SCALABLE SECURITY:

CYBER THREAT INFORMATION SHARING IN THE INTERNET AGE

A THESIS

SUBMITTED TO THE

INTERSCHOOL HONORS PROGRAM IN
INTERNATIONAL SECURITY STUDIES

CENTER FOR INTERNATIONAL SECURITY AND COOPERATION
FREEMAN SPOGLI INSTITUTE FOR INTERNATIONAL STUDIES

STANFORD UNIVERSITY

By

Connor Gilbert

May 2014

Advisors:

Prof. Martin E. Hellman

Dr. Thomas A. Berson

Abstract

The federal government has attempted to foster cyber threat information sharing within U.S. critical infrastructure industries for at least 16 years. Such efforts have produced today's complex constellation of threat analysis centers and sharing organizations in various industries and government agencies, but have brought inconsistent observed improvement in cybersecurity outcomes. Meanwhile, successful but limited sharing relationships have formed separately among private companies and individuals. Objections to information sharing that are raised in the literature and press fall into a limited number of categories, each of which can be refuted relatively easily; the problem is therefore better viewed as a problem of insufficient perceived *benefits* rather than prohibitive *costs*.

While history demonstrates that sharing is a powerful tool for addressing collective threats, new analysis is clearly needed to explain why effective cyber sharing seems so difficult to accomplish effectively.

The *Computational Policy* analytical approach defined herein brings the power of the abstractions used in computer systems design to bear on difficult policy problems, allowing new analytical insights to inform policy choices.

Using this approach and new understanding of the cybersecurity threat landscape, analyses of three archetypal sharing designs—the watch floor, social network, and high-volume data sharing models—are presented and aspects of the current system are questioned.

Cybersecurity programs centered on human analysis are at a disadvantage as networks continue to grow; approaches which harness the power of algorithmic thinking and ever-growing computational resources have better hope of succeeding in the Internet Age.

To Dan Henry Briganti (1932-2013) and
Suzanne Covington Briganti (1932-2014),
who personified selfless service.

Requiem aeternam dona eis, Domine,
et lux perpetua luceat eis.

Acknowledgements

This thesis would never have seen the light of day without the encouragement and expert guidance of my advisors, Prof. Martin Hellman and Dr. Tom Berson, with whom it was a true honor to work.

Neither would it have been completed without the support of the rest of the students in the CISAC Honors Program, who I heartily respect; our "small-unit cohesion" was critical to the success of our mission. I owe a true debt of gratitude to our program co-directors and mentors, Profs. Martha Crenshaw and Chip Blacker; our Teaching Assistant, Shiri Krebs; and our guides through Washington, Amb. Karl Eikenberry and Dr. Thomas Fingar.

Especially because I sometimes felt like I was engaging in international relations while making the journey back and forth between Gates Computer Science and Encina Hall, I am grateful for the assistance of the entire CISAC community. I especially thank Tim Junio, who guided my topic selection, helped me translate between computer science and social science, and provided helpful advice throughout the writing process. I also can't thank Honors Program alumnus Scott Bade enough for telling me that, yes, engineers can apply to the CISAC Honors Program, too!

Thanks are also due to Scott Charney, Gail Kent, Herb Lin, Jane Holl Lute, and others who shared their insights at various stages of this project.

This section would be incomplete without an acknowledgement of the support of my parents, siblings, and other family members. I would have made it neither to Stanford nor to CISAC without their encouragement.

To all of my friends at Stanford and afar who have watched me labor with this project throughout senior year, who are too numerous to list, thank you for understanding and for your (frequent) moral support.

Contents

List of Figures

List of Tables

List of Algorithms

Chapter 1

Context

> “The rapid proliferation and integration of telecommunications and computer systems have connected infrastructures to one another in a complex network of interdependence. This interconnectivity has created a new dimension of vulnerability, which, when combined with an emerging constellation of threats, poses unprecedented national risk.”
>
> President’s Commission on Critical Infrastructure Protection
>
> (October 13, 1997: President’s Commission on Critical Infrastructure Protection, *Critical Foundations: Protecting America’s Infrastructures*, ix)

1.1 Introduction

“Cyber security has gone mainstream.”[1] Despite years of alarming vulnerability, the risks posed by computer network security issues have only gained widespread public attention in the past few years.[2] The United States Government, though, has long been attempting to improve the nation’s cybersecurity posture. As early as 1997, the potential for computer network attacks with *kinetic effects*—those with an impact in the physical world—was recognized as a potentially debilitating national vulnerability in reports to the President.[3] Since then, as governmental inertia and technological progress have both continued unabated, the

[1] Mandiant, *M-Trends: Beyond the Breach (2014 Threat Report)*, 1.

[2] Of course, specialists have been discussing such issues for years. However, only recently have cybersecurity issues gained prominence among the general public. Popular media outlets now frequently report on security vulnerabilities, the behavior of governments in cyberspace, and issues of technology policy and law.

[3] President’s Commission on Critical Infrastructure Protection, *Critical Foundations: Protecting America’s Infrastructures*, 7.

technical threat and the nation's vulnerability to it have grown.[4]

Virtually nothing in computer technology looks or feels the same as it did 15 years ago: computers that filled entire rooms are outperformed by today's tablets and smartphones, laptop computers store more information than mainframes from decades past, and cell phones have faster Internet connections now than entire office buildings did then. The raw capability of processors, disks, and networks has grown at quite literally an exponential pace.[5] Such progress has only just begun to "level off" as technology companies run into fundamental physical and manufacturing limits, but have held true for many years.

There has been widespread agreement since the late 1990s that *information sharing* could improve the defense of U.S. critical infrastructure. Even the earliest reports and strategies suggest that responsible entities might wish to share cyber threat information, and public statements from policymakers often encourage increased sharing and cooperation as a way to improve the nation's overall security posture.

Cyber threat information sharing schemes have been promoted by the US Government for more than 16 years, but even the most high-profile efforts have fallen far short of their objectives. While a few industries have successfully fostered sharing among their member companies, report after report indicates that many critical infrastructure sectors are woefully underprepared to deal with routine or emerging computer and network security threats. The Internet has grown from just four hosts in 1969 to at least a billion in 2014;[6] even since 1998, the Internet has grown by a factor of over 27.[7] This type of remarkable, paradigm-shifting

[4] As early as 2001, informed observers noted this divergence. See Moteff, "Critical Infrastructures: Background and Early Implementation of PDD-63."

[5] Such advances were predicted and are tracked by Moore's, Kryder's, and Nielsen's Laws, respectively. These "laws" are more accurately termed *predictions* that the number of transistors one could manufacture per chip would double approximately every two years; that the density of disk storage would double approximately every 13 months; and that the bandwidth of Internet communications would experience about 50% growth each year.

[6] Internet Systems Consortium, "Internet Domain Survey, January 2014"; Leiner et al., "A brief history of the Internet," 24.

[7] This survey method is likely to undercount the number of devices actually connected to the Internet due to impediments like firewalls and Network Address Translation (NAT) devices. However, it provides concrete numbers with a consistent methodology over time. The number of hosts discovered in the Internet Systems Consortium's survey in July 1998 was 36,739,000; the same survey in January 2014 return 1,010,251,829 hosts. Based on data from Internet Systems Consortium, "Internet Domain Survey, January 2014."

growth has hardly been seen in any other industry. One feature, though, remains nearly identical: the organizational structures and policy strategies for cyber threat information sharing. Only in the last few years have the policy apparatuses of government begun to adjust to a world that has been rapidly changing underneath them for years. However, even the newest policy prescriptions reprise the themes of older ones, suggesting that reexamination of the assumptions and path dependencies now inherent in cyber threat information sharing is needed.

This thesis argues that information sharing efforts in the United States have suffered from fundamental misunderstandings of both the computer security threat environment and the effective means to counter it. Information sharing is accomplished on a much broader scale than is typically implied by the literature or government statements. Most analysis of why information sharing efforts have fallen short of their objectives concentrates on a small number of objections that have been raised consistently for nearly 15 years. These scruples are not the ultimate reason why prospective participants remain reluctant to share. Instead of continuing to attempt to improve the results of participants' cost-benefit analysis by myopically focusing only on attempts to reduce the perceived costs or risks—a task which risk-averse corporate counsel will ensure has no end—a more promising focus is to work to increase the potential benefit.

Existing sharing designs misunderstand the predominant type of threat against targets. The design of predominant institutional sharing organizations improperly assumes that crisis events will be the most pressing threats. This idea, while perhaps appropriate in physical security or law enforcement, simply does not apply in computer network defense; the "slow drip" of a wide spectrum of espionage, data theft, and attack preparations are not best addressed with such a model. And, apparently in an effort to obviate the need for new legislation or regulation, sharing in the United States is currently balkanized by industry, and limited to sectors designated as "critical infrastructure". Meanwhile, attackers with no specific target ensnare critical infrastructure targets, and highly focused attackers aren't likely to be caught by sharing anyway. Many of the most threatening attackers target multiple industries, meaning that such divisions inappropriately disadvantage network defenders against attackers.

These existing institutions are also unsustainably labor- and expertise-intensive—not only do they provide inadequate security benefits, they do not appear ready to scale to handle a larger workload as the network security problem continues to grow.

This thesis presents a novel technique for policy analysis, termed *Computational Policy*. This simple but powerful technique provides a conduit for knowledge to flow between the policy analysis and the computer systems design communities, with the goal of allowing conceptual frames from computer systems to help policy analysts design better solutions to the "wicked problems" of modern policy.[8] By analyzing sharing organizations using insights from computer systems design, we can achieve more—and more valuable—information sharing, and greater cybersecurity success.

1.2 Why Share Cyber Threat Information?

Computer security best practices often call for increased isolation and separation of critical computers or networks from one another.[9] So, why should cyber threat information sharing—which calls for increased spread of potentially sensitive data, and cooperation between those responsible for disparate networks—be so heavily promoted within government and industry?

This section first reviews existing arguments for cyber threat information sharing. It next presents analogous cases of sharing in a variety of arenas to demonstrate that the impulse to share is not without foundation.

[8] Rosenzweig, "Cybersecurity and Public Goods: The Public/Private 'Partnership'," n. 47.

[9] Firewalls and "air-gapped" networks (which are physically and logically separated from other computers by an "air gap") illustrate this tendency. For example, the National Institute of Standards and Technology recommends "incorporating network segregation where appropriate" in its Cybersecurity Framework (National Institute of Standards and Technology, "Framework for Improving Critical Infrastructure Cybersecurity, Version 1.0," 24).

