently a greater tendency to code traffic over the Internet for privacy and economic reasons. Access auditing aims to make sure that users have only a limited scope of abilities on a given system, with the aim of reducing the harm from a rogue user or compromised account.

Other elements of defense in computer networks *are* spatial and these mainly involve the structure – physical and logical – of computer networks. The arrangement of servers, hosts, and other networking components makes a great difference as to how the network performs. It is these spatially-related technologies

and arrangements that are of interest here. Many of the most common technologies have links to counterparts in the physical world, although their meaning may have changed in transition.

A review of some of those most relevant to this research are presented in detail below. Three main types of spatial defenses are considered here: firewalls, bastion hosts, and demilitarized zones (DMZs). Although this is not an exhaustive list, it represents the three most common technologies/constructions within the context of computer networks that explicitly borrow concepts from physical security. These three examples of spatial defenses are important because they establish an existing metaphorical connection between real space and computer networks. In the following section, the metaphors of spatiality are extracted in detail. This ultimately allows for a more refined theoretical connection between defensive structures in real space and defensive structures in computer networks.

2.5.3 Firewalls

Firewalls are perhaps the primary type of spatial defense that typically attempts to restrict and define network traffic between networks, usually a small network (such as a LAN) and the Internet. Although specific types vary, it is essentially a filter, such that all traffic bound in or out of the network must pass through the firewall. Traffic is inspected (in a variety of ways) and compared against a set of rules specified by a security plan. Only selected (i.e., defined) ports and IP address sources/destinations are allowed, thus channeling the flow of data through a network.

In practice, there may be several routes in to and out of a network. For a firewall-type defense to be successful, it must be employed in all routes.

Firewalls can take the form of software or hardware, but are often combinations of both. One example of a hybrid-type firewall is a Network Address Translation (NAT) server (Zwicky, Cooper et al. 2000). NAT servers are essentially routers that utilize a single or small number of IP addresses on behalf of a much larger number of hosts. These hosts are configured with “private” IP addresses with special properties that prevent them from being communicated with directly. The NAT server handles all traffic on behalf of its client hosts, and becomes the “public” face for what can be a much larger network existing “behind” it.

The deployment of firewalls highlights the similarities between the topological space of computer networks and the geographic space in which human settlements exist. During the deployment of a firewall, the space surrounding a network is occupied and a new structure interferes with “normal” traffic and movement. This corresponds quite closely to the construction of walls and gates around settlements. The firewall may even be dynamically configured by other software that monitors for intrusion attempts.

Firewalls (like most other networking components) are subject to attack, and the results are similar to the results of an attack on a real wall. When a firewall is attacked in a computer network, it can be reconfigured (or completely disabled) which allows for unwanted movement through the network. This kind of reconfiguration is essentially the same as a breach on a real wall that may be opened.

In both cases the topological properties of the area are changed. Points that were not directly connected become connected in real space and computer networks when this happens.

To argue that firewalls are spatial constructs makes more sense when describing changes in topology that occur when a firewall is in place than it does when attempting to fully articulate how the fully metric properties of real walls could be represented. As mentioned earlier, the size and shape of walls changed from place to place according to natural terrain and evolved over time as offensive methods became more sophisticated. The “thickness” and “shape” of a firewall is not defined by the topological properties, and thus appears to be open to interpretation. Part of the difficulty in designing useful network security maps is that the appropriate representations of these objects are not intuitive (Anderson, Brackney et al. 2000).

2.5.4 Bastion Hosts

Another method of spatial defense with an equivalent in real space is the construction of a “bastion host”, a public face and the most heavily fortified element of a network (Zwicky, Cooper et al. 2000). In one of the original formulations of the bastion host, it was identified as “a critical strong point in the network’s security” (Ranum 1993). The assumption behind this defense strategy is that the most public face of a network (e.g., the host that offers services such as web hosting, file servers, or terminal access) is the most likely to be attacked. Thus one can focus one’s security efforts on a single host rather than distributing one’s efforts among many hosts.

Bastion hosts were so-named because they “overlook critical areas of defense” (Ranum 1993). In the Renaissance forts in which they originated, bastions were an attempt to provide protected means to defend gates by focusing interlocking fields of fire such that the adjacent two bastions could target any attacker trying to breach a gate or wall. Any bastion, which was itself targeted, could be similarly defended by its two adjacent bastions (De la Croix 1972).

2.5.5 Demilitarized Zones

One final spatial security element sometimes employed in network defense is the construction of a perimeter demilitarized zone (DMZ) network, on which the bastion host might operate (Zwicky, Cooper et al. 2000). The physical separation afforded by the DMZ (sometimes called a perimeter network) prevents intrusion by creating a space between two barriers. One barrier exists between the Internet and the perimeter network while another barrier exists between the perimeter network and the internal network. Thus if an attacker successfully penetrates the first barrier he or she must still overcome another barrier before he or she is able to access the sensitive information restricted to the internal network. The DMZ essentially provides greater depth of defense so that there is not a single point of failure, but rather several reinforcing security mechanisms.

2.5.6 Misleading Analogies

The names for these components of network security (firewalls, bastion hosts, and DMZs) can lead to a misunderstanding of the technologies and/or methods that

they represent, and this drives a wedge between the translation of security concepts developed for physical space to those developed for use in network security. For instance, the term *firewall* is a security term in the physical realm, but not a military one. Its meaning in physical space refers to a wall intended to prevent the spread of fire (Oxford University Press. 2002). Although it does represent a compartmentalization, it is not principally a protection from *attack*.

Bastions, in military history, were certainly heavily fortified, but not because they were to serve as an entry point or exit point for distribution of services. Formal demilitarized zones in military history have been defined in a number of treaties. Perhaps the most famous of which is that created near the 38th Parallel between North and South Korea. While such a region can and does serve as a buffer or frontier, stipulations in the treaty dictate that fortifications and military assembly or maneuver are strictly forbidden within a certain defined area. This seems quite counter to the admonition in network security manuals to pay special attention to the host's security and that it should be the most fortified host given also that bastion hosts should be located in a network DMZ (Zwicky, Cooper et al. 2000).

Despite the differences in terminology, however, there are astounding similarities between the methods of securing computer networks and fortifying cities. This stems from the fact that both are faced with a similar problem to overcome: they have to develop a system that can keep attackers out while allowing legitimate traffic to flow. The metaphor of space and movement permeates through computer terminology as a whole and so network security engineers are on to something when

they spatialize the problem they try to solve. By doing so, they tap in to a wealth of knowledge that physical security engineers have been dealing with for thousands of years.

2.6 Metaphors of Spatiality Within Computer Networks

One way to begin to bridge the gap between defense in computer network security and defense in real space is to propose and discuss how some important spatial factors in real space play out in computer networks. Of the many geographic variables that make a difference to security in the real world, some are: distance, population density, the mobility of the population, how visible the targets are, how much movement occurs, and the degree of interaction. Each of these is examined in detail below as a possible metaphor for understanding security concepts in computer networks.

2.6.1 Distance

Distance, “the amount of separation between two points,” is the defining characteristic of metric space (Woolf 1981; Oxford University Press. 2002). It matters in real space to such a degree that it is posited that while “everything is related to everything else... near things are more related than distant things” (Tobler 1970). Does this general relationship hold true within the specific context of security? Within the scope of human history, it seems likely: given that human mobility was highly limited until the advent of safe long-distance sea-travel and later,

motorized transport and airplanes, it makes sense that people would tend to attack (and defend) against enemies that were closer rather than further away.

If distance is at least relatively well-understood and well-defined in real-space, it is much less-so in computer networks. Huffaker *et al.* (2002) have proposed and tested four distance metrics for the Internet. These distances relate strongly to types of *functional distances* that people use in real-space to describe degrees of separation between points that are not strictly metric. For example, a person may readily answer a question of “How far is it?” by giving a reply in terms of the amount of time needed to get there. Each distance metric that Huffaker *et al.* propose is measured against packet latency, the round trip time (RTT) it takes for packets to travel from host to host under the assumption that longer “distances” should require more time to traverse. In effect, time is the ultimate measure of proximity on the Internet, and the metrics proposed by Huffaker *et al.* are other approximations of it. Surprisingly, Huffaker *et al.* found that geographic great-circle distance was a better predictor of “nearness” on the Internet than was the topological “IP Path” metric that used the number of “hops” (i.e., intermediary routers) as a predictor of RTT. Geographic distance may also be one of the best predictors of source of attack. One U.S. Internet security company reports that more attacks against its clients have their source in the United States than in any other country (Belcher and Yoran 2002).

2.6.2 Density

Population density makes a great difference in the types of weapons and methods of attack and defense that can be employed. One simple example of this is the

difference between rural and urban combat. One of the difficulties of urban combat is the mixing of combatants and non-combatants, so that it is difficult to attack the one while protecting the other.

Another impact of density is the attractiveness of targets. Despite Sun Tzu's admonition to avoid capturing cities, they have proven attractive targets throughout human history. The reason for this is that the bulk of the economic and political power of an entity usually comes from the city (O'Sullivan 1991). Thus, when looking for a way to do the most damage as quickly as possible, high density population centers are ideal. An example of this approach to war can be seen in during the United States Civil War as General Sherman burned many Confederate cities in order to destroy enemy morale and economic support (Addington 1994).

Some theories of war and territoriality rely on density to predict conflict, since competition for scarce resources increases with density (Wright 1965). When density is low, populations (or individuals) may simply choose to move to a new location rather than to fight. As density increases, there are fewer and fewer locations to move to, and fighting and keeping territory becomes a more rational approach. At some point, perhaps, widespread fighting similarly motivates individuals and populations to compromise and cooperate when fighting seems to costly, and fleeing is no longer an option.

In computer networks, density may play a similar role. While attacks on networks have increased over time, so has population density. There are, at the same time, more targets and more attackers. The proliferation of various technologies

(e.g., different operating systems, applications, protocols) means there are more and more avenues of attack to use.

A principal way in which anonymous attackers acquire targets is via network and port scanning (Scambray, McClure et al. 2001). For many Intrusion Detection Systems, scanning constitutes an attack itself, since it generally precedes a more concerted effort, and in any case is unsolicited probing. One of the most popular of scanning tools is *nmap*, a Linux/Unix based tool that users can use to look for computers on the Internet, and examine potential targets more closely for weaknesses. A typical method of use for the attacker is to specify a certain range of IP addresses to scan. IP address allocation (within the subnet) is often *ad hoc* or random, and so in a given subnet not every IP address that can be assigned actually is assigned, and they are not assigned in any sequential order.

So, like geographical spaces, there are different densities at different scales. In the entire range of IPv4 address space, a growing percentage of IP addresses are used. Since IPv4 address space is not distributed uniformly, different segments are more full than others. Large scale efforts to determine whether certain IP ranges are more often probed or attacked than others (in general) are relatively recent. Moore *et al.* (2001) inferred widespread occurrence particular attacks based on backscatter, but did not report findings based on IP destination address. In some cases, such as the spread of certain Internet worms, researchers at the Cooperative Association for Internet Data Analysis (CAIDA) have released reports detailing IP scanning/spreading preferences, but these seem to be related to a bug in the code of

the worm rather than a preference for more densely populated regions of the Internet (Moore, Paxson et al. 2003).

2.6.3 Population Mobility

Historically, the transience of a population has played a large role in their military interactions. Initially, the earliest farmers made themselves a ready target by combining a steady food supply with a fixed location. One theory of political development theorizes that the hunter/gatherers who remained more warlike and mobile became the masters of the agriculturalists who were more sedentary (Keegan 1993). Humans became prey for other humans, rather than simply competitors. The spatial dimension of this theory is relevant here, as it was the immobility of the farmers that both enabled them to farm, and made them into targets.

Later, the lifestyle of the nomads lent itself to military dominance. The development of strong cavalry by the Huns and Mongols enabled them to conquer larger, more advanced, but more *fixed* targets (Keegan 1993). One of the requirements of guerrilla warfare is the hit-and-hide methodology (O'Sullivan 1991). Recently, terrorists have posed a problem for states because transient individuals or organizations that have no formal connection to territory perform the acts of violence, and so reprisal becomes more difficult and politically costly.

There are analogs to this kind of immobility within the context of computer networks. Firstly, most organizations tend to have a static presence in cyberspace, meaning that IP addresses, or at least DNS names, remain the same over time. This is simply because, like physical addresses, it is important for customers (and users in

general) to be able to easily find the sites and businesses that they are looking for. In fact, this is so important, that as the Internet became commercialized during the latter half of the 1990's, "cyber-squatters" would register for names of popular businesses that had not yet created a web presence, and would charge fees to these companies before they would surrender their rights to the name. In this case, even the *anticipated* "location" was crucial.

Another analog of mobility is the allocation of IP addresses. IP addresses used to be allocated centrally by a network administrator, meaning that each IP address was individually and consciously assigned by a human. That address would be static – that is, it would not change – and a host (or computer) would have the same IP addresses each time it was used. Later, this kind of central administration became burdensome, and led to the creation of dynamic addressing schemes, such as BOOTP and DHCP. Under this system, a host would request an IP address from a local server as it was rebooted or at the command of the user, or at some other specified interval of time.

As a result the configuration of the network was more fluid. A given IP address could not reliably be associated with any given user, without other outside means of verification. An open question remains: does this kind of dynamic address allocation make any difference to security? One supposes that it could, although probably it depends someone on the length of time that an IP address "lasts" before it expires and is assigned to a different computer.

2.6.4 Visibility

The visibility of potential targets also makes a large difference in the conduct of military action. Fog, clouds, and other elements of weather can preclude many air assaults because of visibility (Collins 1998). According to traditional military wisdom, guerrilla operations tend to be most successful in areas with heavy vegetation (O'Sullivan 1991; Collins 1998). A decreased ability to see (or more accurately, sense) an enemy tends to make attacks more cautious, and difficult to organize (Collins 1998). One might say that the uncertainty present in low-visibility conflict situations tends to favor the defense since it is (presumably) more adapted to the native conditions.

Visibility in the context of computer networks can take on one of several meanings. First, a certain amount of information transmitted over a computer network can be acquired passively. As discussed previously, a hub is a central connection device that allows for full exchange of information between all interfaces connected to it. It is known (as many other relay technologies are) as a repeater. For example, if computers A, B, and C are all connected to the hub, A's transmission to B is also transmitted to C. This is not unlike the situation of several individuals talking in a room. Although Fred may be talking to George, Sam can hear the conversation, even if he isn't interested in what they are saying. Like such a conversation, computers A, B, and C "hear" all the traffic going to or from any one of them, but only "listen" for traffic directed to them.

Like a hub, a switch moves traffic between hosts connected to it, but unlike a hub it attempts to isolate traffic so that a computer is only sent traffic that is sent to it. This segmentation of traffic allows for improved efficiency since a host only hears a portion of all traffic rather than all of it. It also allows for greater security, as (in the absence of a determined attack) only computers that “need-to-know” are involved in the transmission. The result of this is that users on a hub are more visible in this sense than users connected to a switch, since more information is passively available regarding users on hubs. It should be noted, however, that they are more visible only to other users on the hub. As in the example of the room described above, Fred, George, and Sam can all see each other within the room, but someone else outside the room cannot. Visibility in this sense is not universal, but local.

Another method of assessing visibility lies in the use of public versus private IP addresses. While most of IPv4 address space is public, some blocks are private, and non-routable (Rekhter, Moskowitz et al. 1996). There are three such blocks of Private Address Space: 10.0.0.0 to 10.255.255.255, 172.16.0.0 to 172.31.255.255, and 192.168.0.0 to 192.168.255.255. IP addresses within these blocks are called private because they cannot interface directly with the Internet. In order to connect, they need some kind of intermediary routing that handles the interface on behalf of them. One such method is the Network Address Translation routing discussed previously.

Visibility comes into play here because this private IP space cannot be explored in the same way as public IP address blocks. For instance, any given public IP

address can be tested for accessibility by issuing a “ping”, or by looking for open ports and services using a program like *nmap*. A crude Internet census (or estimate) is accomplished by sampling these IP address, either in full (which is not technologically feasible) or else by statistical sampling. Networks using the private IP address space (i.e., private internets) cannot be sampled in this way, because it cannot be easily determined whether a private internet is “hiding” behind a public address. In theory, almost every public IP address could be hiding a vast Internet by employment of a NAT router or similar device.

Visibility within computer networks is one of the most difficult to fully define. The metaphor of visibility is itself a metaphor for general perception. In this way each of the other five metaphors discussed here may factor into issues of visibility. A “nearer” network may well be more visible, since when scanning times are reduced more detail can be seen. Interaction or movement on a network may make it a more interesting target, and thus in a sense more visible in the same sense as a “vivid” color may be more visible.

2.6.5 Movement

In addition to the metaphor of population mobility discussed previously, it is possible to identify the related spatial metaphor of movement. The importance of movement in a contest of force is one of the more obvious factors that contribute to the outcome. Greater mobility enhances the ability to attack (Biddle 2001). Linear movement (i.e., getting a tactical group from Point A to Point B) is important in deployment of force. The important advantage of railroads during the mid to late-

nineteenth century helped both the United States Army soldiers during the Civil War and the Prussians in Europe during their expansion (Glaser and Kaufmann 1998). The cruising speeds of aircraft carriers and tanks, among a wide variety of other (especially logistic) concerns similarly affect military operations today.

Controlling and limiting the movement of the opponent, while at the same time enhancing one's own ability to move, is the essence of a sound defense. Fortifications aim to limit the movement of enemy troops and munitions, preventing them from reaching their targets. Armor, the personal instance of cover, attempts to stop (or render less harmful) a weapon or munition aimed at the most critical areas of the body.

When assessing spatial defensive features in computer networks, then, it is key to define movement. In this case *traffic* is the principal clue towards locating movement. In the field of Information Technology, traffic refers to the flow of data across a network. Data in this case, takes the form of packets of information that contain header information and content, very much like an envelope in that has an address and content. The movement (or traffic) in a computer network occurs as each packet is sent and received by each computer in the route from source to destination.

Previously, bandwidth was discussed as a limiting factor in the speed of data flows, with differences in the physical method of connection affecting potential bandwidth in a network. The speed of networks is important for security because it

affects the type of attacks that are the most effective. A network's speed also affects how attractive it is as a target.

Bandwidth affects how vulnerable a certain network is to attacks that rely on overwhelming the victim with traffic. For instance, two computers A and B are connected to the Internet. Suppose computer A can send and receive traffic at the relatively fast rate of 100 Megabits per second (the equivalent of *fast Ethernet*, often found on college campuses). Computer B's maximum bandwidth is much smaller, say 500 Kilobits per second (the rough equivalent of a cable-modem connection). One method of attack on the Internet involves *flooding* another computer with data, essentially drowning them in so much data that the target computer can not in a timely matter separate legitimate data from the irrelevant unsolicited data sent by the attacker. Given the maximum bandwidth of Computers A and B in this example, A can overwhelm B with data, but B cannot similarly overwhelm A. The smaller bandwidth of B's connection prevents it from sending out enough data to choke A's connection.

Bandwidth is also thought to play a major role in the attractiveness of a network as a target. Networks with fast connections to the Internet are valued more highly as targets because they offer platforms for attacks like the above. Any computer on that network is then a more likely target, not because of any configuration of its own (such as operating system vulnerabilities) but rather because of the characteristics of the network on which it resides. This might compare to a village or town with high strategic value because of its relationship to its surroundings (i.e., its situation within

terrain) rather than some other (e.g., industrial capability) that it possesses (i.e., its site features).

2.6.6 Interaction

Conflicts do not typically arise out of nothing, but depend upon relationships between combatants. In some cases, these relationships are more developed than others. One way that combatants are related is through their explicit interactions. In physical space, one important set of interactions that tend to produce conflicts are those that involve economic or resource transactions (Wright 1965). These might be called the *nature* of the interaction. The *method* of interaction is also of interest, where this term refers to the mechanism or avenues of the relationship – the interface. These interactions may be examined within the context of computer networks as *services*.

Most people within the United States have probably heard the term *server* in connection with Information Technology. Server in this case refers to a computer that stands ready to offer information or data in some manner, either to the public or else privately. The information or data is bundled in some way, such that file transfer is somewhat different than serving web pages. Technically, each of these different bundles is called a service, and one server may offer (and often does) several different services. Services can also be offered in bundles that correspond to services offered in real space, like banking, stock trading, news updates, mail, and so on.

These kind of “real” services probably also attract attacks on computers on networks just like they do in real space. Thus computers with valuable information or that protect (digital) money are presumably more likely targets than those that do not. More relevant to the current discussion of military-related spatial metaphors in computer networks are the “technical” services that often correspond directly to protocols and ports, distinctly spatial elements of networks.

A server on a network offering some service X typically does so on a standardized port, as discussed previously. Under un-fortified conditions (e.g., no firewall is employed) each open port can be detected. One method to do this is by using whatever application is usual to contact that service/port. For instance, since World Wide Web pages are typically offered on Port 80, one might use a browser and provide it with some (possibly random) IP address to see if that service is offered on that port. A more common method is to use a software scanning tool that automatically contacts a range of IP address and ports to determine the presence of such services.

Of the six metaphors discussed in this section, interaction is most likely the least spatial. Interaction becomes spatial when one realizes that it is important (in both real space and computer networks) *where* the service is provided. The range of port space defines the limits of the location of services. The numerical individuation of each port does not immediately provide a spatial reference, but their sequential ordering does. These spatial references may not be “real” in the sense that any sequence must necessarily be spatial, but the abundance of usage of spatial concepts

within the computer networking literature (e.g., *IP addresses*, *port space*) evokes a spatial ordering within the mind regardless of whether that mental representation is necessary or not.

By virtue of the sequence of their numerical assignments alone, some ports are nearer to each other than others. Port 22 has fewer ports between it and port 25 than between it and port 137. It may be that the sequential ordering of ports is the only sense in which spatiality applies to ports, but even if this is so, it is still worthwhile to consider whether even this conception of the nearness of ports (and therefore, perhaps of services and interaction) is important to the security properties of a host or network connected to the Internet.

2.6.7 Summary

Metaphors of space within computer network literature and discourse abound. Of the metaphors of spatiality given above, some are more spatial than others. The aim in discussing them here is to set the context for a discussion of two specific aspects of military security (i.e., cover and concealment) relevant to computer network defense. Table 2-A (below) lists some of the bridge concepts between geographic space and computer networks as a review.

Geographic Concept	Computer Network Metaphor(s)
Distance / Functional Distance	<ul style="list-style-type: none"> • Time • # of Intermediary Routers (IP Path) • Geographic Great Circle Distance
Population Density	<ul style="list-style-type: none"> • Occupation of IPv4 Space • Subnet Density • # of Computers on Internet • # of People Using Internet
Population Mobility (Transience)	<ul style="list-style-type: none"> • DHCP, BOOTP, & Other Dynamic Addressing Systems • Dialup Users • Shifting ISPs
Visibility	<ul style="list-style-type: none"> • # of Services Run • Standard vs. Non-Standard Port Usage • Bandwidth Usage • Firewall Use
Movement	<ul style="list-style-type: none"> • Data Traffic Flows • Bandwidth Potential
Interaction	<ul style="list-style-type: none"> • Types of Services Run • Standard vs. Non-Standard Port Usage

Table 2-A: Metaphors of Spatiality Within Computer Networks

2.7 Generalizing Spatial Elements Valuable to Computer Network Defense

Based on the previous metaphors of spatiality within computer networks, it becomes possible to return again to two of Clausewitz’s fundamental properties of terrain – cover and concealment – and to provide an argument for their manifestation in computer networks. This section attempts to bring together both physical and network conceptions of security and merge them together using cover and concealment as the unifying concepts. Some of the arguments for cover and

concealment in computer networks discussed below will be empirically tested in the next chapter.

2.7.1 Cover in Computer Networks

Before defining how cover is manifested within computer networks, it is useful to recall that terrain was defined for this study as consisting of “configuration of the medium of conflict,” while cover was defined as “protection from observation and fire.” Each of these definitions is primarily based in the context of physical (largely military) security, but especially in the case of “terrain” an attempt was made to reduce land-centric biases into the concept.

Cover exists in the surrounding environment of the medium of conflict. In addition to the “natural” cover passively supplied by the environment, humans have supplemented terrain with fortification to enhance opportunities for cover. These fortifications have become part of the terrain, blending the human-made with the pre-existing.

Objects are important for cover, but the spatial relationships between the terrain and the combatants are important as well. Walls, for instance, perform their intended function of controlling movement only if their size and shape is appropriate to the technology and methods with which they co-exist. There may be excellent opportunities for cover, but if they are interspaced too widely, a combatant may not be able to fully realize their potential if he or she is vulnerable as he or she moves from one object of cover to the next.

Since cover by definition protects its user from observation and fire, it is not strictly necessary to differentiate the two in this context. It will be shown that concealment does not reap this analytical shortcut. For now, it is enough to say that it can be difficult to separate the two. This is because both scanning and attacking within computer networks involve data transmission of some kind. Since firewalls block data transmission (based on a wide variety of configurable criteria) it makes sense to say that firewalls provide cover.

Even in real space, cover rarely provides *complete* protection from observation and fire, but generally only does so in a limited fashion. A complete outer wall may guard against attack from without but it does not protect from fire originating from the inside. The problems associated with the directional component of cover in real space correspond very closely to those associated with computer networks. For instance, one concern with firewalls is that they can be circumvented by unknown or unexpected changes in topology. An example of this might be an attacker using an unsecured modem and telephone connection to acquire access to the firewalled network. The physical counterpart to this could be someone finding a hidden entrance, tunneling under the walls, or climbing over them.

The counter-offensive concept of *cover-fire* finds almost no employment in computer network defense. An example of this could be a “counter-hack” in which an assailed network administrator discovered the source of his or her attacker and actively attacked them back. The effect, like in physical space, might be that the

instigating attacker is forced to devote resources to defending his or her own network, thus reducing the threat.

High degrees of movement and interaction within a computer network may provide similar difficulties in computer networks as they do in real space. With high amounts of data traffic, the “scalability” of a firewall solution is tested. The resources and techniques that provide cover for small networks with fewer traffic requirements may not work well for large networks. This compares readily to cities that outgrew their circumscribing walls – the protection afforded by the walls also became a bottleneck to movement. Similarly, systems that must offer a greater number of services may not benefit from firewalls to the same degree as smaller networks that offer fewer services. A small network may be able to limit access to and from only a few IP addresses based on personal familiarity, but large sites with very public faces must deal with so much interaction that determining the presence and nature of an attack can take much more time and effort.

Although the general metaphor of firewalls-as-cover seems to fit, it is difficult to apply the metaphor in all cases without more careful study of specific behaviors of both cover and firewalls. One such study is conducted in the next chapter, as firewalls are tested to see whether the firewall displaces attacks to itself when it provides security for another computer. This follows from the earlier hypothesis that a physical object is effective as cover (and thus can only be *considered* cover) if an attack is displaced to it, as a shield absorbs the attack intended for the person wielding it.

2.7.2 Concealment in Computer Networks

Concealment was defined within the scope of this research as “protection from observation but not from fire.” The concept of concealment within computer networks may be somewhat easier to define than cover since most of the conceptual work rests on the concept of visibility. One possible interpretation of concealment within a computer network is simply to not be visible.

Unfortunately, as alluded to earlier, it is difficult to separate attack from observation within a computer network since there is very little in the way of passive perception. In order to “see” what is available on a given a network or computer not on the same hub, a scan must be made, and this requires the transmission of packets to the target computer. It amounts to “probing” and it can be difficult to ascertain whether the probing should be considered an attack or not.

Since interaction between computers and networks generally occurs on standardized ports, one concept of concealment may rest on using the “less traveled” ports for interaction. For example, an attacker may expect the FTP service to be on port 21 and scan that port to see if it is open. If it is not, the attacker may conclude that there is no FTP service and scan another computer for the same port. If the same service is run a different port, the would-be attacker may never find it. It is concealed by existing in an out-of-the-ordinary place.

Is this kind of concealment spatial? There seem to be types of concealment that are not spatial. For instance, someone may conceal the truth from another. Even this, however, relies on spatial separation of some kind, since the knowledge of the

event in question resides within the brain of one individual but not another. The ordered numbering of ports suggests that they are at least minimally spatial. The term “port space” is in widespread use. The standard pairing of services and ports suggests a commonly understood addressing system. All of this suggests that the service/port pairing is spatial, and if so then it seems reasonable to understand the above-mentioned hiding or concealment of ports in terms of that spatiality.

Dynamism may impact concealment as well. Consider for instance a computer network that assigns its IP address dynamically via DHCP. In this circumstance, a given computer acquires a different address within the network at some interval of time – perhaps every day. Computers within this network address space are more mobile, and as a result, less visible. In order to locate the computer, more time and effort are required than if the computer could be reliably found at the same address each day. The protection from observation in this case is provided by the large space in which it can “hide”.

The benefits afforded by this kind of mobility within the network might be hindered by high occupation (i.e., high population density) on that network. If there are fewer open addresses (analogous to open spaces) then less mobility can occur (as might happen on an overcrowded freeway). On the other hand, high densities on networks may provide an attacker with more other targets, so any one computer has a lower chance of being attacked.

With so many possible interpretations of visibility, observation, and therefore concealment, empirical testing is needed to determine which of these interpretations

bears the most similarity to how concealment behaves in real space. In the next chapter, the idea of concealment within computer networks is empirically tested as (a) a “hiding” of well-used services on non-standard, under-used ports and (b) a function of interaction where running more services makes for a more visible and thus less well-concealed computer.

2.8 Summary

Before proceeding to the empirical evaluation of the military metaphor in computer networks, let us review some of major points presented in this chapter. First, the concept of terrain was discussed in the context of military spatial analysis. Terrain was defined as the configuration of the medium of conflict, and the argument for looking at terrain in a dimensionally compressed topological way was presented, using the game of Risk as an illustration of the concept. In this sense the fully continuous, metric conception of terrain is abstracted to what essentially amounts to a network. This network aspect of terrain is what undergirds the application of military spatial security metaphors to computer networks.

Fortification was discussed in terms of the evolution of attack and defense strategies. This discussion, along with the presentation of the structure of computer networks served as the basis for discussing three basic types of spatial network defense: firewalls, bastion hosts, and demilitarized zones. The metaphorical mappings between the computer technologies and structures and their real-world counterpart were examined to reveal their differences and highlight their similarities to one another.

The argument proceeded from these very specific military metaphors to more general concepts of how militarily significant spatial variables might manifest in computer networks. Finally, cover and concealment were examined in detail to show their spatial nature. In effect, employment of either cover or concealment creates a small topological network in which the covering or concealing object exists between the attacker and defender. These simple structures are what allow the military metaphors to be tentatively applied from real space to computer network defense.

The appropriateness of the metaphors, however, can not be judged solely on the basis of this single spatial similarity. The next chapter presents the argument for looking at functional similarities between objects that are used as cover or concealment in real space and objects that are used as cover or concealment in computer network defense.

3 METHODS OF DATA COLLECTION AND ANALYSIS

3.1 *Introduction*

The previous chapter described some of the theoretical connections between cover and concealment in real space and in computer networks. This section details the methods of three experiments designed to provide empirical evidence for cover and concealment within computer networks.

These experiments attempt to determine whether cover and concealment behave or function the same way in computer networks as they do in real space. It is an attempt to begin to ascertain to what degree cover and concealment share common characteristics, and therefore the degree to which each is appropriate as a metaphor. While it is understood that not every characteristic will map between source and target domains in a metaphor, it is an assumption of this thesis that an appropriate metaphor will have met some minimum threshold of such mappings. While the previous chapter established some basic spatial commonalities between network defense topologies, the experiments provide information as to whether these similar spatial structures also give rise to similar behaviors in attacker/target relationships.