1.2.1 Calls for Cyber Threat Information Sharing

Some of the clearest calls for cyber threat information sharing come from senior government officials responsible for some aspect of national computer network defense. Former Federal Bureau of Investigation (FBI) Director Robert Mueller appealed in 2013 for efforts to "[establish] channels to share information swiftly" to "effectively disrupt" cyber threats.[10] Former National Security Agency (NSA) Director Keith Alexander similarly highlighted sharing throughout his tenure, including in a presentation which identified a "Coalition of the Connected" that would marshal "the combined talents, efforts, and capabilities of government, private industry, and allies to secure [cyberspace]".[11] However, calls for such sharing are not new: the primary thrust of Presidential Decision Directive 63 (PDD-63), issued in 1998, was to create "a public-private partnership" to address national vulnerabilities in critical infrastructure; four of its 18 pages were dedicated primarily to information sharing as an integral piece of the nation's critical infrastructure defense plan.[12]

Nor are these exhortations solely the province of policymakers: individuals from private industry often promote sharing. The president of prominent firm RSA Security, for instance, remarked, "We must find a way to increase our sharing and the visibility of networks," at a security conference in 2012.[13] Naturally, vendors offering sharing services are also happy to encourage responsible threat sharing; one promises that "[by] uniting and sharing intelligence" the 'good guys' can "shift the odds back into our favor".[14] Finally, similar admonitions appear in the policy analysis literature, which includes calls to "improve the speed and breadth of sharing", create a "robust system for the sharing of cyber-threat information", and make the private sector into an "active partner" with government.[15]

[10] Mueller, "Working Together to Defeat Cyber Threats," 3.

[11] Alexander, "Securing our Government Networks," slide 7.

[12] "Presidential Decision Directive 63: Critical Infrastructure Protection," 12–15.

[13] Tom Heiser, quoted in Mimoso, "Adequate Attack Data and Threat Information Sharing No Longer a Luxury."

[14] IID, "ActiveTrust."

[15] Nelson and Wise, "Homeland Security at a Crossroads: Evolving DHS to Meet the Next Generation of Threats."

1.2.2 Other Instances of Information Sharing

The idea of information sharing itself is not a creation of the Internet age; the impulse to share information for the common good has manifested itself in a variety of contexts over many years. Indeed, information sharing is a form of cooperation, and cooperation is seen a variety of complex systems in nature and society. However, neither cooperation nor information sharing develop without a reason to do so and favorable terms for prospective participants.

When Information is Distributed Suboptimally

When information is distributed in such a way that those who possess knowledge don't derive value from it, but those who *would* derive value from the information don't (or can't) possess it, sharing can bridge the gap. Financial institutions exhibit this pattern, and have very clear motivations to fix it: money is truly on the line.

The Japanese financial industry's promissory note clearinghouse system demonstrates how financial institutions collaborate to mitigate the effect of irresponsible banking customers. In the clearinghouse—nicknamed the "guillotine" for the severe punishments it doles out—banks share information on dishonored notes, suspending firms with more than one delinquent obligation. The entire system operates without the force of public law; the threat of punishment from all of the country's banks acting in concert is enough to achieve near-total compliance.[16] The incentives to share information are clear, including reduced likelihood of a financial institution being required to absorb the losses of a delinquent business. The system was carefully designed over time to correctly align incentives and encourage all participants to fully share the information they possess; penalties are applied to non-sharing banks in order to ensure that the clearinghouse obtains complete and accurate information.[17]

American readers are likely more familiar with the credit reporting system. This

[16] Matsumura and Ryser, "Revelation of Private Information about Unpaid Notes in the Trade Credit Bill System in Japan," 167.

[17] Ibid., 168.

is, again, a clear example of information sharing: credit bureaus gather information from member institutions and other sources and resell it to those same institutions and other customers, with the goal of gathering more information than each buyer could alone. Nearly all members do, in fact, share information on their borrowers with the major credit bureaus through their hundreds of local offices.[18] Information sharing was previously effected through bilateral exchanges between creditors, but an improved sharing design benefitted the industry: the centralized "hub-and-spoke" sharing that replaced this "crisscrossing" mesh of relationships enabled creditors to offer credit to customers with whom they didn't share an extensive personal history.[19]

ChexSystems, while less commonly known than the major credit bureaus, is a similarly large network of financial institutions—claiming as many as 80 percent of U.S. banks and credit unions as members—dedicated to minimizing checking and savings fraud.[20] On its website, ChexSystems bills itself as a "network. . . of member financial institutions that regularly contribute information" that it then "shares. . . among its member institutions to help them assess the risk of opening new accounts."[21]

Financial applications clearly demonstrate the utility of information sharing against problems that involve imbalances in the possession of information versus its value: a creditor with a delinquent customer obtains no immediate new value simply by sharing that information with other creditors—the debt remains delinquent. Though the threat of such sharing likely reduces the likelihood that a borrower will default for fear of losing credit elsewhere, this effect, too, depends on the action of the institution with which the information is shared. Only when information is reciprocally shared between those who *possess* it (*e.g.,* current creditors) and those who *value* and can act on it (*e.g.,* prospective creditors) does it provide the maximum benefit to the group.

[18] Klein, "Credit-Information Reporting: Why Free Spech is Vital to Social Accountability and Consumer Opportunity," 326.

[19] Ibid., 330.

[20] Bankrate.com, "ChexSystems."

[21] Chex Systems, Inc., "Consumer Assistance."

When Reciprocal Action is Desired

A discussion of information sharing of any kind might seem incomplete without at least a mention of the issue of sharing between law enforcement, intelligence, and national security bureaucracies related to the prevention of terrorism or serious crime.[22] In this section, however, we concentrate on more routine law enforcement tasks like serving warrants and recovering stolen property. The prototypical criminal justice information sharing system is the National Crime Information Center (NCIC), operated by the FBI. Used by local police agencies millions of times each day, the NCIC allows information about stolen items, wanted persons, or other law enforcement tasks to flow between police agencies.[23] This sharing infrastructure calls for law enforcement agencies, in return for the implied promise of others' reciprocation, to act on warrants and other requests from other jurisdictions. The "Driver License Compact", which claims 46 states and the District of Columbia as members, similarly shares information across state lines to assure that drivers are appropriately punished for serious driving infractions, even if they occur in other jurisdictions.[24]

When Specialized Information is Needed

Insurance companies, similarly, found it useful to gradually band together to form Underwriters' Laboratories, the source of the familiar "UL®" logo found on most household appliances. Especially in the early 20th century, when fires posed great personal danger to insurance policy holders and therefore great financial peril to insurers, both parties had an "identical interest in obtaining accurate knowledge of the elements of hazard that must

[22] See, for instance, the strategy articulated in White House, *National Strategy for Information Sharing and Safeguarding*, or any one of the many *"post-mortem"* analyses of the performance of national security bureaucracies before the events of September 11, 2001 or other attacks or plots.

[23] Federal Bureau of Investigation, "National Crime Information Center"; Federal Bureau of Investigation, "When Off-Line is Better: Another Way to Search Crime Records."

[24] The receiving state determines which shared charges will have an effect on their drivers; in most cases, these charges include manslaughter, negligent homicide, driving under the influence of drugs or alcohol, hit-and-run, or felonies involving the use of a motor vehicle. The Compact also ensures that drivers have a single driver license and single driver record. (Pennsylvania Department of Transportation, "Driver License Compact Fact Sheet").

be considered in fixing rates".[25] And, rather than each insurance company investing in its own separate laboratory for safety testing, the industry came to the judgment that working together was a cost-effective and mutually beneficial way to secure important information based on specialized engineering talent.

Arrangements like this demonstrate the value of sharing when specialized knowledge is required, and economies of scale can be obtained through sharing both the costs and the benefits of an effort.

When Information Only Has Value if Aggregated

In the other situations discussed above, participants could have obtained the information they sought in other (though more expensive or time-consuming) ways; however, there are cases where information *only* has value if it is aggregated. In such cases, a sharing organization could be the difference between collective paralysis and collective action.

The federal Affordable Care Act (ACA) set up a risk adjustment mechanism by which money would be shifted between health insurers to account for differences in the health of their health plans' enrollees. (This was meant to encourage insurers to enroll sick people as well as healthy ones.) However, without knowledge about other insurers' enrollment characteristics, each company would be unable to properly account for the amount of the risk adjustment that would be applied to it—a critical issue in an industry with margins low enough that a large risk adjustment could wipe out an entire year's profits. With federal regulators completing a slow process that would produce results after companies would have had to set rates using their assumptions, a private actuarial firm "persuaded insurers in more than 30 states to let it act as a clearinghouse, gathering detailed information from each company, figuring how it fits together and sharing only what's necessary."[26] Even if insurers had with great effort worked out a series of pre-credit-bureau-style bilateral sharing relationships, they would have had to guess about all other insurers; only by aggregating the

[25] Brearly, "A Symbol of Safety: The Origins of Underwriters' Laboratories," 78.
[26] Hancock, "Actuaries In Denver Will Get First Peek At Obamacare's Full Cost."

data would it become sufficiently useful and trustworthy as input to an important calculation.

This theme—that aggregation produces unique insight—could be said to apply to other cases, such as the financial information sharing systems. However, cases like this exhibit it even more clearly because sharing is the *only* clear way to effectively gather the data.

1.2.3 Lessons from Other Sharing Arrangements

Each of the preceding instances of information sharing has addressed a suboptimal distribution of knowledge compared with its value, or an inability to obtain, alone, the knowledge necessary to efficiently conduct business. Cyber threat information sharing suffers from similar structural disadvantages—information possessed by each participant in a sharing scheme may be most valuable to the participants *other* than the one who currently has it, and some important conclusions may not even be reachable without the combination of knowledge contributed by a number of partners. These issues will, of course, be explored further in later chapters; for now, it suffices to note that information sharing has succeeded against a variety of similarly difficult problems. Cyber threat information sharing is therefore not an unreasonable solution on its face, and the existing attempted solutions should not be simply dismissed.

1.3 Theoretical Insights on Sharing

Existing work, in addition to contributing many existing examples of information sharing, also provides more theoretical or general insights on the risks and considerations important while considering sharing schemes.

1.3.1 Risk and Institutional Trust

The business community has long studied inter-firm collaboration, including between actors who are not in equal positions of power. A survey of this literature outlines many roles that a "go-between" can take in a mediated relationship between competitors or others who may otherwise not collaborate.[27] Especially in situations "under asymmetric dependencies"—for instance, when a corporate partner relies on the government for threat signatures[28]— trusted intermediaries can quickly build bridges and increase cooperation.[29]

1.3.2 Data Stakeholder Analysis

A major impediment to cyber threat information sharing policy innovation is the concern of the public about the repurposing of private, and sometimes intensely personal, data for computer network defense. Legislative proposals to address one of the most frequently cited impediments to cyber cooperation—the exemption of private companies' data from public disclosure under laws like the Freedom of Information Act—routinely cause fierce public battles, invoking fears of limitless sharing of citizens' private data with the Department of Homeland Security, the National Security Agency, or other agencies of the U.S. Government.[30]

Fedorowicz et al. formally define roles for the various parties in a data-sharing relationship and apply them to explain the missteps in public relations and interagency coordination made by the California Franchise Tax Board (CFTB) in an early-2000s data-mining

[27] Nooteboom, *Inter-Firm Collaboration, Learning & Networks: An integrated approach*, 116–118.

[28] *Signatures* are technical means by which an attack can be identified, like network traffic to a specific computer or the presence of a certain file. They are one of the primary mechanisms by which classic computer security products operate, though their effectiveness has been mitigated in recent years by increasingly sophisticated evasive mechanisms developed by the writers of malicious software.

[29] Nooteboom, *Inter-Firm Collaboration, Learning & Networks: An integrated approach*, 110.

[30] See, *e.g.*, Rumsey, "Troubling, Broad FOIA Exemptions Not Limited to CISPA" or Timm, "CISPA, "National Security," and the NSA's Ability to Read Your Emails." Note that such an exception already exists for information related to critical infrastructure (General Accounting Office, "Challenges for Critical Infrastructure Protection," 6).

program.[31] In brief, the CFTB sought to find those who were out of compliance with tax filing requirements by combining its own information with data sourced from public records, other state agencies, professional licensing boards, the federal government, and other data sources. The program is almost universally considered to be successful, having yielded over $4 billion in additional revenue for the state, but it raised significant issues about data privacy, the repurposing of data for new uses (for example, using licensing records for barbers to identify tax delinquents), and technical and legal safeguards over personal data.

This and similar cases suggest that, in designing effective cyber threat information sharing regimes, policymakers should carefully analyze all of the stakeholders involved—not just the agencies of government and private companies, but also the citizens who entrust their personal information to these entities. Without careful attention to privacy concerns and clear definition of the use and protection of shared data, initiatives may fail to pass public scrutiny—a necessary condition if legislators are to support them and private companies are to participate in them.

1.3.3 Voluntary Sharing versus Regulation

Under existing law, the federal government cannot compel private companies to take many, if any, actions to improve their network defense posture.[32] Since the issuance of Presidential Decision Directive 63 in 1998, most, if not all, federally sponsored cybersecurity efforts have been voluntary, likely due to this legal problem.[33] While the creation of organizations like the ISACs was strongly encouraged by presidential order and executive branch agencies are responsible for coordinating with them, companies are not required to participate. Even if a company were required to be a member of an ISAC, it could only be compelled by internal

[31] *Data-mining* refers to efforts to gain insight or derive new facts from data sources, or the fusion of data sources. Fedorowicz, Gogan, and Culnan, "Barriers to Interorganizational Information Sharing in e-Government: A Stakeholder Analysis."

[32] This issue is alluded to in Fischer et al., *The 2013 Cybersecurity Executive Order: Overview and Considerations for Congress*, which notes the limits of the President's executive power to impose new requirements on private industry. In particular, see ibid., n. 34 for reference to a White House description of Executive Order 13636 as a "down-payment on further legislation."

[33] PDD-63 will be explored in greater detail in section 2.1.

rules to actually share information; in general, no law or regulation demands such sharing. Though some industries—including chemical, electric, financial, and transportation—are subject to mandatory federal reporting, most industries are not.[34] Defense contractors are also sometimes required to report security breaches.[35]

The success of private, voluntary schemes may, then, affect whether coercive federal regulation is developed—if the industry can handle its own problems, there is no need to make the effort to pass new legislation or develop new executive policy. The Financial Services Information Sharing and Analysis Center (FS-ISAC) articulated this assessment explicitly on its website in 2013: "If the private sector does not create an effective information sharing capability, it will be regulated: this alone is reason enough to join."[36] It is notable that this assertion comes from the financial services industry, which is already one of the most highly regulated of the critical infrastructure sectors.[37] Nevertheless, the financial industry likely would prefer to avoid *still more* regulation, and private companies' incentives to avoid onerous federal regulation may therefore change their calculus about sharing.[38]

1.3.4 Theories of Cooperation

Much of the analysis put forward by economists and policy analysts assumes that the organizations involved in information sharing are monolithic, rational actors seeking to maximize corporate or public good. Such assumptions, whether implicit or explicit, do not appear

[34] Fischer et al., *The 2013 Cybersecurity Executive Order: Overview and Considerations for Congress*, p. 5 & n. 23.

[35] See later discussion on p. 21.

[36] Retrieved on 19 Nov. 2013 from Financial Services Information Sharing and Analysis Center, "FAQs — FS-ISAC." This text has been subsequently modified to concentrate more on the benefits of sharing instead of the costs of *not* sharing.

[37] The "banking and finance" industry accounted for fully half (17) of the laws, regulations, and mandatory standards related to privately owned information technology systems for critical infrastructure operators in U.S. Government Accountability Office, "Information Technology: Federal Laws, Regulations, and Mandatory Standards for Securing Private Sector Information Technology Systems and Data in Critical Infrastructure Sectors," 3, 12.

[38] The rejoinder that companies might simply decide not to look for any cybersecurity problems to avoid penalties or government scrutiny—"we have no problems to report, so there's no need to regulate us"—is easily dismissed by the significant costs faced by many large companies in the wake of incidents that were exposed by third parties, not through internal investigation.

valid in the cyber-threat landscape. Theories of cooperation drawn from various disciplines can therefore extend and enrich this overly simplistic understanding.

An organizational frame for cyber threat information sharing already exists in many cases—the first ISACs were formed over a decade ago, and many units of the Department of Homeland Security are eager to receive and share information on critical infrastructure protection.[39] However, even the existence of these institutions has apparently not created the necessary conditions for mutually beneficial cooperation; many ISACs are reputed to accomplish little, and economic analyses show that their structure may encourage undesirable equilibria.[40] So, an understanding of the ways that cooperation can be brought about through interaction and confidence-building measures could "tip the scales" and overcome existing barriers to institutional success.

Robert Axelrod's *The Evolution of Cooperation* draws insights from biology, computer science, and game theory to give more general conclusions on cooperation. Axelrod's findings are centered around the success of cooperation in *iterated games*—a theoretical abstraction of repeated interactions between actors. His work highlights the centrality of confidence-building reciprocity for the development of continued collaboration between actors, even those who may not initially—or ever—fully trust each other.[41] Other compelling lines of research indicate that "jump-starting" collaboration is possible with even just a small proportion of the actors in a group[42] and that repeated interaction of *specific* individuals is critical to successful collaboration.[43]

Jeffrey Legro's *Cooperation Under Fire* advances an organizational-culture theory of cooperation which appears to more closely match the reality of how cyber threat information is generated and shared.[44] The interaction between companies and the U.S. Government

[39] U.S. Department of Homeland Security, *Information Sharing: A Vital Resource for a Shared National Mission to Protect Critical Infrastructure*.

[40] See detailed discussion in subsection 2.3.3, and Mimoso, "Adequate Attack Data and Threat Information Sharing No Longer a Luxury."

[41] Axelrod, *The evolution of cooperation*, 173.

[42] Ibid., 175.

[43] Ibid., 180.

[44] Legro, *Cooperation under fire: Anglo-German restraint during World War II*.

can be better understood in a way similar to Legro's separation of a state into component bureaucracies with unique cultures.

Only a small component of a company, utility, or non-profit institution ever will, in most cases, be involved in the direct response to a cyber threat or intrusion, and the culture of that subcomponent may practically determine the entire organization's response. For many companies, this group will be part of the Information Technology (IT) staff, headed by a Chief Information Security Officer or a similar executive. While the chief executive or board of a company may sign off on sharing with the government, it is likely that their decision will be based on the culture and beliefs of the organization's security bureaucracy.[45]

Similarly, the federal government should, in practice, not be viewed as a single entity to which threat information can be passed and from which assistance or information can be requested. The constellation of federal agencies responsible for computer network defense includes some responsible primarily for law enforcement (*e.g.*, the FBI, the U.S. Secret Service, or the Defense Cyber Crime Center [DC3]); some responsible only for the defense of classified or military systems (the National Security Agency/Central Security Service Threat Operations Center [NTOC]); and a number of organizations within the Department of Homeland Security that respond to different threats based on the business sector involved or whether the target is part of a critical infrastructure system.[46] This can be maddeningly complex, as illustrated in Figure 1.1.[47]

Dhillon and Backhouse present a comprehensive survey of culture-informed analysis of information systems in a 2001 paper.[48] While such discussion does not directly contribute

[45] Such beliefs might include personal views of DHS's competency in responding to cyber intrusion (*e.g.*, in Google's collaboration with the National Security Agency following the "Aurora" intrusion; see Rosenzweig, "Public-Private Partnerships for Cybersecurity Information Sharing," 1–2) or of an ISAC's ability to provide useful assistance.

[46] Office of Inspector General. U.S. Department of Homeland Security, "DHS' Efforts to Coordinate the Activities of Federal Cyber Operations Centers."

[47] For further examples of the amazingly diffuse and complex distribution of federal cybersecurity responsibility, see Rosenzweig, "Public-Private Partnerships for Cybersecurity Information Sharing"; Booz Allen Hamilton, *National Cybersecurity Center Policy Capture*; and Booz Allen Hamilton, *Commercial/Civil Cyber Community Snapshot* (reproduced here in Figure 1.1). The last two documents were part of the White House's Cyberspace Policy Review: `http://www.whitehouse.gov/cyberreview/documents`.

[48] The term "information systems" almost invariably points to such a style of analysis; the terminology is

to the problem of cybersecurity sharing, it supports the assertion that computer systems do not exist in a vacuum—rather, their development and security are dependent on the culture of the organization they serve, consistent with Legro's intuition about cultural effects on observable behaviors.

1.4 The Way Forward

The idea that information sharing would improve network defense efforts is not unreasonable; calls by prominent officials and industry leaders, examples from other problem spaces, and theoretical reasons for its promise were presented above in Chapter 1. In reality, sharing is executed on a wider scale than is implied by most literature;[49] a broader variety of organizations, companies, and networks by which sharing is effected is described in Chapter 2, followed by a distillation of the findings of existing analyses of their relative lack of success. A better analysis technique is needed to effectively address these issues; the *Computational Policy* abstraction and other theoretical understandings of the computer security problem are presented in Chapter 3. Chapter 4 presents an application of this analysis lens to the cybersecurity problem at hand before Chapter 5 concludes.

almost never used in computer science. Dhillon and Backhouse, "Current directions in IS security research: towards socio-organizational perspectives."

[49] Such work tends to concentrate on the "Information Sharing & Analysis Centers" (ISACs) formed by critical infrastructure industries in response to Presidential Decision Directive 63 (PDD-63). These organizations and this policy will be explored in later chapters.