Cover, in computer networks, is measured by observing the property of displaced attack. If a server offers cover in the same way as geographic objects, it should become attacked more often as it protects other (client) hosts. Conversely, the client hosts should experience a decline in probe and attack during the same period. The experiment on cover seeks to test this hypothesis.

Concealment is measured by comparing the rates of attack on systems running services on either standard or non-standard (i.e., “hidden”) ports. The main hypothesis here is that a computer running services on standardized ports is more visible, and thus is more likely to be attacked.

These two experimentally evaluated concepts of cover and concealment are compared to a quantitative study done by Otterbein (1970) on various military factors in primitive societies. Otterbein’s extensive factor analysis of these societies included defensive measures employed and incidence of attack. Although Otterbein was not particularly interested in how cover and fortification impact rates of attack and the military success of the cultural unit, his extensive data makes such an analysis possible.

The ultimate purpose of these three experiments on cover and concealment is to provide quantitative data to support a comparison between cover and concealment in computer networks and in real spaces.

3.2 Testing Environment and Equipment

3.2.1 Network Environment

The experiments were set up on two separate computer networks. The experiment involving cover took place within an Ethernet network at the University of California at Santa Barbara. IP Addresses were allocated statically throughout the experiment. The experiment involving concealment took place on a cable modem

network in Goleta, California. IP addresses were allocated dynamically by the service provider.

The cable modem network was configured to prevent many incursions into users' computers. As a result of this policy, several inbound ports were blocked at the router level. These ports were: 25, 80, 111, 119, 135-139, 1900, and 27374. These ports correspond to common services that are often attacked, or to ports opened by malicious programs operating unbeknownst to a user.

3.2.2 Hardware and Operating Systems

All computers involved in experiments one and two were standard PCs ranging from Pentium to Pentium 4 class machines. Three PC's (the firewall for the cover experiment and both servers in the concealment experiment) had the Mandrake 9.0 Linux operating system installed. All others used either Microsoft Windows ME, 2000, or XP.

Each computer had only one Ethernet card installed as a network interface, with the exception of the firewall in the cover experiment, which required two.

3.2.3 Software

Libraries *libpcap* (for Linux systems) and *winpcap* (for Windows systems) had to be installed on all participating hosts as a prerequisite for gathering data traffic. Both software libraries are designed for packet capture on a network connection, which means that the packets, the basic units of transmission for all network protocols, are

available for inspection. This allows for complete examination of network traffic by other programs.

Analysis of data flows for intrusion was achieved (for both operating systems) with *Snort* (version 2.0). *Snort* was employed in this study because it is free for academic use, and could be installed on both Windows and Linux machines. The use of *Snort* enabled real-time (i.e., immediate) intrusion detection and logging. *Snort* logs were kept locally on the hosts during the experiment, and were extracted once the experiments were completed for analysis.

Two forms of record of the attack were kept. The first was a text file (*alert.ids*) generated using the *-fast* option to keep disk writing (and thus taxing on the system) to a minimum. The second form of record was generated using the *-b* option, which logs the binary packets to disk. This allowed for more thorough analysis of the packets later, since the text log file could be regenerated with full (instead of fast) information. *Snort* was also configured to analyze the traffic for attacks using the standard rule sets that ship with the software.

Fire Daemon was installed on all Windows machines to assist the automatic loading of *Snort* at boot time. *Fire Daemon* also monitored the status of *Snort*, and would restart it immediately if it shut down for any reason.

The hosts were scanned using the freeware port-scanning tool *nmap* to determine whether they had any unusual open ports that would indicate that they had already been compromised. No hosts showed signs of intrusion at the outset of the

experiment. In order to prevent these scans from being included in the actual attack logs (and thus possibly biasing the results), Snort was disabled prior to running them.

Additionally, Norton Antivirus version 7.6 had previously been installed on all Windows hosts to verify that they had not already been infected with a virus. In order to maintain the integrity of the hosts participating in the study, the latest operating system patches were installed at the beginning and throughout the experiment via the Windows Update Service (for Windows hosts) and Mandrake Update (for Linux hosts).

3.3 Experiment One: Cover in Computer Networks

3.3.1 Design

Cover within military geography has been defined as protection from observation and fire (United States. Dept. of the Army 1986). Previously, it was argued that firewalls are a type of cover in computer networks. One way to test the truth of this assertion is to see if firewalls empirically function the same as traditional physical cover. The most important functional characteristic of cover is that it acts to displace an attack. For example, if someone is shooting at a person, and the person hides behind a tree, the tree obscures the person (i.e., acts to protect from observation) and protects him or her from the bullets. In other words, if cover is used there should be 1) less threat to the covered object, and 2) more threat to the covering object.

If one is to test whether firewalls functionally behave in the same way as physical cover, one must test for this phenomenon of displaced attack. The experiment was arranged conceptually as follows: several computers were attached to the Internet on the same TCP/IP subnet and were observed for incidence of attack. The machines were roughly identical in that none offered any network services. One had the capability to act as a firewall and router, thus shielding other computers from attack. At the outset, however, the firewall and all other participating hosts were (logically, but not physically) equally distant from the router, as shown in Figure 3-A. Machines E1 through E3 are all participating hosts, while M1 through M3 represent non-participating hosts.




Figure 3-A: Logical Diagram of Cover Experiment at Outset




Figure 3-B: Logical Diagram of Cover Experiment at Conclusion

Over time, participating hosts E2 and E3 were placed under the protection of the firewall (E1). Figure 3-B (above) shows the shift by a change in the arrangement in the topology of the network. By the end of the experiment, all traffic for computers E2 and E3 was mediated by the firewall.

Throughout the experiment, rates of attack were logged. If the firewall functioned in the same way as cover does in real space, one would expect that the firewall should have exhibited a higher rate of attack as it protected more and more hosts. Conversely, one would expect that the client hosts should have exhibited a decline in frequency of attacks as they came under the protection of the firewall.

Figures 3-C and 3-D (below) describe this relationship. On the left, Figure 3-C shows the increase of attacks on the computer acting as the firewall. The relationship is graphed as a linear relationship between the independent variable

(Time) and the dependent variable (# of Attacks Expected). It is not expected that the relationship between the independent variable and the dependent variable is strictly linear. Instead, it is only predicted that the rate will increase over time. On the right, Figure 3-D shows a decrease in attack at some time t , at which it is placed under the protection of the firewall. From this point it is expected that the rate of attack will drop significantly.

In Figures 3-C and 3-D relationships were graphed for one independent variable (Time, given in days) and one dependent variable (# of Attacks Detected). These figures offer a simplified expected result for the experiment. Table 3-A (below) lists the actual independent and dependent variables for the experiment and their expected relationship.



Figure 3-C: Testing for Cover - Expected Attacks on Firewall

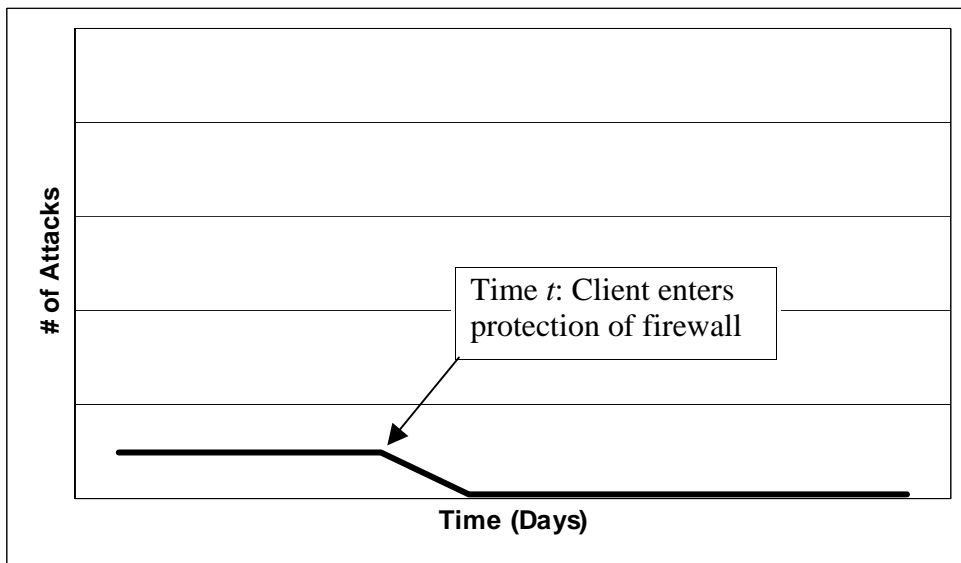


Figure 3-D: Testing for Cover – Expected Attacks on Client

HOST	INDEPENDENT VARIABLE	DEPENDENT VARIABLE	EXPECTED RELATIONSHIP
Firewall	# of Hosts Protected (Phase)	# of Attacks (Normalized to # of Days from start of Experiment)	Direct – Firewall should be attacked more often as it protects more hosts. The dependent variable must be normalized, since hosts did not come under protection at highly regular intervals.
	# of Days from Start of Experiment	# of Attacks	Direct – Since there is a direct relationship between (# of Days from Start of Experiment) and (# of Hosts Protected) there should also be a direct relationship between these. Use of this allows for a higher resolution graphing of attack data.
Non-Firewall	Status of Protection (Phase)	# of Attacks (Normalized to # of Days from Start of Experiment)	Inverse – Hosts should experience a decline in attacks once they are under the protection of the firewall. In this case the dependent variable must be normalized, since each host spent different amounts of time in protected and unprotected states.
	# of Days from Start of Experiment	# of Attacks	Inverse – Since hosts came under protection of the firewall as time increased, and following the reasoning directly above, it was expected that # of attacks should sharply decrease immediately after becoming protected. Other patterns (e.g., variability in attack prior to point-of-protection) could not be predicted.

Table 3-A: Independent and Dependent Variables for Cover Experiment and the Expected Relationships Between Them

3.3.2 Data Collection

Each computer involved in this experiment recorded its own data traffic. This was necessitated by the heavy data flows involved in normal network use (e.g., web browsing, file download, etc.) that would have quickly filled even large local hard disks if all traffic had been logged for analysis at a later time. Data flows were collected using winpcap, and analyzed in real-time with Snort.

The firewall host was set up with two Ethernet interfaces and Mandrake Linux version 9.0 was installed as the operating system. The only services enabled were Network Address Translation (NAT) services using *Shorewall*, a firewall and router software package distributed with the operating system. Shorewall was installed with no customization; all standard rule sets were used.

Two client hosts participated in the study, and were configured with winpcap, Snort, and FireDaemon as described in section 3.2.3.

None of the hosts had serving requirements (e.g., FTP, Telnet, WWW). These hosts were directly connected to the Internet using public IP addresses and were all on the same subnet at the beginning of the experiment.

Data was collected for a total of 194 days, from April 16, 2003 to October 26, 2003.

3.3.3 Method of Data Analysis

The record of attacks that Snort produces is point data – it provided a brief description of the time, date, type, and source IP Address of each individual attack.

Please see Appendix A for an example of a Snort incident record. In order to describe the relationship between the presence of cover and rates of attack, it was necessary to more fully articulate how “rate of attack” was measured. Date and time – as it appears in Snort’s log files – was used and a count of attacks was produced. Rates of attack were calculated for a) each day or b) each *phase*, an interval specified by a modification of the experiment. The important modification for the client hosts in the experiment was their placement behind the firewall, so a count of attacks was produced for the period while the host was unprotected and again for the period while the host was protected. These counts were also normalized by the number of days spent in each phase to allow for easier comparison between hosts. For instance, if a given host spent 10 days unprotected and 20 days protected, and experienced attack counts of 50 and 100 respectively, the normalized rate of attack for the host was the same (i.e., $50/10$ vs. $100/20$).

The important modifications for the firewall were each addition of a client host under its protection, so three phases were calculated: a) The period with zero hosts protected, b) the period with one host protected, and c) the period with two hosts protected. As with the client hosts, the rates of attack for each of these periods were also normalized to the number of days spent in each period.

Statistical measures were used to provide a description of the relationship between cover and attack. Regression analysis was applied to describe the relationship between variables (listed above in Table 3-A). Correlation was then calculated to determine how well the regression curve described the relationship

between variables. The hypotheses were considered supported if the correlation was positive for expected direct relationships and negative for expected indirect relationships.

3.4 *Experiment Two: Concealment in Computer Networks*

3.4.1 Design

Concealment was defined in military geography as protection from observation (Collins 1998). Previously it was argued that one way in which concealment exists in computer networks is the hiding of services on non-standard ports. This hypothesis can be tested by looking for a similarity between how concealment functions in real spaces and how concealment functions in computer networks. The principal function of concealment is to hide the target, thereby reducing its risk of being attacked. One can roughly infer the degree of success of the concealment by looking at the rates of attack on the target, assuming all other factors are equal.

In order to measure concealment of computer networks one can look at rates of attack on two computers in computer networks whose only significant difference is that they offer services on different ports. The difference in attack profile, if present, provides evidence for difference in visibility and concealment.

The experiment was designed based on four port scanning patterns which tend to make certain ports scanned more often than others (Scambray, McClure et al. 2001; Belcher and Yoran 2002). The most basic scanning pattern uses the ICMP (“Ping”)

protocol to detect whether a host is on and connected to the network. A slightly more detailed scan attempts to look for the most common ports that services are run on, including Port 80 (WWW service), Ports 137-139 (Windows File Sharing services), and those involved in this experiment: FTP on Port 21, SSH on Port 22, and Telnet on Port 23 (Belcher and Yoran 2002; Doctor Electron 2002). Many port scanners only scan all ports between 1 and 1024 by default (Scambray, McClure et al. 2001). Nearly all provide an option to scan all 65,536 ports.

Since scanning all ports is much more time intensive than just scanning the well-used ports, it is further hypothesized that hiding a service on a higher-numbered port is generally more effective concealment than hiding a service on a lower-numbered port. This effectively means that there is a range in the quality of concealment and that some concealment in computer networks is more effective than others. The hypothesis is considered correct if fewer attacks occur on the computer running services hidden with this “better” concealment than on a computer running the same services hidden with “worse” concealment.

In order to assess these differences in concealment, an experiment was set up to compare one computer connected to the Internet that ran its services on standard ports with another computer that ran its services first on a port in the middle range of port space (15001 to 15003) and then ran the same service in the lower range of port space (1001 to 1003). Table 3-B (below) lists the progression of the experiment where “Phase” shows the number of different configurations as the experiment progressed. “Days in Phase” shows the number of days spent in each phase.

“Server(S)” represents the server running its services on standard (S) ports.

“Server(NS)” represents the server running its services on non-standard (NS) ports.

Phase	Days in Phase	Server(S) Added Services and Ports	Server(NS) Added Services and Ports
1	10	ICMP & Default	ICMP & Default
2	6	ICMP & Default FTP Service on Port 21	ICMP & Default FTP Service on Port 15001
3	9	ICMP & Default FTP Service on Port 21	ICMP & Default FTP Service on Port 1001
4	8	ICMP & Default FTP Service on Port 21 SSH Service on Port 22	ICMP & Default FTP Service on Port 1001 SSH Service on Port 15002
5	13	ICMP & Default FTP Service on Port 21 SSH Service on Port 22	ICMP & Default FTP Service on Port 1001 SSH Service on Port 1002
6	8	ICMP & Default FTP Service on Port 21 SSH Service on Port 22 Telnet Service on Port 23	ICMP & Default SSH Service on Port 1001 FTP Service on Port 1002 Telnet Service on Port 15003
7	7	ICMP & Default FTP Service on Port 21 SSH Service on Port 22 Telnet Service on Port 23	ICMP & Default SSH Service on Port 1001 FTP Service on Port 1003 Telnet Service on Port 1003

Table 3-B: Progression of Concealment Experiment

Both machines began the experiment running two services enabled by a default installation of the operating system. These were: sunrpc on port 111 and kdm on port 1024. Additionally, both were responsive to ICMP “Pings” throughout the experiment. This was also a default condition of the installation of the operating system. After ten days of initial monitoring, Server (S) began to host FTP service on port 21, while Server (NS) began to host FTP first on port 15001, then on 1001.

After another fifteen day interval, Server (S) began to host the SSH service on port 22, while Server (NS) began to host SSH first on port 15002, then on 1002. Finally, after an interval of twenty-one days telnet services were enabled. Server (S) ran these services on the standard port 23, while Server (NS) ran telnet on port 15003 initially, and then on port 1003. Monitoring continued until the conclusion of the experiment, a period of fifteen days.

In the domain of computer network security each server is acting as a *honeypot*. Honeypots are computers that act as bait (and sometimes traps) for attackers. A honeypot is any system that is set up with the intended purpose of being attacked (Webopedia 2003). In some cases, this may be done to divert (and possibly catch) attackers without endangering important network resources (SecuritySearch 2004). Many times, as in this experiment, the attackers face no penalty for attempting to hack the system – their movements are merely recorded for further study (Webopedia 2003).

Expected results are graphically shown in Figures 3-E and 3-F (below). Figure 3-E shows the expected increase of attack on the host running services on standard ports. Figure 3-F also shows the expected increase in attacks on the host running services on non-standard ports, but at a much slower rate.

In Figures 3-E and 3-F (below) relationships were graphed for one independent variable (Time, given in weeks) and one dependent variable (# of Attacks Detected). These figures offer a simplified expected result for the experiment. Table 3-C

(below) lists the actual independent and dependent variables for the experiment and their expected relationship.



Figure 3-E: Expected Attacks on Server Running Services on Standard Ports (Server(S))




Figure 3-F: Expected Attacks on Server Running Services on Non-Standard Ports (Server(NS))

Hypothesis	Independent Variables	Dependent Variables	Expected Relationships
A – Attacks increase as services increase	# of Services	# of Attacks	An increase in # of services run should increase visibility of each host, and this should translate to an increase in attack and data traffic
		# of Packets	
B – The server running its services on standard ports will be attacked more often than the server running its services on non-standard ports.	Presence/Absence of Concealment	# of Attacks	Each of the independent variables shows a different resolution of the data from lowest resolution to highest resolution. It is expected that Server(S) should show more attacks and data traffic than Server(NS) for each of the independent variables.
	Phase	# of Packets	
	Time		
C – Upper ranges of Port Space offer better concealment than lower ranges of port space	Range of Port Space Used (Lower or Upper Range)	# of Attacks	It is expected that concealment in higher port ranges should result in fewer attacks and less data traffic.
		# of Packets	

Table 3-C: Independent and Dependent Variables for Concealment Experiment and the Expected Relationship Between Them

3.4.2 Data Collection

Unlike in the previous experiment, Snort was not installed on either of the two hosts involved directly in the experiment. In the previous experiment, data transfer rates were likely to be quite high. As a result, data monitoring had to be done in real time. Any attempt to log all the packets and inspect them later would have resulted in a large amount of disk space used. However in this experiment, only minimal traffic was anticipated. As a result, all traffic could be logged, and inspected later. This process had the added benefit of being able to more thoroughly inspect the traffic, since every packet sent to or from each host was saved.

The data logging in this case was performed by a third host. It utilized the traffic logging application Ethereal (Version 0.9.11), a Windows version of the tcpdump program used often in Unix and Linux for network diagnostics. Ethereal, in combination with winpcap, logged all traffic going to or coming from the two servers.

This third machine was connected to the same hub as the participating servers, but was not configured with an IP address. This minimal network installation allowed for all packets to be captured, but prevented nearly all types of Internet attack. Because of this configuration, it was very much safe from any intrusion, and could quite quickly and efficiently log data centrally. A graphical representation of the experiment is below in Figure 3-E, where again “Server(S)” represents the server

offering services on standard ports, “Server(NS)” represents the server offering services on non-standard ports, and “Monitor” represents the interface responsible for packet logging.

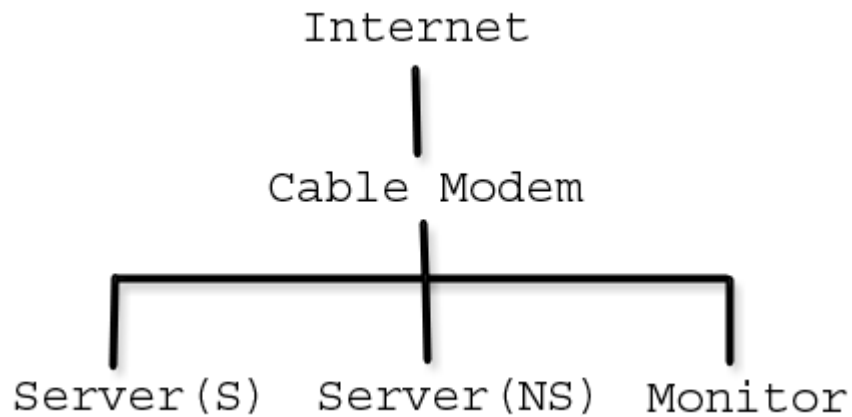


Figure 3-G: Logical Diagram of Concealment Experiment

At the conclusion of the experiment, the binary log files captured by the monitor with Ethereal were filtered to produce two groups of logs – one group for each computer rather than a single large log for both. These were processed with Snort to produce attack logs in the same way as the previous experiment. The text log file of attacks (alert.ids) was examined to produce rates of attack for each computer in the same way as described previously for the experiment involving cover. Since all traffic was logged for this experiment, it was also possible to produce a record of bandwidth as an indirect measure of intensity of attack. Each of these was also

normalized by dividing by the number of days-in-phase for each independent variable.

3.4.3 Method of Data Analysis

Statistical measures were used to provide a description of the relationship between cover and attack. Regression analysis was applied to describe the relationship between variables (listed above in Table 3-C). Correlation was then calculated to determine how well the regression curve described the relationship between variables. The hypotheses were considered supported if the correlation was positive for expected direct relationships and negative for expected indirect relationships.

3.5 *Experiment Three: Cover and Fortification in Geographic Spaces*

3.5.1 Design

In order to provide a similar quantitative benchmark for comparison, a statistical analysis was performed on using data published in Otterbein's (1970) classic study on warfare in primitive societies. Otterbein's method to produce data for analysis was to code fifty (primitive) societies based on twenty-six factors, among which were: type of cover, presence/absence of various fortifications, use of protection (shields or body armor), and success of the military organization. These factors were coded by Otterbein, based on ethnographic sources. While Otterbein looked for the relationships between these various factors and the degree of political sophistication, the interest here is in how cover, fortification, and protection impact military

operations. Of particular interest is a comparison of how various defensive measures impacted both the frequency of attack on a political groups, and how successful (militarily) that group was overall.

In order to determine how these defensive measures affected frequency of attack, factors were sorted into independent variables involving cover and fortification and dependent variables involving rate of attack and military success. These are listed in Table 3-D (below). The independent variables were chosen because they related strongly to cover and/or concealment. The dependent variables were chosen because they corresponded to either frequency of conflict or whether the cultural unit was considered a “military success.” Each of these independent and dependent variables and its relevance to the investigation into the military metaphor in computer networks is discussed in more detail in the following section on Data Collection.

Independent Variables	Dependent Variables
Degree of Cover (High or Low) Use of Field Fortifications Presence of Village Fortification	Frequency of War Frequency of Attack Military Success

Table 3-D: Independent and Dependent Variables Selected From Otterbein’s (1970) Factors

The importance of this experiment lies in its ability to quantify the effect of cover and fortification factors on rates of attack in geographic spaces. As a result, the discussion the military metaphor in computer networks benefits because one can readily compare the same factors in both computer networks and real space. Further,

unlike contemporary investigations into terrain, Otterbein coded his factors based on pre-modern societies. This allows for a more robust analysis of the effects of terrain in general, since the importance and use of terrain has changed over time.

3.5.2 Data Collection

The data used for analysis was taken directly from Otterbein's (1970) data. It is important to briefly describe how Otterbein selected the fifty societies that he used in his study. It is also necessary to explain what each of the selected seven variables are, and how each is relevant to the study.

The fifty societies represented in Otterbein's data were a stratified random sample chosen to represent 60 "culture areas" identified by the "Ethnographic Atlas," a standard reference within the field of anthropology. Random selections were discarded for any of four reasons mostly centering on accessibility of accurate information. As a result, the data set represents a wide spectrum of human primitive societies.

The independent variables chosen for examination as part of this analysis were selected on the basis of how well they related conceptually to ideas of cover and concealment. The coding followed Otterbein's own coding with two modifications. The first change was to sort the choices such that "more" of the given factor was coded with a greater value. In this way the coding followed a more intuitive understanding so that a "better" defense (for example) was assigned a higher number. Second, factors that were originally coded with more than two options were compressed such that only two options were coded. These modifications and the

justification for them are provided below. This compression was generally done to allow for more reliable statistical testing since there were so few cases.

Three factors were tested as independent variables, and all were related to defense. The first factor (Cover) was most directly relevant to this thesis, as it specifically coded for the geographic cover provided in the area of the society. Poor cover was coded as “0” and good cover was coded as “1”. The second factor (Field Fortification) was chosen because it represents one type of geographic-scale cover. Societies that employed no field fortifications were coded “0” and societies that did employ field fortifications were coded “1”. The third factor (Village Fortification) represented the other type of geographic-scale cover. Otterbein originally distinguished between those societies that were located in a defensive position and those that employed defensive structures, but in order to provide more intelligible and direct results, these have been compressed. As such, societies employing neither defensive positions nor fortifications were coded as “0” and societies employing either of these were coded as “1”.

There were three factors chosen as dependent variables. The first, “Frequency of War”, was compressed such that societies either “never” or “infrequently” at war were coded “0” and societies at war “frequently” or “continually” were coded “1”. The second, “Frequency of Attack,” was coded exactly the same so that societies “never” or “infrequently” attacked were coded “0” and societies attacked frequently or more were coded “1”. The third and final factor, “Military Success” was defined by the territorial expansion or contraction of the cultural unit. Those with

contracting borders were considered less militarily successful by Otterbein, and were coded with a “0”. Those societies with stable or expanding borders were considered more militarily successful and were coded as “1”. Table 3-E (below) summarizes this coding scheme.

Type	Variable Name	Coding Scheme
Independent Variables	Cover	1 – Poor Cover
		2 – Good Cover
Dependent Variables	Field Fortifications	1 – No
		2 – Yes
	Village Fortifications	1 – No
		2 – Yes, or Defensive Site
	War Frequency	1 – Never or Infrequent
		2 – Frequent or Continuous
	Attack Frequency	1 – Never or Infrequent
		2 – Frequent or Continuous
Military Success	1 – No (Shrinking Borders)	
	2 – Yes (Stable or Expanding Borders)	

Table 3-E: Recoding Of Otterbein’s Variables

3.5.3 Method of Data Analysis

The method of analysis also followed in a similar manner as Otterbein. In order to test for significant statistical relationships between factors, both the chi-square statistic (χ^2) and the phi coefficient (Φ) were used. The chi square statistic attempts to determine independence of variables. The phi coefficient is a correlation measure used to describe relationships between two dichotomous variables. Although the variables were coded in such a way that the numeric representation roughly describes

the magnitude of the variable, there is no clear “breakpoint” between variables – they remain highly tied to the subjective assessment of the original researchers. As it happens, since there are only two values for each variable, the calculated phi coefficient will be the same as other measures of correlation, including Pearson’s r . Because of the imperfect coding, and following Otterbein’s own method, the breakpoint for statistical significance of the p-value for the χ^2 test is less than .10 (i.e., there is less than a 10 percent chance that the results were due to chance).

3.6 Summary

The methods of this section provide a means for detecting how cover and concealment actually affect attack patterns in computer networks, and for providing quantitative descriptions of how cover and fortification behave in real space as a means for comparison. The results of these experiments are presented in the next chapter.

4 RESULTS AND DISCUSSION

4.1 Introduction

This chapter details the findings of the experiments conducted to test and support the theory that some features of computer network defense can reasonably be treated as either cases of, or metaphors for, spatial cover and concealment. Of particular interest is an assessment of whether changes in concealment and cover within computer networks impacted the rate of attack on servers and clients. The results offer support for the idea that cover and concealment manifest in very similar ways in both real geographic space and within computer networks. Overall, however, the results indicate that in some ways the hypotheses of both cover and concealment were too simplistic, and that the relationships between independent and dependent variables are more complex than was anticipated.

The results for each experiment are presented below sequentially, followed by an overall discussion that concludes this chapter.

4.2 Cover in Computer Networks

4.2.1 General Results

There were two episodes during which the firewall host logged no data record of attack. These occurred between days 45 and 74 and again between days 91 and 159.

It should be emphasized that this experiment represents only one case out of many instances of firewalls on the Internet. The results should not, therefore, be interpreted to be representative of the larger population of all Internet servers. The results are only suggestive – not conclusive - that the hypotheses are true.

There were also three instances of outlier events. One was a data collection problem where the intrusion detection system recorded a series of 1566 bad packets during the course of one day (Day 58) on Client2. Further investigation revealed that these bad packets were almost certainly legitimate traffic rather than an attack. As a result, these attack events were removed from the attack log, and only the legitimate attacks remained in the data record for that day.

Two other outlier events were recorded. Both of these were days during which an abnormally large number of attacks were recorded for one of the computers involved in the study. On Day 44 there were 193 attacks on the Firewall, and on Day 56 there were 179 attacks recorded on Client2. Investigation of the data revealed that these were legitimate attacks rather than an error, and so they remained in the data set.

After gathering the data it was noted that there were a large number of reports generated by a worm that attacks Microsoft's SQL database port. These were recorded as "MS-SQL Worm Propagation Attempts" and were caused by the continued spread of the Sapphire worm on the Internet (Danyliw 2003). Worms are automated attacks that generally attack random IP addresses, looking for particular weaknesses. Since the research was interested in "active" attacks rather than the

random attacks generated by worms or viruses, the results were filtered into “Raw” and “Active” attacks. “Raw Attacks” represents all attacks recorded by Snort. “Active Attacks” represents all attacks recorded by Snort *except* MS-SQL Worm propagation attempts. It may be the case that other attacks recorded by Snort were random attacks, but this filtering eliminated only those that were almost surely randomly targeted.

Evidence compiled during the study supported the hypothesis that the firewall computer would be attacked more often as it protected an increasing number of clients. Two analysis methods are used to provide support for the hypothesis. One was to see how the *number of clients protected* influences *rate of attack*, and is described by the term *in phase*. The other was to see how the *day* (as a higher-resolution surrogate of *number of clients protected*) influences *rate of attack*.

4.2.2 Attacks on Firewall by Phase

There were three relevant phases in the progression of this experiment. These are referred to as phases 0, 1, and 2, each corresponding to the number of clients protected by the firewall. In Phase 0, no clients were supported, in Phase 1 Client1 only was protected, and in Phase 2, Client1 and Client2 were both protected. Phases were defined as periods between days where all three computers were able to gather data. Phase 1 lasted between Days 1 and 44. Phase 2 lasted between days 76 and 90. Phase 3 lasted between days 160 and 194.

Attacks were summed for each phase and divided by the total number of days spent in phase to produce a rate of attack during the phase. Each phase has two rates

of attack – one for Raw Attacks and one for Active Attacks. Finally, the number of clients protected in each phase and rate of attack during phase were correlated to determine whether there was a positive correlation between the two. Table 4-A summarizes these results. In both cases (Raw and Active Attacks) a strong positive correlation (0.77) was found between number of clients protected and # of Attacks recorded in phase thus indicating support for the hypothesis.

# Of Clients Protected (Phase)	Days in Phase	# Of Raw Attacks During Phase	Raw Rate of Attack During Phase	# Of Active Attacks During Phase	Active Rate of Attack During Phase
0	44	508	11.5	362	8.2
1	26	98	3.8	28	1.1
2	35	1710	48.9	1505	43

Table 4-A: Raw and Active Attacks on Firewall by Phase

4.2.3 Attacks on Clients by Phase

While attacks against the firewall were expected to increase as it protected an increasing number of clients, the corollary hypothesis was that attacks on the clients would decrease as they were placed behind the firewall. This element of the greater hypothesis of cover in computer networks was relatively trivial since it was almost certain that it would be the case. If it were not, there would be no reason at all to employ firewalls. As predicted, attacks on clients decreased markedly as they were moved behind the firewall. A closer examination of the data illustrates the dramatic

shift in attack pattern resulting from the change in location. The following information is summarized in Table 4-B.