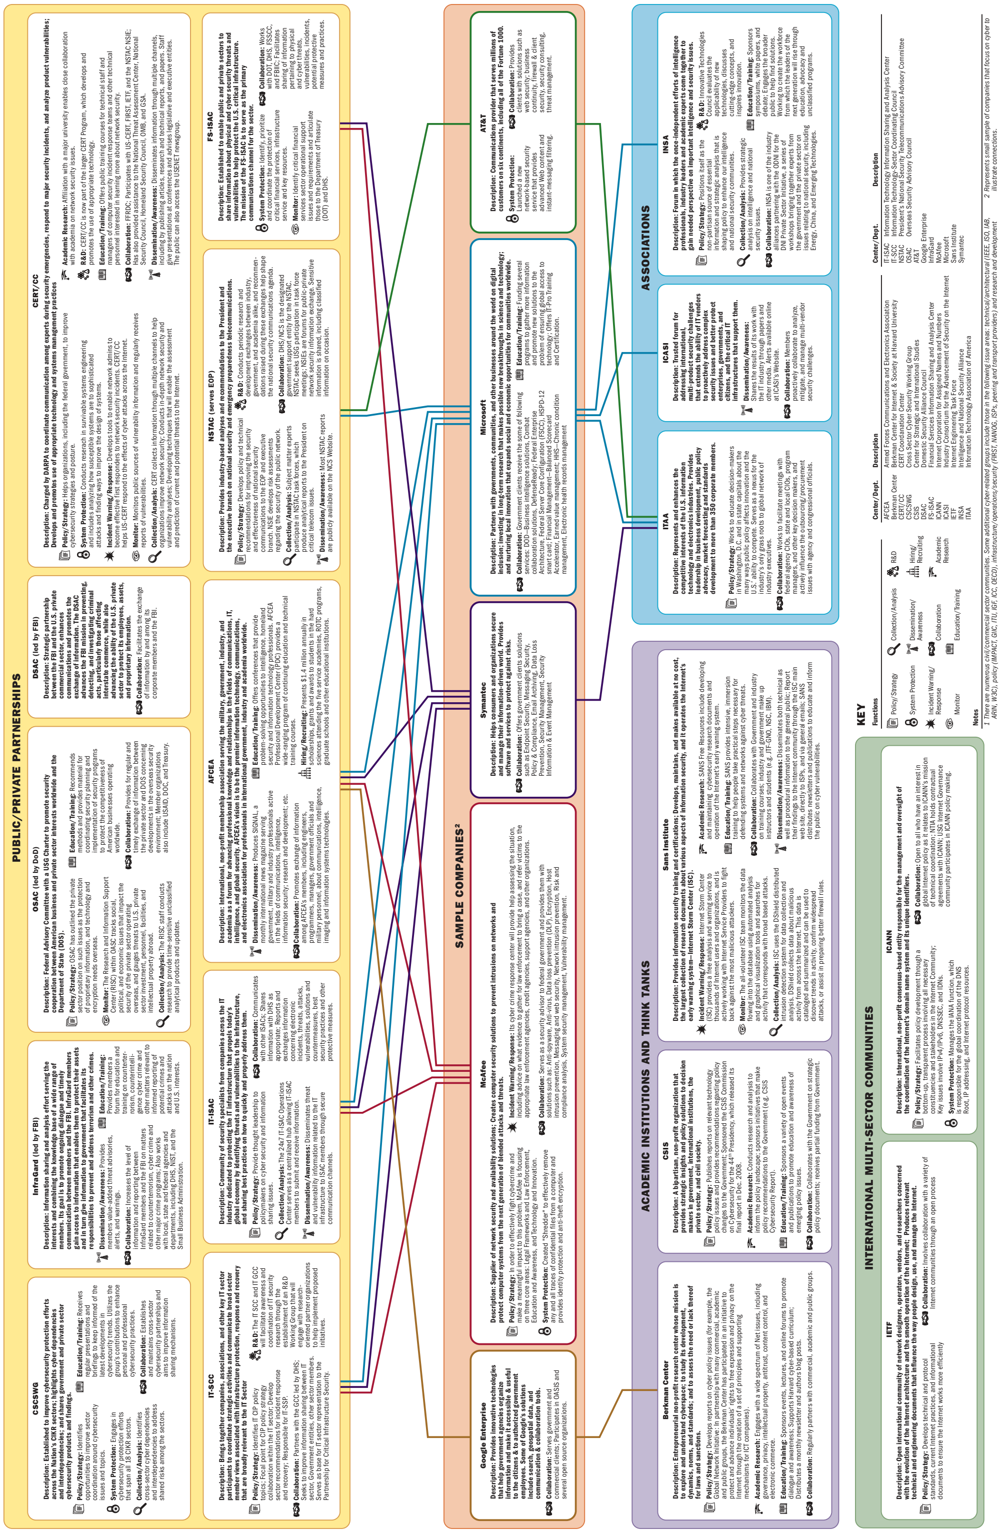

Figure 1.1: Cybersecurity Policy "Snapshot"
A document submitted to the White House's 2009 Cyberspace Policy Review, reproduced here to illustrate the maddening complexity of federal cybersecurity policy. (Booz Allen Hamilton, *National Cybersecurity Center Policy Capture*)

Chapter 2

Current Initiatives and Analyses

> "Instead of just building better defenses, we must build better relationships. And we must overcome the obstacles that prevent us from sharing information and, most importantly, collaborating."
>
> Robert S. Mueller, III, then FBI Director
>
> (Feb. 28, 2013: Mueller, "Working Together to Defeat Cyber Threats")

This chapter presents a history of cyber threat information sharing efforts, both government-sponsored and privately organized. Because cyber threat information sharing is a relatively unfamiliar topic for most readers, this chapter presents a broader survey of sharing efforts than will be necessary for the ensuing analysis. Information sharing is broadly defined—and a popular goal to claim that one's organization is accomplishing. This overview covers the most important types of sharing relationships and structures, including some that are usually forgotten. Following this overview, a summary of existing analyses of cyber threat information sharing is presented.

2.1 Government-Sponsored Sharing Efforts

Formal cyber threat information sharing began in earnest with Presidential Decision Directive 63 (PDD-63), issued in 1998 by President Clinton in response to concerns over the nation's increasing reliance on complex privately owned cyber-physical systems for critical infrastructure.[1] A number of policy initiatives and legislative agendas have attempted to change the sharing landscape since PDD-63. Executive actions have been incremental, often responding more to changes in the government's bureaucratic structure rather than the ever-evolving threat landscape. Legislative efforts have been audacious but abortive, often falling prey to concerns over civil liberties, government overreach, or the imposition of costs on businesses. But, promising efforts are emerging in the Defense Industrial Base and the critical infrastructure community, and with the development of voluntary standards.

2.1.1 Executive Strategy for Critical Infrastructure

PDD-63 encouraged the creation of private-sector organizations to share information relating to physical and cyber security of critical infrastructure, leading to the establishment of industry-based Information Sharing and Analysis Centers (ISACs). The directive designated sector-specific agencies (SSAs) which would take "clear accountability" for their sector and directed the appointment of one private-sector Sector Coordinator for each industry who would work with a single government Sector Liaison Official to assess, plan, and coordinate security measures.[2] The ISACs that emerged are still among the primary formal mechanisms by which physical and cyber threat information is shared among companies and with the federal government. There are now at least 16 ISACs, most founded in the early 2000s, but some later (like the Oil and Natural Gas ISAC [ONG-ISAC], which was founded in 2014).[3] Beyond individual ISACs, the ISAC ecosystem also includes the National Council of ISACs,

[1] *Cyber-physical systems* are "physical and engineered systems whose operations are monitored, coordinated, controlled and integrated by a computing and communication core" (Rajkumar et al., "Cyber-physical Systems: The Next Computing Revolution"). For PDD-63, see: Moteff, "Critical Infrastructures: Background and Early Implementation of PDD-63"; "Presidential Decision Directive 63: Critical Infrastructure Protection."

[2] "Presidential Decision Directive 63: Critical Infrastructure Protection," Section IV & Annex B.

[3] ISACs now exist in the following industries: communications; electricity; emergency services; financial services; health; information technology; maritime security; "multi-state" (state, local, tribal, and territorial governments); nuclear energy; public transit; real estate; research and education; supply chain; surface transportation; water; and oil & natural gas, according to the National Council of ISACs (National Council of ISACs, "Member ISACs"). The ONG-ISAC does not list a creation date, but its domain name (`ongisac.org`) was registered on January 24, 2014, and it did not appear on the National Council website in November 2013.

which meets monthly and attempts to bridge the gaps between its member ISACs, since each only deals with a single sector. The ISAC ecosystem is not uniform: the amount of funding and its source, the level of government involvement, and the level of cybersecurity success varies significantly between ISACs.[4]

Each successive presidential administration has issued a new policy on critical infrastructure protection. President Bush issued Homeland Security Presidential Directive 7 (HSPD-7), which partially superseded PDD-63, in December 2003.[5] HSPD-7 did not make significant changes to the structure set up by PDD-63, but did bring the directive up to date following the passage of the USA PATRIOT Act of 2001 and the Homeland Security Act of 2002, and gave overarching responsibility for critical infrastructure security to the new Department of Homeland Security. In February 2013, President Obama in turn superseded HSPD-7 with Presidential Policy Directive 21 (PPD-21). As before, PPD-21 is very similar to HSPD-7, but also specifically encourages the development of information systems, interoperable formats, and redundant systems, demonstrating greater understanding of the logistical problems faced "in the field" 16 years after PDD-63's issuance.[6] Despite these specific differences, the presidential policy structure for information sharing remains largely the same as it was in 1998.

2.1.2 Programs in the Defense Industrial Base

While ISACs have remained somewhat static, the Department of Defense has since about 2012 been moving quickly—by regulatory standards at least—to establish sharing programs with its contractors, which are collectively called the Defense Industrial Base (DIB).[7] In

[4] General Accounting Office, "Improving Information Sharing with Infrastructure Sectors," 7–10.

[5] "Homeland Security Presidential Directive 7: Critical Infrastructure Identification, Prioritization, and Protection."

[6] "Presidential Policy Directive 21: Critical Infrastructure Security and Resilience," "Three Strategic Imperatives" #2.

[7] This renewed push comes in the wake of high-profile computer security breaches of trusted defense contractors including Lockheed Martin (see Hodge and Sherr, "Lockheed Martin Hit By Security Breach") and Booz Allen Hamilton (see R.L.G., "Hackers strike at a foe"), and a growing realization that contractor networks filled with government information are targets just as government networks are (see Lee and

particular, the DIB Enhanced Cybersecurity Services (DECS) program provides automated sharing of unclassified and classified threat indicators and intelligence among its partners.

However, while this program is nominally bidirectional, it seems structured mainly to facilitate the unidirectional sharing of information *from* the government *to* industry, and not the reverse. Regulations require that participating companies report any compromise of defense information to the Defense Cyber Crime Center (DC3) within 72 hours of the incident's discovery, but makes further sharing (including potentially the most useful data: "forensic analyses, mitigation and remediation, and . . . damage assessments") voluntary.[8] And, while the regulation allows the redistribution of anonymized versions of this information with other DIB participants, it is not clear that the government will be able to do so quickly and effectively.[9]

Reports indicate that current signatures distributed through the program may be of little use to sophisticated network defense teams—a number of early participants appear to have chosen to leave the program and reallocate resources to other security initiatives, though accounts are inconsistent.[10] Even so, regardless of the program's overall enrollment trends, it still has the potential to raise the level of security at the worst-prepared companies since these companies might derive more new, valuable information from the shared signatures and countermeasures.