Client1 spent 72 days unprotected (Phase 0) and 122 days protected by the firewall (Phase 1). During Phase 0 it was attacked a total of 499 times, of which 235 were active attacks. During Phase 1, no attacks were recorded on Client1 at all. The rate of attack (# of Attacks in Phase / # of Days in Phase) in Phase 0 was 6.9 for Raw Attacks and 3.3 for Active Attacks. The rate of attack for both Raw and Active Attacks in Phase 1 was 0. Since attacks decreased in Phase 1, and there were only two observation points, there was a perfect inverse relationship (correlation -1.00) between protection and both types of rate of attacks in phase. Thus, the data supports the hypothesis for Client1.

Client2 spent 99 days unprotected (Phase 0) and 95 days protected (Phase 1) by the firewall. During Phase 0 it was attacked a total of 998 times, of which 596 were active attacks. During Phase 1, 5 attacks were recorded on Client2, all of which were active attacks. The rate of attack in Phase 0 was 10.1 for Raw Attacks and 6.0 for Active Attacks. The rate of attack in Phase 1 was 0.05 for both Raw and Active Attacks. As with the Client1, Client2 showed a perfect inverse correlation (-1.00) between amount of protection and both types of rate of attacks in phase. Thus, the data supports the hypothesis for Client2.

Client	Phase	Length	# Raw Attacks	# Active Attacks	Raw Rate of Attack	Active Rate of Attack
1	0	72	499	235	6.9	3.3
	1	122	0	0	0	0
2	0	99	998	596	10.1	6.0
	1	95	5	5	0.05	0.05

Table 4-B: Raw and Active Attacks on Clients by Phase

4.2.4 Attacks on Firewall and Clients by Day

Another way to look at the data is to see how many attacks there were on each host for each day. This method yields a higher resolution visualization of the data. Figure 4-A shows the daily raw attack rate and Figure 4-B shows the daily active attack rate on all three computers in the study. Figures 4-C and 4-D show the respective linear regressions and correlations for the data shown in Figures 4-A and 4-B. The reader should note that the scale has been broken on Figures 4-A and 4-B to allow better resolution of the data for most of its range (0-100). Further, the scale has been extended below zero on Figures 4A-4D to allow for better visualization of the attack patterns for Clients 1 and 2, both of which had attack rates equal to zero for a significant portion of the experiment.




Figure 4-A: Raw Attacks on Hosts by Day




Figure 4-B: Active Attacks on Hosts by Day




Figure 4-C: Linear Regression Plot of Raw Attacks by Day




Figure 4-D: Linear Regression Plot of Active Attacks by Day

The graphs showing daily attack rate are characterized by high daily variability for each of the hosts. The lack of data for firewall shows up in the large discontinuity in the time-series. When data was again gathered on Day 160, there was clearly a higher level of attack than observed previously. It is only from this point on that data was available for all three computers while both clients were behind the firewall.

Linear regression analysis using the least-squares method shows the overall trend. The results are nearly identical for both Raw and Active Attacks. In both cases, Client1 and Client2 show a linear relationship with a decreasing slope, with Client2 having a higher Y-intercept. This is almost certainly a result of the later date at which it entered protection relative to Client1. The positive slope for the regression line of the firewall is due to the higher attack rate observed towards the end of the experiment.

Table 4-C shows summary data for the regression lines, R^2 values, and correlations between Day and (Raw and Active) Attacks on hosts. All correlations were positive for the Firewall and negative for Client1 and Client2. The data thus supports the hypothesis that attacks on the firewall would increase and attacks on clients would decrease as the location of clients was shifted from next to behind the firewall.

Host	Attack Type	Linear Regression Eq.	R ² Value	Correlation
Client1	Raw	$y = -0.0728x + 12.265$	0.0834	-.57
	Active	$y = -0.0269x + 3.8384$	0.1605	-.40
Client2	Raw	$y = -0.0528x + 7.7519$	0.3255	-.28
	Active	$y = -0.0452x + 7.5023$	0.0358	-.19
Firewall	Raw	$y = 0.2494x + 2.0526$	0.3645	.60
	Active	$y = 0.2315x - 0.7787$	0.3284	.57

Table 4-C: Regression and Correlation results for Cover Experiment Hosts

4.2.5 Analysis and Summary of Cover Experiment

The results from the experiment on cover all supported the hypothesis that a firewall would be attacked more often as it protected a greater number of hosts. Correlations and linear regressions all showed this hypothesized direct relationship. Similarly, all correlations and linear regressions showed the expected inverse relationship for clients between protection and frequency of attack.

An important caveat stems from the highly irregular data record. The periods during which no data could be gathered, particularly the second, preclude general continuity in the data set that would lend credibility to its analysis. The second break in data gathering was crucial in that it occurred during the back-to-school period of September and October. It seems likely that at least part of the observed increase in attack rates was due to an increased number of attackers in the field. This confounding of factors could be alleviated by an experiment conducted over a longer period of time (in which no breaks in data gathering occurred), or possibly one in which this factor was explicitly controlled for.

4.3 Concealment in Computer Networks

Two independent variables were posited to be examples of concealment within computer networks. The first of these corresponded to visibility - the inverse of concealment – and described the number of services offered by a server. Both servers increased the number of services offered, and both were expected to experience an increase in the number of attacks. The second variable tested for was more strongly related to concealment itself. One server ran its services on standard ports, while the other “hid” the same services in the upper ranges of port space. It was expected that as a result, the computer running its services on standard ports would experience a higher rate of attack than the server that ran its services on non-standard ports.

4.3.1 General Results

Two computers gathered data for the experiment: Server(S), which ran its services on standard ports, and Server(NS), which ran its services on non-standard ports. As in the previous experiment, both servers recorded a very large number of MS-SQL Worm propagation attempts. Of 538 total (Raw) attacks observed, only 221 were not MS-SQL Worm propagation attempts. As a result only these Active Attacks were used in analysis of the data.

Additionally, since all traffic was recorded, bandwidth was used as an additional surrogate measure of attack. This was justified since the servers had no sanctioned users – therefore any interaction was unsolicited. Some of these “unsolicited”

packets were, in fact, legitimate data exchanges for IP address acquisition. However, since these exchanges were relatively rare, they were left as part of the data record.

Further investigation of the traffic records showed that a large amount of data traffic was sent to port 6346, a port commonly associated with the file sharing network *gnutella*. This most likely occurred as a result of the dynamic IP-addressing system used by the local ISP. A previous user of the IP address reassigned to one of the computers involved in this experiment was likely using a *gnutella* client. When the IP address was reassigned, *gnutella*-bound traffic went instead to one of the servers connected with the experiment. In total, 15546 packets between port 6346 were filtered from the data. Additionally, there were three other data flows not connected with attack that were mistakenly logged into the data record. These included contact of two webpages and the traffic associated with the installation of the *proftp* server.

A comparison between bandwidth and raw attacks showed a correlation of 0.34 for Server(S) and -0.003 for Server(NS). Figures 4-E and 4-F (below) show raw attacks and bandwidth usage for both servers. It is evident from visual inspection that during some periods attack and bandwidth are highly correlated, while in others they are not. This is generally due to the fact that Snort is designed to sift through all traffic to detect intrusion – as a result, some of it is classified as legitimate traffic. There were, for instance, several failed attempts to use the FTP server on Server(S). However, since this could quite plausibly be a legitimate user on any given system, Snort did not record these attempts as attacks. The result is that analysis of attack by

bandwidth is likely to provide quite a different perspective than that offered by analysis of attack patterns recorded by Snort alone.

Also evident from figures 4-E and 4-F is the interruption in data from Day 53 to Day 76. During this period both servers were off, stemming from a power supply failure in Server(S). Once Server(S) was repaired, both computers were turned on and began gathering data. Since both computers spent nearly identical amounts of time on during the study, there should be no bias in the data stemming from the hardware failure observed with Server(S).

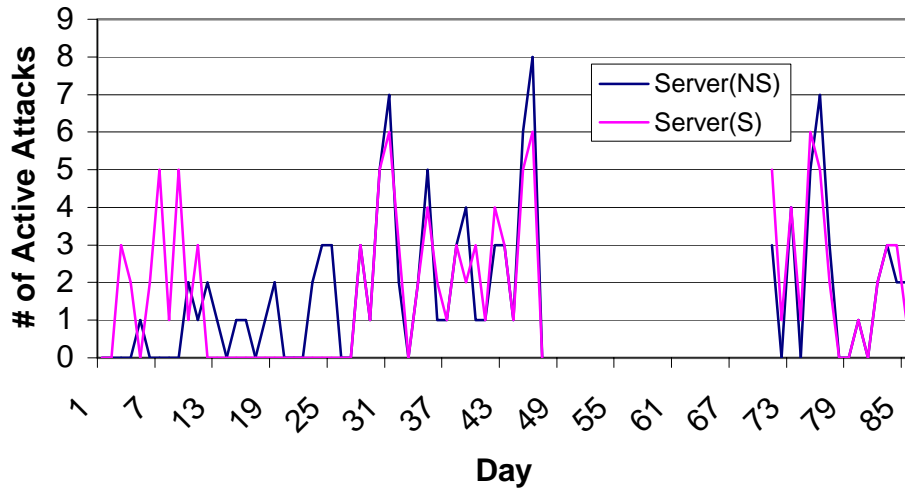


Figure 4-E: Active Attacks by Day on Concealment Experiment Hosts




Figure 4-F: Packets Transmitted by Day on Concealment Experiment Hosts

In addition to looking at daily rates of attack, attack patterns were examined with respect to *phase* of the experiment, where a phase represents some modification of the servers. These phases were described in detail in the previous chapter, and are summarized in Table 3-B.

Finally, it should be restated that the data and analysis stemming from this experiment should be treated as a case rather than a statistical sample. As a result, generalizations should be treated with caution, and the data should be taken to support, and not prove, the hypothesis.

4.3.2 Impact of Number of Services Offered

A preliminary hypothesis of this work on concealment was that attacks on both servers would increase over the course of the experiment as a result of each offering an increasing number of services. Since each service is assigned a unique port, an increase in services increases the number of open ports – apertures into the system.

Since this was only a preliminary hypothesis – not directly connected to the work on concealment – the experiment was not set up with answering this question in mind. Both servers had an increasing number of services added at the same time. As a result, there is no control group with which to compare, and so the results only answer this question on a very limited and suggestive basis.

Table 4-D (below) summarizes the data in response to this hypothesis. Column (B) shows the number of services for each server, with columns (C) and (D) showing the attack rates during each of these periods. The Active Attack Rate shown in column (C) was calculated by summing the number of active attacks observed during

the period and dividing by the total number of days spent in that period. The Packet Transmission Rate shown in column (D) was calculated by summing the total number of packets transmitted to or from each server during the period and dividing by the total number of days spent in that period.

The correlations shown in column (E) were between the number of services and the Active Attack Rates for that period for each server. The correlations shown in column (F) were between the number of services and the Packet Transmission Rate for that period for each server. All correlations were positive, showing support for the hypothesis that as services offered increase one can expect attacks to increase. Conclusive evidence for this hypothesis could be provided by observing a larger number of cases in conjunction with the employment of control groups.

(A) Server	(B) # of Services	(C) Active Attack Rate	(D) Packet Trans. Rate	(E) R ² Correlation (B with C)	(F) R ² Correlation (B with D)
Server(S)	0	1.9	128.5	0.372	0.215
	1	0.2	293.8		
	2	1.1	224.8		
	3	2.6	189.9		
Server(NS)	0	0.4	80.8	0.993	0.562
	1	1.1	200.6		
	2	1.7	200.1		
	3	2.7	162.8		

Table 4-D: Active Attack Rates and Correlations

4.3.3 Use of Concealment

The principal hypothesis connected to concealment within computer networks was a prediction that the hiding of services on non-standard ports would result in fewer attacks. To put it in more familiar geographic terms, the question becomes: *does the location of services impact the rate of attack?* There are three ways in which this question was addressed. First, attack totals are presented for each server for the entire study period. Second, attacks are charted by day to show a high-resolution view of the data. Finally, attacks are shown by phase, as each modification was made to the servers.

Table 4-E (below) shows the cumulative results for the study. The results at this level generally support the hypothesis, with Server(S) showing a higher number of both Active Attacks and Packets Transmitted.

Server	# of Active Attacks	# Packets Transmitted
Server(S)	111	12981
Server(NS)	110	10562

Table 4-E: Attack and Data Traffic Data for Concealment Experiment

Figures 4-E and 4-F (below) show daily rates of Active Attack and Packets Transmitted, respectively. The correlation of Active Attacks with Day for Server(S) was 0.23, and 0.29 for Server(NS). The correlation of Packets Transmitted with Day for Server(S) was 0.002, and 0.30 for Server(NS). These results do not support the

hypothesis, since attacks broadly increased on Server(NS) through the progression of the experiment, while attacks on Server(S) tended to increase at a slower rate.

It is worth noting that during some periods there were strong correlations in both Active Attacks and in Packets Transmitted between both servers. From Day 26 to 47, for instance, the correlation between Server(S) and Server(NS) was 0.93. The correlation for Packets Transmitted between the two during the same period was 0.86. This correspondence is suggestive that some other factor or set of factors may play a strong role in determining attack patterns.

Another way to look at the data is by examining attack rates during phases of the experiment. This provides more direct relationship than looking at daily attack rates since it controls for differences in phase length by normalizing attack counts to days spent in phase. Figures 4-G and 4-H show graphs for Active Attack Rates (# of attacks during phase / number of days spent in phase) and Packet Transmission Rates (# of packets transmitted to or from hosts / number of days spent in phase) for the seven phases of the study.

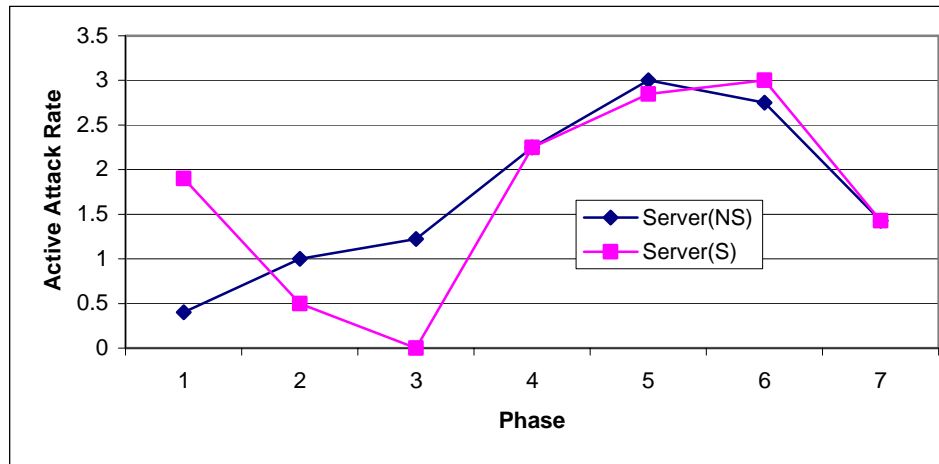


Figure 4-G: Active Attacks by Phase on Concealment Experiment Hosts




Figure 4-H: Packets Transmitted by Phase on Concealment Experiment Hosts

The pattern of active attack is quite similar for both servers. Active attacks initially drop for Server(S), while they initially steadily increase on Server(NS). Both, however, exhibit an absolute maximum somewhere between phase 5 and 6, before dropping somewhat during phase 7. A linear regression plotted for both servers shown below in Table 4-I illustrates more precisely the close correspondence

between Server(S) and Server(NS). Active attacks were moderately correlated with phase for both servers (0.44 for Server(S) and 0.67 for Server(NS)).