The DECS program was expanded in 2012: it was moved to a new primary home in the Department of Homeland Security (DHS) as simply the Enhanced Cybersecurity Services program.[11] It is now being expanded to all companies defined as "critical infrastructure"

Schlanger, "Department of Defense Expands Defense Industrial Base Voluntary Cybersecurity Information Sharing Activities," 1).

[8] Office of the Secretary of Defense, "Cyber security information sharing," Sec. 236.5b.

[9] See further discussion later in this chapter regarding complaints from government partners about the government's reluctance or inability to release useful information, for instance in the Conficker case in subsection 2.2.2.

[10] Lim, "Pentagon Cyber-Threat Sharing Program Lost Participants"; Reed, "DoD-DHS' info sharing program on cyber threats isn't shrinking (updated)."

[11] Note the loss of the "DIB" prefix. The program was also called the Joint Cybersecurity Services Program (JCSP) before the transition began.

by the Secretary of Homeland Security in accordance with Executive Order 13636.[12] The bureaucratic and legal structure for this sharing is somewhat convoluted, requiring customer companies to purchase ECS from trusted Commercial Service Providers who receive the threat information from DHS. Operationally, this is implemented by the distribution (by DHS) of threat signatures derived from analysis by the National Security Agency/Central Security Service (NSA/CSS) Threat Operations Center (NTOC), the DHS Office of Cybersecurity and Communications (CS&C), and other government agencies.[13] The signatures specify malicious addresses, programs, or types of traffic to block, and can prescribe limited countermeasures to disrupt adversarial behavior.[14] Signatures are passed to Commercial Service Providers—telecommunications companies with special security clearance to receive and redistribute this information—and then provided or sold to participating companies.

The program supports limited two-way sharing: the providers may, with permission, report "limited, anonymized, and aggregated" metrics back to DHS to allow DHS to ascertain the effectiveness of indicators. DHS indicates that it shares this information with other U.S. Government agencies that have cybersecurity responsibilities.[15] While the program is currently only open to critical infrastructure companies, recent regulatory actions suggest that the government is exploring possibilities for more universal participation in government-sponsored cybersecurity efforts that previously only targeted critical infrastructure.[16]

[12] "Executive Order No. 13636: Improving Critical Infrastructure Cybersecurity," Sec. 4. The primary government point of contact for DIB companies will remain the Department of Defense. (Defense Cyber Crime Center, "DIB Enhanced Cybersecurity Services (DECS)").

[13] Fischer et al., *The 2013 Cybersecurity Executive Order: Overview and Considerations for Congress*; U.S. Department of Homeland Security, *Privacy Impact Assessment for the Enhanced Cybersecurity Services (ECS)*.

[14] These countermeasures are, though the system is still being developed at the time of writing, limited to relatively low-hanging fruit: Domain Name System (DNS) sink-holing, which sends traffic destined for malicious servers to a "sink- hole" instead of allowing it to reach its malicious target; and email filtering, which would automatically intercept and quarantine or block malicious attachments. The Department plans to increase deployed capabilities as the system develops. See U.S. Department of Homeland Security, *Privacy Impact Assessment for the Enhanced Cybersecurity Services (ECS)*.

[15] Ibid., 5.

[16] U.S. Department of Commerce, "Incentives To Adopt Improved Cybersecurity Practices."

2.1.3 Law Enforcement-Based Efforts

Cybercrime and other computer-related law enforcement efforts often employ some type of information sharing. The FBI's InfraGard program and the National Cyber-Forensics and Training Alliance (NCFTA) are often praised in public speeches and news releases, but it is difficult to ascertain how successful they are from external or academic sources, how they function, or what type of information is routinely shared among members.[17] The websites of some of the nearly 100 InfraGard chapters show troubling signs of inactivity among members and a lack of focus on cyber threat information sharing among the many other potential threats facing critical infrastructure operators.[18] With overall reported membership numbering in the tens of thousands, it is likely that many participants are of relatively low value to national security or law enforcement.[19] However, the existence of the partnership still may benefit participants and the government by providing a clear way to report crime or security issues. The NCFTA similarly appears to consist of a small number of FBI agents augmented by relationships with important partners, some of whom routinely share with each other. It is unclear what role InfraGard and NCFTA play in operational computer network defense activities compared with law enforcement investigations, and this thesis does not attempt to rigorously assess them.

[17] Illustrative references include: Federal Bureau of Investigation, "InfraGard: A Partnership That Works"; Federal Bureau of Investigation, "The NCFTA: Combining Forces to Fight Cyber Crime"; Mueller, "Working Together to Defeat Cyber Threats."

[18] The websites examined were: San Francisco Bay Area InfraGard Chapter, "SF Bay InfraGard Meetings"—showing quarterly meetings with approximately half devoted to computer security; InfraGard National Capital Region Members Alliance, "Washington, D.C. – National Capital Region Members Alliance Chapter Website"—which seems similarly focused on other issues, but with a Cyber Special Interest Group that meets monthly; and North Texas InfraGard Members Alliance, "InfraGard – North Texas Chapter Events Home Page"—featuring a desperate plea for members to log in at least once per year and pay $20 dues.

[19] Membership totals are obtained from Federal Bureau of Investigation, "InfraGard: A Partnership That Works." With only one or two special agents assigned to most chapters, the FBI clearly must not be handling frequent reports from even a significant fraction of the InfraGard members.

2.1.4 Enduring Security Framework

The federal government also has the ability to bring a variety of players to the table to solve long-term, strategic cybersecurity problems. An effort called the Enduring Security Framework, while more accurately termed "joint problem solving" than "information sharing", nevertheless has reportedly addressed some critical, systemic vulnerabilities in U.S. infrastructure.[20] And, news reports indicate that computer manufacturers were only prompted to implement reasonable security measures to ensure the integrity of computers' Basic Input/Output System (BIOS)—security-critical code that initializes hardware when a computer boots—by a government briefing that "scared the bejeezus out of them".[21] While some dispute the level of vulnerability addressed in that case or the unique position of the government to address it,[22] the power of the Framework to "identify critical cyber vulnerabilities and mobilize experts to address the risks" suggest that such efforts may play an important role in strategic, long-term information sharing efforts to improve computer security.[23]

2.1.5 Voluntary Efforts and Research

Executive Order 13636 also directed the National Institute of Standards and Technology (NIST) to develop a "Cybersecurity Framework" for voluntary adoption by U.S. industry.[24] This framework was released in February 2014, and the Department of Commerce (NIST's cabinet-level parent agency) published an inquiry soliciting input from the public on how to incentivize adoption of the voluntary framework.[25]

The NIST Cybersecurity Framework offers a promising avenue for government lead-

[20] Mueller, "Working Together to Defeat Cyber Threats."

[21] Gjelten, "Cyber Briefings 'Scare the Bejeezus' Out Of CEOs."

[22] Sometimes in spectacularly strident form, for instance, Graham, "How we know the 60 Minutes NSA interview was crap."

[23] Coviello, "Written Testimony to the U.S. Senate Committee on Commerce, Science, & Transportation," 7.

[24] "Executive Order No. 13636: Improving Critical Infrastructure Cybersecurity," Sec. 7.

[25] National Institute of Standards and Technology, "Framework for Improving Critical Infrastructure Cybersecurity, Version 1.0"; U.S. Department of Commerce, "Incentives To Adopt Improved Cybersecurity Practices."

ership on cybersecurity, but does not bear the force of law or regulation. This may not relegate it to perpetual obscurity, though, since government standards commonly become a mark of approval with which products and services are marketed. Many of the most popular and well-regarded encryption products, for instance, are those issued as standards by NIST through a collaborative, open process.[26] In addition, there is precedent for "voluntary" actions becoming mandatory for companies involved in government contracting, so this voluntary framework may be the model for much of the federal government's cybersecurity and information sharing policy in the future.[27]

2.2 Private Sharing Efforts

Private industry actors, including those not covered under the federal government's definition of "critical infrastructure", have recognized the benefits of sharing in situations in which existing social networks can be tapped for sharing or in which significant technical expertise is required. Some eschew formal coordination mechanisms for more flexible, trust-based collaboration, to great observed success. And, some collaborate purely on a technical basis through contractual service provider relationships.

2.2.1 Ad hoc Collaboration Between Victims

A representative and public example is the Silicon Valley technology industry's response to a coordinated attack on their networks. The compromise was a "watering hole" attack, meaning that the attacker compromised a website that many targets visited; one compromised website therefore was a convenient, semi-targeted way to infect users in many distinct companies that shared an interest in mobile application development. Facebook, the first

[26] Such standards include the Data Encryption Standard (DES), Triple-DES (3-DES), Advanced Encryption Standard (AES), the Digital Signature Algorithm (DSA), and the Secure Hash Algorithm (SHA) family of cryptographic hash functions. These are the most commonly used cryptographic functions in use today.

[27] See, for instance, the False Claims Act disclosure that is voluntary under that law but is made mandatory for contractors under the Federal Acquisition Regulation (Lee and Schlanger, "Department of Defense Expands Defense Industrial Base Voluntary Cybersecurity Information Sharing Activities," 5).

company to discover the intrusion, immediately contacted a number of its industry peers—including its competitors—and these companies continually shared technical information throughout their response to the incident.[28] Existing social networks likely facilitated this sharing; organizations like the Bay Area CSO Council facilitate professional contact among security executives in related companies, and "active participation" is required for membership.[29] The group is credited with facilitating sharing before and especially in the wake of the massive 2009 "Aurora" attack on Google and many Bay Area companies.[30]

2.2.2 Ad hoc Technical Working Groups

The Conficker Working Group—a collaboration of security researchers, technology companies, and others—was hastily organized to tackle a particularly worrisome piece of malicious software in 2008. Conficker was a notably virulent and resilient threat that eventually infected millions of computers, forming a massive "botnet" that could be used to overwhelm websites, launch other attacks, send spam, or harvest private information. The worst fears about Conficker's potential for harm were not realized, in large part due to the coordinated response by the security community. The working group demonstrated that ad hoc, mostly private collaborations could be effective in countering significant generic cybersecurity threats.[31] In fact, the group's example was so compelling that the Department of Homeland Security commissioned a study to analyze the reasons it functioned so effectively. The prominent example of the Conficker Working Group shows the promise of technical, issue-based collaboration for dealing with significant cyber-threats, and, as with Aurora and

[28] Facebook Security Team, "Protecting People On Facebook."

[29] CSO is an acronym for Chief Security Officer (see Appendix A). The organization's name does not expand the acronym. This executive position is also sometimes called the Chief Information Security Officer. In some companies, there is no separate CSO, and such functions are carried out by the Chief Information Officer (CIO).

[30] Higgins, "'Operation Aurora' Changing the Role of the CISO." Note that Google also collaborated with the NSA on its response (Nakashima, "Google to enlist NSA to help it ward off cyberattacks," cited in Rosenzweig, "Public-Private Partnerships for Cybersecurity Information Sharing," 1–2), but that NSA was not meant to serve as a hub for collaboration with other private companies.