Figure 4-I: Linear Regression Plot of Active Attacks by Phase on Concealment Experiment Hosts

When attacks are measured by the unsolicited transmission of packets, in all but one of the 7 phases, Server(S) had more packets transmitted to and from it than did Server(NS). This information tends to support the hypothesis. It should be noted, however, that a linear regression of Packets Transmitted by Phase shows a much stronger positive relationship between Server(NS) and Phase than between Server(S) and Phase (Figure 4-J, below) . From this perspective, attacks actually drop off over time on Server(S) and increase over time for Server(NS).

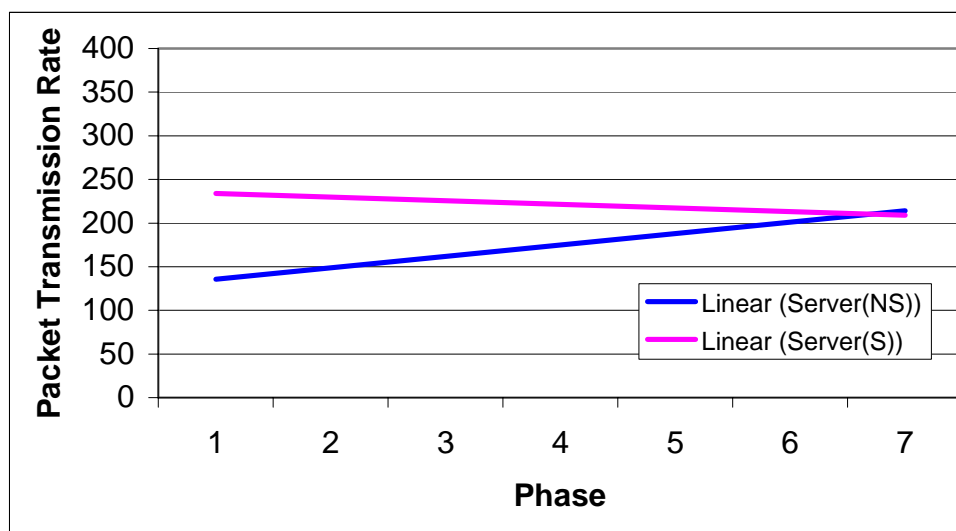


Figure 4-J: Linear Regression Plot of Packets Transmitted by Phase on Concealment Experiment Hosts

This observed increase in packet transmission may be due, in part, to a tendency of the attackers to expend more effort searching for open ports on servers utilizing a concealment strategy. Since the host would respond to ICMP Pings, attackers could easily determine the presence of a host, even if they could not locate viable services. Since scanning (searching for services) is construed as an attack by Snort, stronger concealment led to increased attacks because it promoted this searching behavior.

4.3.4 Quality of Concealment

This experiment provides information as to whether or not there is such a thing as different gradations or qualities of concealment. There was no indication that attack patterns on the computers as a whole varied according to whether services were run on high or low numbered ports. Interaction with ports 1001-1003 (1 packet

received) and ports 15001-15003 (0 packets received) were overall fairly comparable as neither showed heavy interaction.

A look at scanned patterns as a whole shows a slightly different story, consistent with the hypothesized differences in port range interaction. Figures 4-K through 4-P (below) show plots of port contact and degree of interaction (i.e., number of unsolicited inbound packets) for the three ranges of port space. Figures 4-K and 4-L show the full range of port space, and a clustering of contact is shown roughly in the range of ports between 1 and 10,000. Figures 4-M and 4-N show a closer view of this range, with both servers showing a clustering of contact between 1 and 2,000, and Server(S) showing a high degree of contact between roughly between ports 4,250 and 6,250. Figures 4-O and 4-P show port contacts between 1 and 2,000 for both servers. It is apparent from this view that Server(S) experienced both more contacts (i.e., number of different ports contacted) and more intense probing at those contacts, in general.

This data tends to support the hypothesis that in general, higher numbered ports offer greater concealment. The difference in quality of concealment does not arise from any inherent properties in the ports themselves, but from a difference in the interaction patterns between ranges in port space. This high degree of interaction seems to create something similar to a path in real space, where certain segments of the space are traversed more often to reach familiar and common destinations, but other segments of the space are left untrodden. Just as in real space, it makes sense

to attempt to “hide” things in these spaces where less interaction commonly takes place.




Figure 4-K: Distribution of All Ports Contacted – Server(S)




Figure 4-L: Distribution of All Ports Contacted - Server(NS)

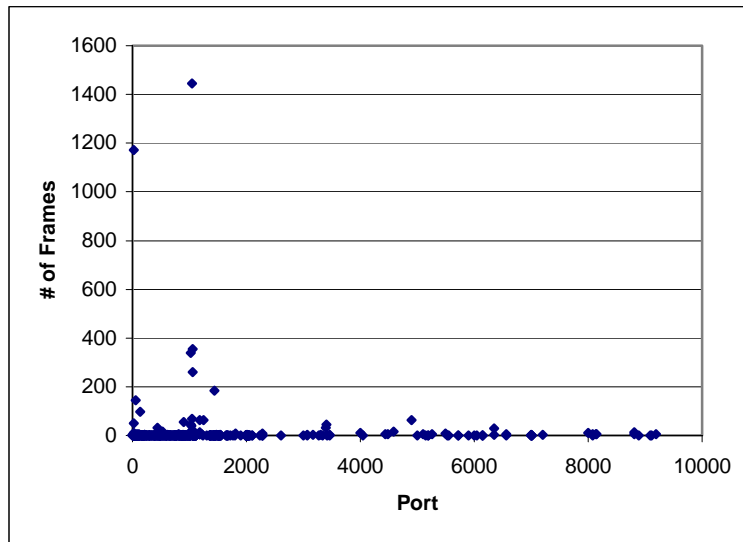


Figure 4-M: Distribution of Ports Contacted (1-10,000) - Server(S)




Figure 4-N: Distribution of Ports Contacted (1-10,000) - Server(NS)




Figure 4-O: Distribution of Ports Contacted (1-2,000) - Server(S)

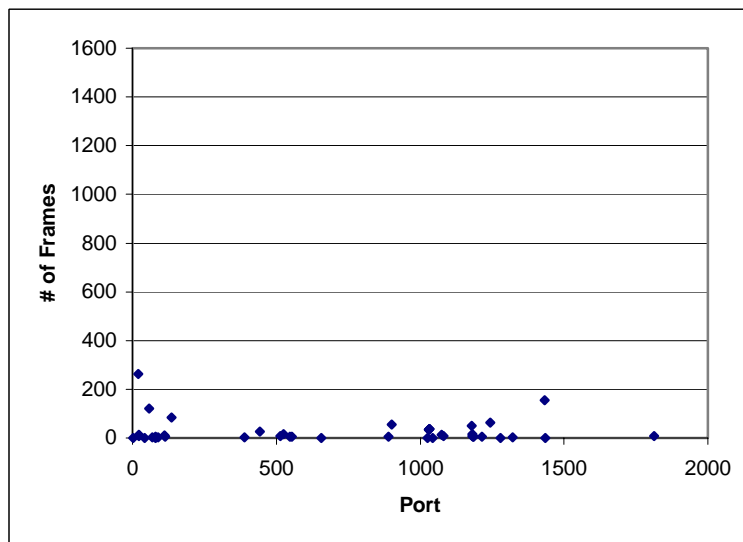


Figure 4-P: Distribution of Ports Contacted (1-2,000) - Server(NS)

4.3.5 Searching Behavior

One of the ways in which to interpret some of the findings of this experiment is to characterize some of the interaction as searching behavior. Many of the intrusions that Snort logged are actually port scans. Port scans are not generally harmful in themselves, but represent (generally unwanted) reconnaissance of targets. The hypothesis regarding concealment predicted that a computer that utilized concealment would be attacked less often. The two cases observed here suggest that this is likely to be true. However, active attacks on the two servers over the course of the experiment were very nearly equal. Additionally, attacks increased on the server running services on non-standard ports at a faster rate than the server running its services on standard ports, even though they were lower in absolute numbers during six of the seven phases of the study.

The original formulation of the hypothesis did not include one other aspect of concealment in real space during the quest for the analogues in computer networks – namely that searching behavior tends to increase as targets become harder to find. When searching is interpreted as attack – as it is with Snort and many other Intrusion Detection Systems (IDS) – it in large part counter-acts any reduction in more legitimate attacks.

One approach to separating searching from legitimate attacks is simply to exclude them from the attack record. There are two reasons why this is not generally advisable. Searching is widely considered an attack because it *is* a danger – it

represents possible future incursions, but it may even be the only clue that a systems administrator receives that he or she is under attack, since new and unique attacks may not be detected by an IDS at all. In the case of this experiment, the overwhelming number of attacks were either scan attempts of various kinds and extents, or MS-SQL Worm propagation attempts. When these two main types are filtered from the Snort attack record, only four attacks remain on each server. It seems that a longer observation period would be required in order for this method to be useful in lending support for or against the hypothesis.

One way in which to look at “real” attacks in the collected data is to look at interaction on ports used in this study (i.e., 21-23, 1001-1003, 15001-15003). Table 4-F summarizes the packets transmitted to these ports. Although port 21 was probed on Server(NS), there were over four times as many packets transmitted to Server(S). Many of these packets are accounted for by authorized logon attempts. There were sixty unauthorized FTP logon attempts on Server(S) compared with none on Server(NS). Clearly hiding FTP services helps obviate attacks on the FTP service itself.

It appears, however, that hiding services may also help to curb unwanted interaction as a whole. Table 4-G summarizes other measures of searching behavior. Server(S) exhibited a wider range of ports scanned, a higher number of ports scanned, and more ping probes than Server(NS). From this perspective then, Server(S) not only experienced more actual attacks, but also more intense searching.

This suggests that attackers search more intensely as their searches increase in effectiveness. In other words, the more services they find, the harder they look.

Port	Packets Transmitted	
	Server(S)	Server(NS)
21	1172	264
22	50	7
23	11	12
1001	0	0
1002	0	0
1003	1	0
15001	0	0
15002	0	0
15003	0	0

Table 4-F: Packets Transmitted To/From Ports Used in Concealment Experiment

Interaction	Server(S)	Server(NS)
Number of Pings	925	816
# of Different Ports Contacted	494	92
Lowest Port Contacted	2	2
Highest Port Contacted	65301	63000
Unauthorized FTP Logon Attempts	60	0

Table 4-G: Other Measures of Searching Behavior on Concealment Experiment Hosts

4.3.6 Summary of Concealment Findings

Although the data for this experiment were quite complete in comparison with the previous experiment, the results are somewhat more ambiguous. The data record generally supports the hypothesis: there were more attacks against the server running its services on standard rather than non-standard ports. The predicted relationship

held true for various measures of attack: there were more recorded attacks, packets transmitted, and unauthorized logon attempts on Server(S) than on Server(NS). Additionally, it was found that searching behavior on Server(S) may have increased as a result of its higher visibility.

4.4 Cover and Fortification in Geographic Spaces

In order to provide a means of quantitative comparison between real space and computer networks, geographic variables relating to the surrounding terrain and the military activity of fifty primitive societies found in Otterbein's (1970) classic study on warfare were analyzed for comparable relationships. The evidence for statistically significant and strong relationships is, however, scant. Only one chi-squared analysis of factors (out of nine) had a P-value of less than 0.10. The independent variable "Cover" and the dependent variable "Frequency of Attack" failed the chi-squared test for independence (indicating a non-random relationship) in the direction expected. The correlation was negative and of moderate strength ($\Phi = -0.33$). Use of natural cover then, tended to correspond with attacks on that society occurring infrequently or never more than chance would suggest. Although no causal relationship can be determined from the data, it is at least plausible that utilization of cover created an atmosphere where strong defense tended to discourage attack. Table 4-H gives the P-values for each of the nine pairings, with the pairing of "Cover" and "Frequency of Attack" shown in bold for emphasis. Table 4-I (below) gives the contingency table for that pairing.

	Frequency of War	Frequency of Attack	Military Success
Cover	0.73	0.06	0.27
Field Fortification	0.32	0.24	0.24
Village Fortification	0.69	0.43	0.39

Table 4-H: Chi-squared Test P-Values for Cover and Fortification in Otterbein's (1970) Data

		How often is the society attacked?	
		Never or Infrequently	Frequently or Continually
What type of cover is offered by the terrain?	Poor Cover	6	12
	Good Cover	18	11

Table 4-I: Contingency Table for Cover vs Frequency of Attack

As a whole, this data generally shows independence for the variables tested. Use of geographic cover did not widely correspond to lower rates of attack or warfare or to military success. In the test related most directly related to two of the variables under investigation in computer networks (i.e., cover and rates of attack), a statistically significant moderate correlation was found.

4.5 Discussion

This thesis poses three research questions in order to explore how the military metaphor might apply to computer network defense. These questions are:

- What are the network-based topologies that humans have used as part of their defenses in military affairs?
- How have computer network defenses been organized spatially?
- How do specific uses of cover and concealment compare between real space and computer networks?

The first two questions are largely historical. The third is answered by development of a system of spatial metaphors that transfer the military metaphor from real space to computer networks. The empirical results presented in this chapter provide an opportunity to directly test elements of the system.

The empirical work shows how the use and disuse of cover and concealment affects attack rates on servers, and whether these patterns resemble those we would expect to find with analogous features in real space. The results of the first two experiments were consistent with expectations. Hiding services on higher-numbered ports (a suspected equivalent of concealment) did tend to reduce attack rates. Additionally, a firewall was observed to not only reduce attack rates on its clients but to actually displace attacks onto itself – a defining trait of an object behaving as cover. The third experiment evaluated Otterbein's (1970) data on warfare in primitive societies to see if the expected relationships between spatial defense features and attack rates could be reliably found in physical space. Although a statistically reliable ($P < .10$) moderate correlation ($\Phi = -0.33$) in the expected

direction was found between use of cover and frequency of attack, all other eight variable combinations (representing other types of defense factors and attack rates) were shown to be statistically independent.

Taken as a whole, two lessons can be drawn from these results. First, there is functional similarity between computer network components and physical defense systems. This is a highly significant result, because it provides evidence that the military metaphor is an *appropriate* one as applied to computer networks. Taken in combination with the spatial similarities discussed in chapter two, we can assert with confidence that given the strong similarity both in structure and in behavior between defensive features in real space and computer networks, the military metaphor can be justifiably applied. Essentially, these results show that the concepts of cover and concealment map well from real space to computer networks.

However, these results do not directly establish the *usefulness* of thinking of network defenses in terms of real-world military defenses. Such usefulness is hinted at by the frequency with which terms from military discourse are already applied (and occasionally misapplied) to network defense, but could be empirically tested for by incorporating the military metaphor more fully into visualizations of network security information (as in Fisk, Smith, *et al.* 2003) and testing whether human subjects find them helpful in understanding network security problems.

The second lesson that can be drawn from the results is that if expected relationships between spatial defenses and attack rates do not hold up in physical space, one must assume that the relationships are more complex and much less

deterministic than they seem. Although the expected relationships were found with cover and concealment in computer networks, they may not be found in other network security tests simply because of confounding variables. Additionally, the classification of variables in Otterbein's (1970) study may have occurred at too wide a scale to be applicable to specific circumstances; attack rates and spatial defense variables were classified for the entire span of a civilization's existence. This, in contrast to experiments with a duration of only a few weeks does create problems of comparison. Ideally, information about the effect of spatial defenses on attack rates should be measured in real space with an eye toward keeping variables of time and scope as commensurate to network security measurements as possible.

5 CONCLUSIONS AND OUTLOOK

5.1 *Conclusions*

The goal of this thesis is to diagram a comparison between the two very separate worlds of computer network security and military geography. Although computer network security techniques, technologies, and language all to some extent borrow from concepts of physical security, a rigorous exposition of their shared spatial commonalities may be of some use in each of the communities. As governments increasingly attempt to project their power into cyberspace and continue to maintain an overwhelming military presence in real spaces, a synthesized concept of security that overlaps both elements is crucial to developing better strategies that can be employed against a wide variety of threats.