[31] The *generic* nature of the threat should be emphasized. Since the attack affected many, many companies, government agencies, and individuals, there was little or no competitive risk to admitting infection or working to stop the botnet. See further discussion about the types of computer network attacks in section 3.2.

the Facebook watering hole attack, was enabled by social networks of familiarity and trust that existed before the response began.[32]

However, this effort was not an unqualified success. Participants complained of unidirectional sharing—though, unlike in other cases, this time the private sector was sharing and the government was reportedly "[leeching] off the process".[33] And, the group eventually had trouble deciding how to divide into subgroups and how to distribute any decision-making authority. More frighteningly, as the main email list's size reached over 300, the trust implicit in its original design (from either direct personal experience or based on the assertion of a mutual friend) broke down and some participants even wondered if Conficker's malicious author might be listening to their communications.[34]

2.2.3 Computer Security Vendors

Computer security companies also form a bridge between traditional "islands of responsibility"—companies, agencies, and institutions that would otherwise not be concerned with each other. Antivirus companies use their customers to form distributed networks, sometimes literally calling them "intelligence networks", through which to detect threats and gain global situational awareness.[35] In practice, malicious executables found in one network are discovered and analyzed, then protections are distributed to all other customers—a perfect sharing scenario. An especially desirable quality of this arrangement is that the security vendor has no obvious way of sharing the identity of the company or network which initiated its analysis, so

[32] The Rendon Group, *Conficker Working Group: Lessons Learned*, 28.

[33] Ibid., 40. This was particularly upsetting to some researchers who found that their work had been briefed within the government with no attribution of its source. Such concerns shouldn't be dismissed as simple pettiness—they could certainly sour a relationship and inhibit further sharing, especially once a crisis event is over. And, the report notes that un-cited repetition of findings could cause researchers to incorrectly believe that their work had been independently verified, which could cause significant problems in a fast-moving, precise response effort.

[34] Ibid., 21.

[35] For example, Symantec boasts of a "Global Intelligence Network" (`http://www.symantec.com/deepsight-products`). Symantec also provides academic access to the "Worldwide Intelligence Network Environment", which includes millions of binaries and information on when and where they were discovered (Shou and Dumitraş, "Worldwide Intelligence Network Environment (WINE): Symantec's Data Sharing Environment," 5).

"reports" are inherently anonymized.[36] Even competing antivirus companies share samples of virulent malware, and threats often are eventually detected by all major products.[37] In recent years, highly sophisticated—and probably state-sponsored—actors have been foiled by these companies without the use of information from government or classified sources.[38] Perhaps the best known recent example is the discovery and analysis of the Stuxnet family of malware, a "brilliantly executed" suspected state-sponsored attack on Iranian nuclear enrichment, which involved a number of security companies and "extraordinary" sharing among them.[39] Security companies like Symantec and Mandiant have faced down the "Elderwood Gang" and "Advanced Persistent Threat 1", respectively, both suspected Chinese computer network exploitation teams.[40] Since their business model relies on knowing how to identify and counter as many threats as possible, these companies have very real economic incentivizes to spread threat information as widely as possible once they obtain it. Security companies therefore form an effective information sharing network that transcends industrial and even national boundaries.

Such incentives also extend to open-source and public computer security efforts like the Snort intrusion detection and prevention system. Snort is an intrusion detection product, but also features "rulesets"—sets of patterns to block known or suspected malicious activity. In a way similar to antivirus updates, such rulesets can be automatically updated on deployed devices, providing more comprehensive defense over time. Different ruleset versions are available under open-source licenses and paid licenses, demonstrating the collaboration of a lead company (Sourcefire) and a community of independent contributors.[41]

[36] There is a slight risk that an attacker could detect which security products a target is running by observing whether their malicious software is later released as a signature by various vendors. However, this risk is rather remote and it is unclear that it is serious enough to inhibit any sharing.

[37] Kaspersky, "The contemporary antivirus industry and its problems."

[38] See, for instance, Mandiant, "APT1: Exposing One of China's Cyber Espionage Units," n. 3.

[39] Kushner, "The Real Story of Stuxnet."

[40] Mandiant, "APT1: Exposing One of China's Cyber Espionage Units"; O'Gorman and McDonald, "The Elderwood Project."

[41] See `http://snort.org/snort-rules/`.

2.3 Existing Analyses

Cyber threat information sharing is a relatively understudied area in the cybersecurity policy discourse: in a recent 58-page comprehensive bibliography on cyber conflict issues, fewer than ten entries even considered sharing, and only three of these even briefly considered information sharing between entities other than national governments.[42]

The literature on cyber threat information sharing is mostly found in economic analysis, policy analysis, trade press, and investigative reports. These sources identify three main obstacles to successful cyber threat information sharing: legal uncertainty, reputational risk, and insufficient incentives. While some of these sources explicitly refer to the design or implementation of ISACs, many of the conclusions would apply to any similarly structured sharing organization. In addition, many of the risks and benefits identified apply to any kind of sharing. This section considers each of the major objections in turn and argues that they could be overcome with more effective sharing strategies.

2.3.1 Legal Uncertainty

Legal uncertainty for participants in information sharing organizations has been recognized since Presidential Decision Directive 63 (PDD-63)'s inception—the Directive commissioned studies on issues including "liability issues arising from participation by private sector companies in the information sharing process" and "existing legal impediments to information sharing".[43] The issue is still present today.

Potential information sharers have expressed concerns over legal liability, worrying that they might inadvertently distribute information that they are not legally permitted to share. For instance, sharing network traffic might plausibly be prohibited by the Electronic Communications Privacy Act (ECPA), its state or local equivalents, or other privacy-

[42] Fossum, *Cyber Conflict Bibliography*. The relevant entries are: The Netherlands (2011); European Network and Information Security Agency (ENISA) (2012); and Yannakogeorgos, P. A. & Lowther, A. (2013).

[43] "Presidential Decision Directive 63: Critical Infrastructure Protection," Annex B.

protection laws. However, Paul Rosenzweig calls assertions of uncertainty due to such laws "overblown", noting the broad exemptions for the protection of a service provider's interests and the ability of a company to simply obtain the consent of its customers before sharing, for instance through service contracts.[44]

Companies also have expressed concerns that, by sharing computer security information with partners, they could violate federal antitrust law. The antitrust issue has been brought up in various recent contexts, including the widely reported recent release of a favorable joint opinion from the Federal Trade Commission and the Justice Department's Antitrust Division.[45] However, this objection has long been dismissed as a "canard";[46] the recent opinion directly cites and reaffirms legal opinions issued in 2000.[47] Similarly, Rosenzweig rejects antitrust concerns, noting that legal interpretation "first announced 100 years ago" restricts antitrust prohibitions on information sharing only to that which is an "'unreasonable' restraint on trade".[48]

2.3.2 Risk of Disclosure

Participants in information sharing organizations may reveal information about vulnerabilities in their software products, breaches of their corporate networks, or other sensitive or embarrassing topics. PDD-63 called for, along with the studies mentioned above, investigation into "the methods and information systems by which threat and vulnerability information can be shared securely while avoiding disclosure or unacceptable risk of disclosure to those who will misuse it," and "the improved protection. . . of industry trade secrets and confidential business data,. . . material disclosing vulnerabilities of privately owned infrastructures and apparently innocuous information that, in the aggregate, it is unwise to

[44] Rosenzweig, "Cybersecurity and Public Goods: The Public/Private 'Partnership'," 14–16.

[45] For example, the antitrust "news" was reported as a totally new development in Fung, "Washington is making is easier for businesses to swap notes on hackers."

[46] Brenner, "Cyber Threat Information and the Antitrust Canard."

[47] Department of Justice and Federal Trade Commission, "Antitrust Policy Statement on Sharing of Cybersecurity Information," 4.

[48] Rosenzweig, "Cybersecurity and Public Goods: The Public/Private 'Partnership'," 16, n. 74.

disclose."[49]

Without assurance that they will remain anonymous and their secrets will be protected, institutions may simply opt not to share. Gal-Or and Ghose note that information sharing brings both costs and benefits to the revealing firm, and that the most significant risk comes from the potential embarrassment that would result from unauthorized disclosure.[50] Companies do not perceive the risk of such reputational damage to be hypothetical or minor: Gal-Or and Ghose's study reports that a software company completely severed ties with the Computer Emergency Response Team (CERT), a federally funded research and development center at Carnegie Mellon University, over allegations that CERT improperly passed vulnerability information to third parties.[51] In a time of heightened awareness of insider threats in the wake of the unauthorized disclosures by Edward Snowden, such fears may take on an extra salience.

Public knowledge of a network security breach can also have significant negative repercussions for companies. Data breaches make headlines in popular press outlets, and studies indicate that the loss of current and potential customers is now the predominant cost associated with data breaches.[52]

On this count, too, it appears difficult to sufficiently reassure industry players. Twelve years ago, due in part to ISAC complaints about the possibility that their information could be released under public disclosure laws, the Homeland Security Act of 2002 specifically protected relevant critical infrastructure information with "information use and disclosure restrictions."[53] Nevertheless, industry players seem to still be preoccupied by the possibility of information disclosure.

[49] "Presidential Decision Directive 63: Critical Infrastructure Protection," Annex B.

[50] Gal-Or and Ghose, "The Economic Incentives for Sharing Security Information."

[51] Ibid., 187.

[52] Gow, "Data security breaches: More reach and frequency requires more diligence."

[53] General Accounting Office, "Challenges for Critical Infrastructure Protection," 6, 58.

2.3.3 Insufficient Incentives

Companies, pursuant to their duty to maximize shareholder value or their private motivations to obtain profits, must evaluate the costs and benefits before participating in information sharing regimes. This, combined with the sheer number of separate networks and actors involved, suggests that a unique market dynamic will develop. For these reasons, a small group of economists have studied existing cyber threat information sharing mechanisms to discover the economic incentives—both for and against sharing—that they create.

The findings of these separate analyses agree: information sharing organizations, as they are organized today, do not economically incentivize sharing. Gordon, Loeb, and Lucyshyn find, encouragingly, that sharing can decrease the total cost of information security and lead to better outcomes. But, they conclude that existing incentives are not adequate to achieve benefits at a per-firm or society-wide level.[54]

Again, PDD-63 remains strikingly current; though its call was not limited to the problem of promoting information sharing, the Directive identified the insufficient incentives for critical infrastructure providers to secure themselves. Among the commissioned studies was one that was to address "the potential benefit to security standards of mandating, subsidizing, or otherwise assisting in the provision of insurance for selected critical infrastructure providers."[55]

Gordon et al.'s research also indicates that characteristics of the firm receiving the threat information affects the balance of incentives, with firms that have more inherent ability to respond to shared information incentivized to invest less and share less with others. Hausken presents little information not found in the previous studies, but builds on Gordon et al. to assert that information sharing is positively correlated with interdependence between firms. This confirms the implicit intuition of current policy that different indus-

[54] Gordon, Loeb, and Lucyshyn, "Sharing information on computer systems security: An economic analysis."

[55] "Presidential Decision Directive 63: Critical Infrastructure Protection," Annex B.

tries may require tailored solutions.[56] While the problems articulated by each researcher led them to similar conclusions—that the ISAC mechanism does not sufficiently incentivize self-sustaining and mutually beneficial behavior—these weaknesses may not be fatal: the Financial Services ISAC (FS-ISAC) is reputed to be one of the most effective sharing organizations,[57] even though a specific study of its structure also found that the equilibrium outcomes would be the absence of sharing, or pervasive attempts to free-ride on the disclosures of others.[58]

2.3.4 A Question of Cost or Benefit?

Many of these issues have been raised in one form or another since 1998, but are *still* brought up today. Is this due to the policy community's total inability to solve them adequately, or does it suggest that something else might be the true reason companies are reluctant to share?

The following chapters argue that these issues of potential *cost* are actually less important than the inadequate perceived *benefits* such sharing would provide. Rosenzweig concludes, "In the end, what really restricts cooperation are the inherent caution of lawyers who do not wish to push the envelope of legal authority and/or policy and economic factors such as proprietary self-interest that limit the desire to cooperate."[59] If they see insufficient potential benefit from sharing, companies have little incentive not to inflate potential risks—almost no risk, however low, might be considered acceptable. There will certainly be no strong counterargument when the institution's counsel or others warn of the potential downsides of sharing.

[56] Hausken, "Information sharing among firms and cyber attacks." As an illustrative example of industries with different levels of interdependence, compare the electricity industry with the water industry: while the electric power grid is highly interconnected and cascading blackouts are all too common, water is almost always distributed in separate systems for each municipality or local political subdivision.

[57] Mimoso, "Defenders Still Chasing Adequate Threat Intelligence Sharing"; PricewaterhouseCoopers LLP, *Key findings from the 2013 US State of Cybercrime Survey*, 7.

[58] Liu, Zafar, and Au, "Rethinking FS-ISAC: An IT Security Information Sharing Model for the Financial Services Sector."

[59] Rosenzweig, "Cybersecurity and Public Goods: The Public/Private 'Partnership'," 12.

2.4 Conclusions

The large number of information sharing initiatives, organizations, and processes suggests that cooperation in computer network defense is a promising idea. Indeed, a total lack of sharing would raise the cost of defense, since no information learned by one defender would be shared with another. An attacker would have significant advantage over each of these defenders because the same tactics, techniques, and procedures would have to be identified and counteracted independently by each defender. For example, "an attacker can send the same spearphishing message to various companies—some of which may catch it, and others not."[60] Such a system would be analogous to a degeneration of the credit reporting system discussed in Chapter 1 in which no creditors shared information with each other; in such a system, a delinquent borrower could obtain additional credit and default on debt from each creditor independently. A system with a complete lack of sharing thus would almost certainly produce suboptimal outcomes; this intuition is confirmed by literature cited earlier in this work.[61]

Often, analyses of cyber sharing success have offered only incremental changes. While this is perhaps a productive mechanism by which to improve the existing system, it implicitly endorses the idea that only fine-tuning is necessary to craft a successful comprehensive cyber threat information sharing system. While the ultimate goal of cyber threat sharing schemes is clear—an enhanced cybersecurity posture for participants, collectively and individually—and existing efforts are made in good faith toward that goal, more fundamental changes may be needed. The institutional structure of existing organizations exposes certain assumptions that may not line up with the reality of today's cyber threat.

[60] White House senior adviser Rand Beers, quoted in Fung, "Washington is making is easier for businesses to swap notes on hackers."

[61] For example: "The level of information security that would be optimal for a firm in the absence of information sharing can be attained by the firm at a lesser cost when computer security information is shared." (Gordon, Loeb, and Lucyshyn, "Sharing information on computer systems security: An economic analysis," 461) and "information sharing by firms can act as a deterrent for hackers" (Gal-Or and Ghose, "The Economic Incentives for Sharing Security Information," 187).

We established in Chapter 1 and this chapter that sharing is reasonably believed to increased computer security. The question now remains: *what* should be shared, and *how*, to achieve the greatest security impact?

Chapter 3

Analytical Approach

> "Computational thinking builds on the power and limits of computing processes, whether they are executed by a human or by a machine. Computational methods and models give us the courage to solve problems and design systems that no one of us would be capable of tackling alone."
>
> Jeannette M. Wing, Professor of Computer Science, Carnegie Mellon University
>
> (March, 2006: Wing, "Computational Thinking," 33)

In previous chapters, we have discussed the current landscape of cyber threat information sharing and identified problems with existing organizational structures and efforts. This chapter presents *computational policy*, a novel abstraction for the analysis of policy prescriptions that, through the application of insights from computer science, allows analysts to design and rigorously analyze solutions to difficult policy problems. Then, it presents the problem of cyber threat information sharing in the view of this new lens and briefly describes one way to categorize of the types of computer security threats facing network defenders today. These theoretical contributions lay the groundwork for the analysis in the following chapter.

3.1 Computational Policy

Computer scientist Peter J. Denning writes of the great advances being realized in many disciplines as biologists, social scientists, artists, and others "[discover] information processes in the deep structures of their field." Underlying these discoveries is a realization that computing "is not—in fact, never was—a science only of the artificial"; rather, computing principles underpin or can be applied to many of the difficult problems faced by an advanced society.[1] Jeannette M. Wing similarly identifies the promise of the wide application of a "computational mindset": not only should computing help solve the existing problems of other fields faster, but it should change the way those social scientists, natural scientists, analysts, and others think about solving problems.[2] To this, we add that those in computing should recognize their responsibility to be ready to collaborate and to apply their knowledge in new ways to pressing problems outside of the comfortable realm of "the artificial".

Although existing analyses of cybersecurity policy often employ theories of organizational dynamics, economics, or political science, it is rare to find analysis based on principles of computer systems design. This is not unreasonable, and the object of this thesis is explicitly *not* to argue that such efforts are inappropriate or misguided. However, each discipline deals with the problem with a distinct analytical focus, one that can be enriched by *also* analyzing problems through an approach informed by the principles used in the design of computer systems. The lack of this perspective so far represents the failure of the computer science research community to fully take responsibility for the complex systems upon which society has become increasingly reliant.

Computer systems are indeed some of the most complex creations of humankind, and the research community has spent decades learning to build systems that tolerate faults, provide guarantees of reliability and security, and implement ideas of concepts like fairness among users.[3] Such applications indicate the promise of insights from economics and other social sciences to computer systems[4]—there is no reason the reverse should not be true.

[1] Denning, "Computing is a Natural Science," 13–14.

[2] Wing, "Computational Thinking," 33–34.

[3] Fairness, for instance, is an explicit measure used by the networking community to evaluate performance of protocols like the Transmission Control Protocol (TCP) that are meant to prevent network overload or the unfair domination of limited resources by a few users. Operating systems similarly implement methods to prevent one program from using up all of a limited resource, like memory or processor time, while another experiences "starvation".

[4] A compelling example of this cross-pollination is VMWare's widely deployed ESX Server product, which uses a resource management framework that is explicitly based in markets and taxation to implement the allocation of a limited resource (physical memory) to separate virtual machines based on their importance

Policy analysis and computer science should not remain disconnected, especially in a world where technology is being infused into nearly every facet of daily life.

We now present an introduction to the key principles necessary to effectively apply algorithmic analysis and identify how these principles apply to real-world policy problems.

3.1.1 Introduction to Algorithmic Analysis

One of the most important foundations of computer science is the theory of *algorithms*, which are sets of defined steps taken to solve a problem.[5]

Often, it is critical to determine the *performance* of an algorithm, which is usually measured in terms of the amount of time, data storage, memory, or network communication needed to produce a result. Some algorithms, especially those that compute solutions to very difficult problems, are also measured by how closely their results match the optimal result. In both cases, these results frequently are *theoretical bounds*, giving provable minimum or maximum values for the metric being measured.[6]

The performance of algorithms is commonly expressed using *asymptotics*. The term is based on the word *asymptote* because asymptotics describe the performance of an algorithm as the size of its input—usually denoted n—grows large, approaching positive infinity.[7] Often, the most useful bound is an *upper bound*—a theoretical maximum amount of time, memory, or other resource that will be required to complete the algorithm.[8] Proving an upper bound $f(n)$ is equivalent to proving that the algorithm requires *no more than* $f(n)$

(Waldspurger, "Memory Resource Management in VMware ESX Server," Sec. 5).

[5] A standard algorithms textbook defines algorithms as "any well-defined computational procedure that takes some value, or set of values, as input and produces some value, or set of values, as output," (Cormen et al., *Introduction to Algorithms*, 5). For an approachable discussion of algorithms and their applicability to concrete problems traditionally solved by non-technical means, see Diakopoulos, *Algorithmic Accountability Reporting: On the Investigation of Black Boxes*, 3–9.

[6] Note that, in contrast with some other disciplines, computer science uses the word "theoretical" as a stronger qualifier than "practical"; a "theoretical bound" is guaranteed to hold in *all* cases for which a proof is presented.

[7] See, *e.g.*, Cormen et al., *Introduction to Algorithms*, 43–53 for a more detailed description.

[8] Unless otherwise specified, we (for brevity, without loss of generality) use time as the measured resource in the rest of the chapter. Time is the most common metric to analyze and is the most relevant to the forthcoming analysis of cyber threat information sharing.