This thesis represents initial work in such a direction. Specifically, a comparison between military geography and computer network security was developed using the overarching concept of *networks* as the bridge between the two. This is appropriate because the principal idea behind terrain (as it relates to security) is that of an underlying structure on which (or in which) conflict takes place. Although often thought of as merely the land, it actually encompasses a variety of structural elements (e.g., fortification) as well as more complex relationships (e.g., the spatial relationships represented by topology).

Viewing computer networks from a military security perspective makes sense not only because many of the structures are topologically similar, but also because

thinking of them as similar may allow people to visualize, and therefore understand network conflicts better. It may also allow for incorporation of many real world terrain analysis features into computer network analysis to locate vulnerabilities more quickly. In the long run, thinking of computer networks from a military geography perspective may contribute to the development of stronger network defenses.

Beyond the general argument for treating computer networks in military terms, this thesis presented several specific examples of how categories used in traditional terrain analysis might be applied to computer security. The first step of this line of reasoning was an examination of some of the most important borrowed concepts including firewalls, bastion hosts, and demilitarized zones. More robust analogues for computer networks were suggested (but not tested) via the militarily significant measures of distance, density, population mobility, visibility, movement, and interaction. Finally, two specific elements of military terrain analysis – cover and concealment - were examined in detail and experimentally tested to describe functional similarities between counterparts in physical and computer security.

These theoretical connections were empirically tested by examining the performance of several servers and client computers within two experiments. Both the experiment testing for firewalls-as-cover and the experiment testing for service-hiding-as-concealment provided evidence that many of the characteristics of cover and concealment do translate between the real world and computer networks. A third experiment using historical ethnographic data on warfare in primitive societies

suggests that relationships between geographic terrain factors and military success is not deterministic.

Ultimately, this thesis shows that the military metaphor for computer network security can be justifiably applied. This is true because both military operations and computer networks have good analogs with each other that share functional characteristics between them. This means that at least some features established as common are not only *spatially* similar, but also have similar implications for and effects on attack patterns. Although functional similarity may *indicate* appropriateness, it does not firmly establish it – in fact no metaphor can *a priori* be judged in such a way. The life in a metaphor comes also from its use. The current employment of military terminology in computer network security – misapplied though it may sometimes be – is enough to establish that the physical security metaphor for computer networks as at least minimally useful. These two together – a strong argument for the *appropriateness* and the current widespread *use* of the metaphor – constitute a very good reason for believing that thinking of (and ultimately visualizing) computer networks in military terms is an invaluable approach towards solving the security dilemma.

5.2 Outlook

Further work on the theoretical connections between military geography and computer network security should ideally concentrate on testing for the full suite of analogs suggested. This work may help computer network defense by encouraging researchers to develop security strategies that draw on millennia of physical security

experience rather than recent experience. While differences between physical security and computer network security abound, the similar underlying spatial arrangements of both realms may provide clues for estimating the likely success and failures of strategies.

The work here was also intended to justify geographically-based visualizations of computer networks (as in Fisk, Smith, *et al.* 2003). Such visualizations may simplify network security management by utilizing pre-existing ideas of terrain, cover, and concealment. Since many humans have experience with some forms of spatial strategy (from lessons in history, game play, and other sources), visualizations utilizing spatially strategic elements (e.g., cover and concealment) may aid in processing of security information. The cognitive and security benefits remain to be tested, but justification for their employment can, in part, be found in some of the arguments presented in this thesis.

BIBLIOGRAPHY

- Addington, L. H. (1994). The Patterns of War Since the Eighteenth Century.
Bloomington, Indiana University Press.
- Albert, R., H. Jeong, et al. (1999). Diameter of the World-Wide Web. Nature: 130.
- Albert, R., H. Jeong, et al. (2000). Error and attack tolerance of complex networks.
Nature. 406: 378.
- Alford, L. C. L. D., Jr. (2001). Cyber Warfare: A New Doctrine and Taxonomy,
United States Air Force.
- Anderson, R. (2001). Security engineering : a guide to building dependable
distributed systems. New York, Wiley.
- Anderson, R. H., R. Brackney, et al. (2000). Advanced Network Defense Research,
Hawaii, RAND.
- Arquilla, J., D. F. Ronfeldt, et al. (1997). In Athena's camp: preparing for conflict in
the information age. Santa Monica Calif., Rand.
- Arquilla, J., D. F. Ronfeldt, et al. (1999). The emergence of noopolitik : toward an
American information strategy. Santa Monica CA, Rand.
- Atherton, M., J. Zhuang, et al. (2003). "A functional MRI study of high-level
cognition. I. The game of chess." Cognitive Brain Research Vol. 16(No. 1):
pp 26-31.
- Avolio, F. M. (1999). The Castle Defense. Performance Computing Magazine.
- Avolio, F. M. and M. J. Ranum (1994). A Network Perimeter with Secure External
Access. ISOC NDSS Symposium.
- Barnett, R. D. (1963). "Xenophon and the Wall of Media." Journal of Hellenic
Studies 83: 1-26.
- Barnett, R. W. (1998). Information Operations, Deterrence, and the Use of Force,
United States Navy.
- Barringer, R. E. and R. K. Ramers (1972). War: patterns of conflict. Cambridge
Mass., MIT Press.
- Beaulieu, P.-A. (1993). "An Episode in the Fall of Babylon to the Persians." Journal
of Near Eastern Studies Vol. 52(No. 4): 241-261.
- Beeler, J. H. (1956). "Castles and Strategy in Norman and Early Angevin England."
Speculum Vol. 31(No. 4): pp. 581-601.
- Belcher, T. and E. Yoran (2002). Riptech Internet Security Threat Report, Riptech,
Inc.
- Bellovin, S. M. (2000). Distributed Firewalls, Cisco World.
- Bellovin, S. M. and W. R. Cheswick (1994). Network firewalls. IEEE
Communications Magazine. 32: 50.
- Benedikt, M. (1991). Cyberspace : first steps. Cambridge Mass., MIT Press.
- Bennett, P. G. and Institute of Mathematics and Its Applications (1987). Analysing
conflict and its resolution : some mathematical contributions : based on the
proceedings of a conference organized by the Institute of Mathematics and Its

- Applications on the mathematics of conflict and its resolution, held at Churchill College, Cambridge in December 1984. Oxford, Oxford University Press.
- Bernstein, A. H., M. Libicki, et al. (1998). "High-tech: The future face of war? A debate." Commentary V105(N1): 28-34.
- Biddle, S. (2001). "Rebuilding the Foundations of Offense-Defense Theory." J Politics 63(3): 741-774.
- Blaze, M. and S. M. Bellovin (2000). Tapping on my Network Door. Communications of the ACM. 43: 136.
- Brewer, R. J., G. C. Boon, et al. (2000). Roman fortresses and their legions. London Cardiff, Society of Antiquaries of London, National Museums & Galleries of Wales.
- Broido, A. and k. claffy (2001). Internet topology: connectivity of IP graphs, CAIDA.
- Carol, H. (1960). "The Hierarchy of Central Functions within the City." Annals of the Association of American Geographers Vol. 50(No. 4): pp. 419-438.
- CERT Coordination Center (2003). CERT/CC Statistics 1988-2002, Carnegie Mellon University.
- Chen, X., D. Zhang, et al. (2003). "A functional MRI study of high-level cognition: II. The game of GO." Cognitive Brain Research Vol. 16(No. 1): pp 32-37.
- Cheswick, B. and H. Burch (2001). Internet Mapping Project, Lumeta Corporation.
- Clausewitz, C. v. (1968). On war. Baltimore, Penguin Books.
- Clausewitz, C. v. (1984). On War. New York, Alfred A. Knopf.
- Cohen, R., K. Erez, et al. (2001). "Breakdown of the internet under intentional attack." Physical Review Letters V86(N16): 3682-3685.
- Cohen, R., K. Erez, et al. (2001). "Comment on "Breakdown of the Internet under intentional attack" - Reply - art. no. 219802." Physical Review Letters V8721(N21): 9802.
- Collins, J. M. (1998). Military geography for professionals and the public. Washington D.C., Brassey's.
- Couclelis, H. (1998). Worlds of Information: The Geographic Metaphor in the Visualization of Complex Information. Cartography & Geographic Information Systems. 25: 209.
- Cox, K. R. (1995). "Concepts of space, understand in human geography, and spatial analysis." Urban Geography 16(4): 304-326.
- Critical Infrastructure Assurance Office (2003). About CIAO, CIAO.
- Croft, B. and J. Gilmore (1985). Bootstrap Protocol (BOOTP): RFC 951.
- Curtin, M. and M. J. Ranum (2000). Internet Firewalls: Frequently Asked Questions.
- Danyliw, R. (2003). CERT® Advisory CA-2003-04 MS-SQL Server Worm, CERT/CC.
- De Blij, H. J. (2000). Wartime encounter with geography. Lewes Sussex, Book Guild.
- De la Croix, H. (1972). Military considerations in city planning: fortifications. New York, G. Braziller.
- Delano-Smith, C. and R. J. P. Kain (1999). English maps : a history. Toronto,

- University of Toronto Press.
- Denning, D. (2000). Cyberterrorism. Special Oversight Panel on Terrorism, Committee on Armed Terrorism. Washington, DC, Georgetown University.
- Doctor Electron (2002). Computers Connected to IPv4 Address Space, Global Services.
- Doctor Electron (2002). Internet Host Behavior Statistics by Port, Global Services.
- Dodge, M. (1998). Accessibility to Information within the Internet: How can it be Measured and Mapped, Center for Advanced Spatial Analysis.
- Dodge, M. and R. Kitchin (2001). Mapping cyberspace. London ; New York, Routledge.
- Droms, R. (1993). Interoperation Between DHCP and BOOTP: RFC 1534.
- Dumas, A. and D. Coward (1991). The Man in the Iron Mask. Oxford ; New York, Oxford University Press.
- Dumas, A. and D. Coward (1995). The Vicomte de Bragelonne. Oxford ; New York, Oxford University Press.
- Dumas, A. and D. Coward (1995). Louise de la Vallière. Oxford ; New York, Oxford University Press.
- Eade, J. (1996). Chess for dummies. Foster City, CA, IDG Books Worldwide.
- Eckmann, S. T., G. Vigna, et al. (2002). "STATL: An Attack Language for State-based Intrusion Detection." Journal of Computer Security 10(no. 1/2): 71-104.
- Eick, S. G. (2001). "Visualizing Online Activity." Communications of the ACM 44(8): 45-50.
- Everard, J. (2000). Virtual states : the Internet and the boundaries of the nation state. London ; New York, Routledge.
- Fabrikant, S. I. and B. P. Buttenfield (2001). "Formalizing Semantic Spaces For Information Access." Annals of the Association of American Geographers 91(2): 263-280.
- Ferrill, A. (1997). The origins of war : from the Stone Age to Alexander the Great. Boulder, Colo., Westview Press.
- Fisch, E. A. and G. B. White (2000). Secure computers and networks : analysis, design, and implementation. Boca Raton Fla., CRC Press.
- Fisk, M., S. A. Smith, et al. (2003). Immersive Network Monitoring.
- Forno, R. and R. Baklarz (1999). The art of information warfare : insight into the knowledge warrior philosophy. Parkland? Fla., Universal Publishers/UPUBLISH.COM.
- Forrest, S., S. A. Hofmeyr, et al. (1997). "Computer immunology." Communications of the Acm V40(N10): 88-96.
- Frederick and J. Luvaas (1999). Frederick the Great on the art of war. New York, Da Capo Press.
- Giddens, A. (1987). The nation-state and violence. Berkeley, University of California Press.
- Gien, M. and J.-L. Grangé (1979). Flow control in computer networks : proceedings of the International Symposium on Flow Control in Computer Networks,

- Versailles, France, February 12-14, 1979. Amsterdam ; New York
New York, North-Holland Pub. Co.
sole distributors for the U.S.A. and Canada Elsevier North-Holland.
- Glaser, C. L. and C. Kaufmann (1998). "What is the offense-defense balance and can we measure it?" International Politics 1998(v22 n4): p44(39).
- Golledge, R. G. and G. Rushton (1976). Spatial choice and spatial behavior : geographic essays on the analysis of preferences and perceptions. Columbus, Ohio State University Press.
- Gray, C. S. (1996). The continued primacy of geography. Orbis. 40: 247.
- Gray, C. S. (1997-98). "RMAs and the Dimensions of Strategy." Joint Force Quarterly(No. 17): 50-54.
- Haeni, R. E. (1997). Firewall Penetration Testing, The George Washington University, Cyberspace Policy Institute.
- Hayes, R. E., G. F. Wheatley, et al. (1997). Information warfare and deterrence. Washington D.C.?, National Defense University Institute for National Strategic Studies.
- Hopf, T. (1991). "Polarity, The Offense Defense Balance, and War." The American Political Science Review Vol. 85(No. 2): pp 475-493.
- Horan, T. A. (2000). Digital places : building our city of bits. Washington, D.C., ULI-the Urban Land Institute.
- Houghton Mifflin Company. (2001). Webster's II New College Dictionary. Boston, Houghton Mifflin Company.
- Howard, J. D. (1997). An Analysis Of Security Incidents On The Internet 1989-1995. Engineering and Public Policy. Pittsburgh, Carnegie Mellon University.
- Howard, M. E. (1976). War in European history. London ; New York, Oxford University Press.
- Huffaker, B., M. Fomenkov, et al. (2002). Distance Metrics in the Internet. IEEE International Telecommunications Symposium, Rio De Janeiro, Brazil.
- Internet Assigned Numbers Authority (2003). IP Address Services.
- Internet Assigned Numbers Authority (2003). Internet Protocol V4 Address Space.
- Isard, W., C. Smith, et al. (1988). Arms races, arms control, and conflict analysis : contributions from peace science and peace economics. New York, Cambridge University Press.
- Janelle, D. G. and D. C. Hodge (2000). Information, place, and cyberspace : issues in accessibility. Berlin ; New York, Springer.
- Jenkins, A. and D. Pepper (1985). The Geography of peace and war. Oxford Oxfordshire ; New York NY USA, B. Blackwell.
- Jervis, R. (1984). The illogic of American nuclear strategy. Ithaca, Cornell University Press.
- Johnson, A. (1983). Roman forts of the 1st and 2nd centuries AD in Britain and the German Provinces. London, A. & C. Black.
- Johnston, R. J. and P. Claval (1984). Geography since the Second World War : an international survey. London
Totowa N.J., Croom Helm

- Barnes & Noble.
- Jones, S. B. (1959). "Boundary Concepts in the Setting of Place and Time." Annals of the Association of American Geographers 49(3): 241-255.
- Karamustafa, A. T. (1992). Military, Administrative, and Scholarly Maps and Plans. Cartography in the Traditional Islamic and South Asian Societies. J. B. Harley and D. Woodward. Chicago, University of Chicago Press. 1: 209-227.
- Katzner, D. W. (1983). Analysis without measurement. Cambridge Cambridgehire ; New York, Cambridge University Press.
- Keegan, J. (1978). The face of battle. Harmondsworth, Penguin.
- Keegan, J. (1993). A History of Warfare. New York, Vintage Books.
- Keegan, J. (1999). The book of war. New York, Viking.
- Keeley, L. H. (1996). War before civilization. New York, Oxford University Press.
- Kemmerer, D. and G. Vigna (2002). "Intrusion detection: A brief history and overview." Computer(SUPPS): 27-30.
- Kemmerer, R. A. (2002). Computer Security. Encyclopedia of Software Engineering. J. J. Marciniak. New York, John Wiley.
- Keuhl (1997). Defining Information Power, National Defense University Strategic Forum.
- Khalilzad, Z. M., J. P. White, et al. (1999). Strategic appraisal : the changing role of information in warfare. Santa Monica CA, Rand.
- Klare, M. T., D. C. Thomas, et al. (1994). World security : challenges for a new century. New York, St. Martin's Press.
- Klaus, C. W. and M. J. Ranum (1996). Does scanning for vulnerabilities mean your firewall is safe? InfoWorld. 18: 79.
- Koprowski, G. J. (2003). Internet Protocol for the Future: IPv6 Poised for Adoption. TechNewsWorld.
- Krepinevich, A. F. (1994). "Cavalry to computer; the pattern of military revolutions." The National Interest(n37): 30-43.
- Kristof, L. K. D. (1959). "The Nature of Frontiers and Boundaries." Annals of the Association of American Geographers 49(3): 269-282.
- Krol, E. and E. Hoffman (1993). FYI on 'What Is The Internet?': RFC 1462.
- Lakoff, G. (1992). The Contemporary Theory of Metaphor. Metaphor and Thought. A. Ortony, Cambridge University Press.
- Lakoff, G. and M. Johnson (1980). Metaphors we live by. Chicago, University of Chicago Press.
- Langworthy, C. N. (1995). The geography of war. Oklahoma City Okla., Cooper House.
- Lawrence, A. W. (1969). Fortified trade-posts: the English in West Africa, 1645-1822. London, Cape.
- Lawrence, A. W. (1979). Greek aims in fortification. Oxford New York, Clarendon Press Oxford University Press.
- Lawrence, A. W. and R. A. Tomlinson (1996). Greek architecture. New Haven, Yale University Press.

- Lawrence, T. E. (1935). Seven pillars of wisdom; a triumph. Garden City, N.Y., Doubleday Doran & company inc.
- Leigheb, M. (1998). The Wall of Aurelian, Comitato di Azione sul Territorio.
- Levy, J. S. (1984). "The Offensive/Defensive Balance of Military Technology: A Theoretical and Historical Analysis." International Studies Quarterly Vol. 28(No. 2): pp. 219-238.
- Levy, S. (2001). Crypto: The story of how a group of code rebels saved your privacy on the internet. Newsweek: 42.
- Libicki, M. (1996). The emerging primacy of information. Orbis. 40: 261-276.
- Libicki, M. (1998). "Ghosts in the Machines?" U.S. Foreign Policy Agenda 3(4).
- Libicki, M. C. (1997). Information Dominance, National Defense University Strategic Forum.
- Libicki, M. C. (1997). "Information warfare: A brief guide to defense preparedness." Physics Today V50(N9): 40-45.
- Lieber, K. A. (2000). "Grasping the Technological Peace." International Security 25(1): 71.
- Lucas, H. C., Jr. (2001). Information Technology and Physical Space. Communications of the ACM. 44: 89.
- Lynch, K. (1979). The image of the city. Cambridge Mass., Technology Press.
- Macfadyen, M. (1998). The Game of Go: Learning and Mastering the Most Challenging Game In the World, Carlton Books Limited.
- Mansbach, R. W. (1994). The global puzzle : issues and actors in world politics. Boston, Houghton Mifflin.
- Mathis, M., G. L. Huntton, et al. (2001). Traffic Dynamics Testbed, CAIDA.
- McKay, N. (1998). China: The Great Firewall, Wired.com.
- McLendon, C. J. (2001). Information Warfare: Impacts and Concerns, United States Air Force.
- Merriam-Webster Inc. (2003). Merriam-Webster's collegiate dictionary. Springfield, Mass., Merriam-Webster Inc.
- Mitchell, W. J. (1995). City of bits : space, place, and the infobahn. Cambridge, Mass., MIT Press.
- Molander, R. C., United States. Dept. of Defense. Office of the Secretary of Defense, et al. (1998). Strategic information warfare rising. Santa Monica, CA, RAND.
- Molander, R. C., P. Wilson, et al. (1996). Strategic information warfare : a new face of war. Santa Monica CA, Rand.
- Montello, D. R., S. I. Fabrikant, et al. (2003). Testing the first law of cognitive geography on point-display systems. Spatial information theory: Foundations of geographic information science. Proceedings of COSIT '03. Lecture Notes in Computer Science 2825, Berlin, Springer.
- Moore, D., K. Keys, et al. (2001). The CoralReef software suite as a tool for system and network administrators, CAIDA.
- Moore, D., V. Paxson, et al. (2003). The Spread of the Sapphire/Slammer Worm, CAIDA.

- Moore, D., G. Voelker, et al. (2001). Inferring Internet Denial-of-Service Activity, CAIDA.
- Mundy, B. E. (1998). Mesoamerican Cartography. Cartography in the Traditional African, American, Arctic, Australian, and Pacific Societies. D. Woodward and G. M. Lewis. Chicago, University of Chicago Press. 3: 183-256.
- Napoleon and J. Luvaas (1999). Napoleon on the art of war. New York, NY, Free Press.
- National Research Council (U.S.). Computer Science and Telecommunications Board (2002). Cybersecurity today and tomorrow : pay now or pay later. Washington D.C., National Academy Press.
- Nicolle, D. (1999). "Medieval Warfare: The Unfriendly Interface." The Journal of Military History Vol. 63(No. 3): pp. 579-599.
- O'Brien, P. (1975). "L'Embastillement de Paris: The Fortification of Paris during the July Monarchy." French Historical Studies Vol. 9(No. 1): pp. 63-82.
- Oliver, M. (1999). TCP/IP FAQ, Sun Microsystems, Inc.
- Olmstead, A. T. (1922). "The Rise and Fall of Babylon." The American Journal of Semitic Languages and Literatures Vol. 38(No. 2): 73-96.
- Olsen, F. (2002). "The Growing Vulnerability of Campus Networks." The Chronicle of Higher Education 48(27): A35.
- O'Sullivan, P. (1992). "Geography and Political Power - the Geography of Nations and States - Slove,P." Annals of the Association of American Geographers V82(N1): 180-181.
- O'Sullivan, P. M. (1991). Terrain and tactics. New York, Greenwood Press.
- O'Sullivan, P. M. (2001). The geography of war in the post Cold War world. Lewiston, N.Y., Edwin Mellen Press.
- O'Sullivan, P. M. and J. W. Miller (1983). The geography of warfare. London, Croom Helm.
- Otterbein, K. F. (1970). The evolution of war; a cross-cultural study. n.p., Human Relations Area Files Press.
- Otterbein, K. F. (1994). Feuding and warfare : selected works of Keith F. Otterbein. Langhorne, Pa., Gordon and Breach.
- Oxford University Press. (1999). The Concise Oxford English Dictionary. Oxford, Oxford University Press.
- Oxford University Press. (2002). The Unabridged Oxford English Dictionary. Oxford, England, Oxford University Press.
- Payne-Gallwey, R. (1995). The book of the crossbow. New York, Dover Publications.
- Peltier, L. C. and G. E. Percy (1966). Military geography. Princeton N.J., Van Nostrand.
- Pfleeger, C. P. (1997). Security in computing. Upper Saddle River NJ, Prentice Hall PTR.
- Phillips, T. R., Sunzi, et al. (1940). Roots of strategy; a collection of military classics. Harrisburg, Pa., The Military service publishing company.
- Postel, J. (1980). User Datagram Protocol: RFC 768.

- Postel, J. (1981). Internet Control Message Protocol: RFC 792.
- Poterba, J. M. and National Research Council (U.S.). Board on Science Technology and Economic Policy (1997). Borderline case : international tax policy, corporate research and development, and investment. Washington D.C., National Academy Press.
- President's Critical Infrastructure Protection Board (2002). The National Strategy to Secure Cyberspace (Draft), White House.
- Quester, G. H. (2003). Offense and defense in the international system. New Brunswick N.J., Transaction Publishers.
- Ranum, M. (1993). Thinking about Firewalls, Digital Equipment Corporation.
- Ranum, M. and S. Mace (1998). Finding your firewall. Byte. 23: 96NA3.
- Ranum, M. J. (1992). A Network Firewall, Digital Equipment Corporation.
- Rekhter, Y., B. Moskowitz, et al. (1996). Address Allocation for Private Internets.
- Richard, J. (2001). Configuring Snort, MySQL and ACID on Windows NT, SANS.
- Rosecrance, R. (1995). "The Obsolescence of Territory." New Perspectives Quarterly Vol. 12(No. 1): pp. 44-50.
- Rosecrance, R. (1996). The rise of the virtual state. Foreign Affairs. 75: 45.
- Sack, R. D. (1983). "Human Territoriality: A Theory." Annals of the Association of American Geographers 73(1): 55-74.
- Sahlins, P. (1990). "Natural Frontiers Revisited: France's Boundaries since the Seventeenth Century." The American Historical Review Vol. 95(No. 5): pp. 1423-1451.
- Scambray, J., S. McClure, et al. (2001). Hacking exposed : network security secrets & solutions. Berkeley Calif., Osborne/McGraw-Hill.
- Schneider, F. B. and National Research Council (U.S.). Committee on Information Systems Trustworthiness (1999). Trust in cyberspace. Washington D.C., National Academy Press.
- Schwartz, W. (1996). Information warfare : cyberterrorism--protecting your personal security in the electronic age. New York Emeryville CA, Thunder's Mouth Press Distributed by Publishers Group West.
- SecuritySearch (2004). Honey Pot: A SearchSecurity Definition, Techtargget.com.
- Shubik, M. (1983). Mathematics of conflict. Amsterdam ; New York New York N.Y., North-Holland Pub. Co. Sole distributors for the U.S.A. and Canada Elsevier Science Pub. Co.
- Siamwalla, R., R. Sharma, et al. (1998). Discovering Internet Topology, Cornell University.
- Singer, J. D. and P. F. Diehl (1990). Measuring the correlates of war. Ann Arbor, University of Michigan Press.
- Smith, B. (1997). The Cognitive Geometry of War. Aktuelle Fragen politischer Philosophie : Gerechtigkeit in Gesellschaft und Weltordnung : Akten des 19. Internationalen Wittgenstein-Symposiums, 11. bis 18. August 1996 Kirchberg am Wechsel (Österreich). P. Koller and K. Puhl. Wien, Hölder-Pichler-Tempsky: 394-403.

- Smith, B. and A. C. Varzi (2000). "Fiat and bona fide boundaries." Philosophy and Phenomenological Research 60(2): 401-420.
- Socolofsky, T. and C. Kale (1991). A TCP/IP Tutorial: RFC 1180.
- Sunzi, M.-c. Sawyer, et al. (1996). The complete art of war. Boulder Colo., Westview Press.
- Takuan, S. o. o. and W. S. Wilson (1986). The unfettered mind : writings of the Zen master to the sword master. Tokyo ; New York
New York N.Y., Kodansha International
Distributed in the U.S. by Kodansha International/USA through Harper & Row.
- Taylor, T. C. (1990). "Principles of Modern Tactics - the n-Square Law and Beyond." Defense Analysis 6(3): pp. 277-288.
- Tobler, W. (1970). "A computer movie simulating urban growth in the Detroit region." Economic Geography 46: 234-240.
- Toffler, A. (1970). Future shock. New York,, Random House.
- Toffler, A. (1980). The third wave. New York, Morrow.
- Toffler, A. (1983). Previews & premises : an interview with the author of Future shock and The third wave. New York, W. Morrow.
- Toffler, A. and H. Toffler (1993). War and anti-war : survival at the dawn of the 21st century. Boston, Little Brown.
- Treece, C. D., T. Ledoux, et al. (1999). U.S. Army Europe's Information Assurance Program, United States Navy.
- Tufte, E. R. (1983). The visual display of quantitative information. Cheshire, Conn. (Box 430, Cheshire 06410), Graphics Press.
- Tufte, E. R. (2001). The visual display of quantitative information. Cheshire, Conn., Graphics Press.
- United States. Department of Defense (2003). DOD Dictionary of Military and Associated Terms, United States. Department of Defense.
- United States. Dept. of State. Office of the Geographer (1961). Boundary concepts and definitions. Washington D.C., Office of the Geographer U.S. Dept. of State.
- United States. Dept. of the Army (1986). Operations. Washington DC, Headquarters Dept. of the Army.
- United States. Dept. of the Army (1990). Terrain Analysis. Washington DC, Headquarters Dept. of the Army.
- United States. Dept. of the Army (1990). Manual of Common Tasks. Washington DC, Headquarters Dept. of the Army.
- United States. Dept. of the Army (1992). Infantry Rifle Platoon and Squad. Washington DC, Headquarters Dept. of the Army.
- Vigna, G. A Formal Model for Firewall Testing. Milano, Italy, Politecnico di Milano: 1-10.
- Vigna, G. (1996). A Topological Characterization of TCP/IP Security. Milano, Italy, Politecnico di Milano: 1-27.
- Vigna, G., S. Eckmann, et al. (2000). Attack Languages. IEEE Information Survivability Workshop, Boston.

- Vigna, G., S. T. Eckmann, et al. (2000). The STAT Tool Suite. DISCEX, IEEE Press.
- Vigna, G. and R. A. Kemmerer (1998). NetSTAT: A Network-based Intrusion Detection Approach. 14th Annual Computer Security Application Conference, Scottsdale, Arizona.
- Vigna, G. and R. A. Kemmerer (1999). "NetSTAT: A Network-based Intrusion Detection System." Journal of Computer Security 7(1).
- Vigna, G., R. A. Kemmerer, et al. (2001). Designing a Web of Highly-Configurable Intrusion Detection Sensors. 4th International Symposium on Recent Advances in Intrusion Detection, Davis, CA.
- Waltz, K. N. (1959). Man, the state, and war: a theoretical analysis. New York, Columbia University Press.
- Webopedia (2003). What is a honeypot?, Webopedia Computer Dictionary.
- Wenger, G. R. (1991). The story of Mesa Verde National Park. Mesa Verde National Park, Colo., Mesa Verde Museum Assoc.
- Wheeler, E. L. (1993). "Methodological Limits and the Mirage of Roman Strategy: Part I." The Journal of Military History Vol. 57(No. 1): pp. 7-41.
- Whittlesey, D. (1935). "The Impress of Effective Central Authority upon the Landscape." Annals of the Association of American Geographers Vol. 25(No. 2): pp 85-97.
- Winter, F. E. (1971). Greek fortifications. Toronto, University of Toronto Press.
- Woodcock, S. (1995). The "Recognizing Strategic Dispositions" Thread, gameai.com.
- Wolf, H., Ed. (1981). Webster's New Collegiate Dictionary, G. & C. Merriam Co.
- Wright, Q. (1965). A study of war. Chicago, University of Chicago Press.
- Yadin, Y. (1963). The art of warfare in Biblical lands in the light of archaeological study. New York,, McGraw-Hill.
- Yamamoto, T. and W. S. Wilson (1983). Hagakure : the book of the samurai. Tokyo New York, Kodansha International distributed in the United States by Harper & Row.
- Ye, N., J. Giordano, et al. (2001). A Process Control Approach to Cyber Attack Detection. Communications of the ACM. 44: 76-82.
- Yee, C. D. K. (1994). Chinese Maps in Political Culture. Cartography in the Traditional East and Southeast Asian Societies. J. B. Harley and D. Woodward. Chicago, University of Chicago Press. Book Two: 71-95.
- York, H. F. and John F. Kennedy School of Government. Center for Science and International Affairs (1987). Does strategic defense breed offense? Cambridge Mass.
- Lanham MD ;, Center for Science and International Affairs Harvard University University Press of America.
- Zalenski, R. (2002). "Firewall technologies." IEEE Potentials 21(1): 24-29.
- Zilinskas, R. A. (2000). Biological warfare : modern offense and defense. Boulder CO, Lynne Rienner Publishers.
- Zwicky, E. D., S. Cooper, et al. (2000). Building Internet firewalls. Beijing ;

Cambridge Mass., O'Reilly.

APPENDIX

As part of this thesis, several software packages were used as part of the experimental process. These packages, along with their location on the world wide web are listed below:

- *Firedaemon* was used to start the intrusion detection software running during boot, and to restart it automatically if it crashed during startup. It can be found on the web at <http://www.firedaemon.com>.
- *Libpcap* and *winpcap* were used to capture data packets from the network.
 - *Libpcap* is on the web at <http://sourceforge.net/projects/libpcap/>.
 - *Winpcap* is on the web at <http://winpcap.polito.it/>.
- *Snort* was used to detect and record intrusions. It can be found on the web at <http://www.snort.org>.