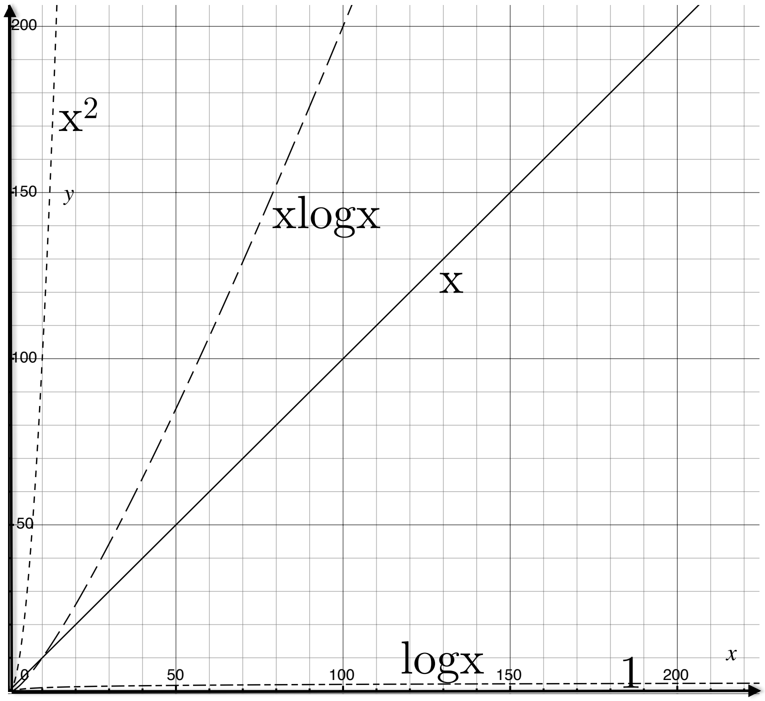

Figure 3.1: Graph of typical asymptotic base functions as input x grows.

time. Conversely, a *lower bound* indicates a theoretical minimum—the algorithm *must* take *at least* $f(n)$ time.[9]

Once proven, asymptotic bounds can be used to classify an algorithm by its complexity and to determine whether the algorithm is an appropriate and feasible solution to a problem. Examples of these complexity classes include the *NP-hard* problems: those believed to be the hardest to solve, which for even reasonably small input might not complete in the amount of time left before the predicted death of the universe.[10] More reasonable algorithms are generally linear (they complete in time less than some multiple of n); quasilinear (they complete in time less than some multiple of $n \log^k n$); or in some cases quadratic (they com-

[9] Upper bounds are generally expressed in "Big-O" notation; saying an algorithm's runtime "$t(n)$ is $O(f(n))$" is equivalent to saying $t(n) \leq cf(n)$ for all $n \geq n_0$ for some positive constants n_0, c. Lower bounds are generally expressed in "Big-Omega" notation; saying an algorithm's runtime "$t(n)$ is $\Omega(f(n))$" is equivalent to saying $t(n) \geq cf(n)$ for all $n \geq n_0$ for some positive constants n_0, c. Intuitively, this means that, after a certain value of n, before which the algorithm's runtime may fluctuate, it will *always* be above or below a fixed constant multiple of the function $f(n)$. Combining the two yields "Big-Theta", which bounds $t(n)$ to be within some constant multiple of $f(n)$ as both an upper and a lower bound. See Cormen et al., *Introduction to Algorithms*, 47–48.

[10] This is not an exaggeration. See Denning, "Computing is a Natural Science," 16.

plete in time less than some multiple of n^2). Algorithms for which the performance does not change based on n are said to be constant-time, or $O(1)$. Examples of these function types are shown in Figure 3.1. Asymptotic bounds do not provide a strict, algorithm-by-algorithm, total order by performance; since they "hide" the effect of *constant factors*, two algorithms in the same complexity class may perform very differently in practice. (For instance, an actual runtime of $2n$ is in the same class as a runtime of $10n$; the constant factors are 2 and 10—*constant* since they do not depend on n and *factors* since they are multiplicative factors of the runtime.) However, if the size of the input is expected to grow, the first criterion in the choice of algorithm for a scalable system should not be the measured performance on a test set of data, but the asymptotic performance potential of the algorithm as the input grows.

These theoretical efforts are not simply academic exercises: they guide system designers in choosing the right algorithms to solve problems. A system designer optimizes the performance scalability, most commonly measured in runtime and memory requirements, with the difficulty of implementation and maintenance of the system.[11] In most cases, using asymptotically attractive algorithms is possible and is critical to performance. As introductory computer science students are instructed, "total system performance depends on choosing efficient algorithms as much as on choosing fast hardware."[12]

3.1.2 Applying Algorithmic Analysis to Policy

It's clear, at least to computer scientists, how such an analysis process could apply to well-defined algorithms for solving problems on computers. But how can algorithmic analysis inform reasoning about policy prescriptions? We consider each of the key concepts from the previous section and suggest how algorithms and policy correspond. The mapping presented here is certainly not the only way such a correspondence could be described; we look forward

[11] For instance, it is probably not worth it to use a difficult-to-implement but asymptotically attractive algorithm if the input is guaranteed to be of bounded small size. It would also be questionable to highly optimize the system for a small improvement in performance at the expense of future maintainability.

[12] Cormen et al., *Introduction to Algorithms*, 13.

to refining it in the future as computational policy analysis is applied to other problems.

Fundamental Operations

In the analysis of algorithms, there are certain necessary assumptions about the cost of fundamental operations. In a computer, these tasks might include reading from or writing to memory, computing an arithmetic operation like a sum, or moving on to a different part of an algorithm. Understanding the cost of these operations is important because they underly many of the more complex tasks a program carries out—analyses must take them into account or lose accuracy. For example, if a computer took more time to read a number from memory depending on its value or on the total amount of memory used, an analysis that assumed that all memory accesses took the same amount of time would be incorrect.

In policy, such fundamental operations might include various basic office tasks like sending emails, reading documents, or answering phone calls.

The more positive side of fundamental operations, alluded to above, is that if the fundamental operations run faster in real-time, so does the algorithm, even though the difference is not reflected in the asymptotic performance bounds.

Asymptotic Bounds

Asymptotic bounds, as discussed, describe how the performance of an algorithm scales with its input. However, they can be applied in a variety of *cases*, the most common of which are *best*, *worst*, and *average*. These correspond to types of inputs for which the algorithm might have a particularly easy or difficult task ahead of it. For instance, an algorithm that is about to sort a list of numbers might find that the list is, in fact, already sorted—this is a *best* case input for that algorithm. An algorithm to recover a lost or unknown password for which only an irreversible "hash" is known might be implemented in a way that some input would cause it to find the answer at the last possible moment—a *worst* case.[13] Average-case analysis

[13] Properly implemented cryptographic hash functions are meant (and generally believed) to be irreversible, so this sort of guessing process is how most password "crackers" are implemented. Many are

gives a bound on the function's expected value (using the definition from statistics, that is, an average of all possible values weighted by their probability of occurrence). Useful analytic results are a subset of the combinations of case type and bound type; for instance, an upper bound on the worst-case performance gives an idea of the absolute worst performance to expect, and a lower bound on a best case gives an idea of the absolute best, but an upper bound on the best case is significantly less meaningful—there might be only one best case, for example, or it might complete in much less time than the upper bound indicates.

In policy, such case types will largely depend on the problem at hand, but follow the same general theme as the above examples. For instance, an office responsible for reviewing petitions for exception from an academic requirement might have a list of situations in which the exception is pre-approved. Upon quick verification of those conditions, the petition might be automatically accepted; this would be a best case. Conversely, a petition that triggered a complicated, extensive review might correspond to a worst case. In the network security problem space, sharing an alert about a routine malicious program or common type of email-borne attack might be a best case, while the response to a truly novel threat might require significantly more time to analyze and coordinate among the members of a sharing organization.

Constant Factors

We discussed above how asymptotic analyses "hide" constant factors that multiply the runtime. While this helps make it clear how algorithms will scale as input grows and makes the analysis easier, it does remove critical information about real-world performance: an algorithm that takes 1,000,000n time will almost certainly perform worse in practice, for all but the largest inputs, than one that takes n^2 time.[14]

In computer systems, these constant factors often come from inefficient code written

deterministic, so there is actually a possibility that the cracker would get through all but the last possible input of a certain size before finding that this last input is the correct answer.

[14] Specifically, only those inputs larger than approximately 1,000,000 would be large enough.

by programmers. With extra attention and tools that measure where the computer spends most of its time while executing a program, a programmer can reduce the effect of constant factors on runtime. The programmer might also inefficiently use memory, for instance by reserving more space than necessary for each item in a list—while this might only show up as an ignored factor of two in an asymptotic analysis, it could in practice cause the program to need more memory than the computer can provide, and a simple fix could change the program from impractical to practical.

In policy analysis, these constant factors might correspond to how efficiently a bureaucratic process is implemented. For instance, if filing a report requires printing a form from a hard-to-find internal website then scanning and emailing it to someone, the constant factor involved would be much greater than that of a simple web-based form that automatically saved results to a database. In a more general sense, factors like a bureaucracy's inertial resistance to change, the difficulty of reassigning personnel from one type of task to another, or the flexibility and responsiveness of a purchasing process might manifest themselves as constant factors.

A critical insight is that constant factors must be weighed against asymptotic performance (as defined by the bounds above). Effort spent reducing the constant factors inherent in an asymptotically expensive process might be a worse choice than efforts to reduce the asymptotic complexity of the process. But, an asymptotically well-designed process might perform significantly better if constant factors are reduced—especially those at the most basic level of the process, since these factors "bubble up" if the operation is repeated asymptotically many times.[15]

This is one of the clearest examples so far of the power of the computational policy technique—this is a case when traditional analysis would generally indicate that analysts ought to ascertain in detail the reason for a process's performance. Applying the algorithmic analysis approach, we know that we can disregard many of the "implementation details"

[15] For example, if a process requires $O(n^2)$ data-entry tasks to be executed, streamlining that basic process to take half the time would reduce the original runtime t to $(.5)^2t = .25t$, a reduction of 75%.

of a process and concentrate on the bigger picture until we want to analyze the "constant factors" of a process we have already found to be asymptotically acceptable.

Resource Complexity

Most algorithmic analysis concentrates on runtime, but other resources are often measured too. The most common of these other measurements is the memory required to execute the algorithm—its *space complexity.*[16] Memory usage is measured using the same types of asymptotic bounds used for other quantities. Often, space complexity and time complexity are opposed: one can be optimized (sometimes significantly) but at the expense of the other.

Space complexity is fairly easy to understand in computer systems. Fast-running algorithms often rely on saving some intermediate results in memory to avoid having to do the same work multiple times. If space in memory is limited, the computer might have to spend extra time duplicating work that could otherwise have been saved.

In policy, such resource complexity constraints might correspond to more traditional resource limits, like personnel or office space. This is a slightly less direct correspondence than the previous three, but is still reasonable, since memory requirements often essentially measure the amount of working "scratch" space required to complete an operation or the number of tasks that can be "tracked" at once. Also, constraints on both computer memory and its policy analogues can, for the most part, be ameliorated by spending more money, for instance to simply buy more memory chips or hire more personnel.

3.1.3 Defining Computational Policy

With the material of previous sections in mind, we now define *computational policy* as an analysis process that includes:

- Specifying a policy problem as a computational goal;

[16] For further theoretical reference on this topic, see Sipser, "Space Complexity."

- Specifying proposed solutions as precisely as possible, including identification of:
  - Actors,
  - Interactions between these actors, and
  - The steps required to reach the goal, including a specification of which actor performs the task and its approximate complexity; and
- Reasoning about the performance of the solutions in anticipated use cases and as the problem size grows asymptotically large.

This methodical approach to policy emphasizes scalability and performance, which often take the backseat to feasibility and ease of legal or regulatory implementation in current policy debates.[17] It is a change in perspective and in emphasis, not a wholesale reinvention of the policy analysis process, but it promises to formalize and systematize many of the important considerations implicit in current policy decisions, and include some helpful concepts that are missing.

Ideas for the application of computing to policy have begun to appear in recent years, most often in the context of improving transparency or the ease of governance. Technology publisher Tim O'Reilly advocates for "algorithmic regulation" in a chapter for Code for America's electronic book on reimagining local government, suggesting that regulations be specified as series of measurements, defined steps to take in response to new information, and periodic performance analyses, rather than as static rules.[18] This welcome—if questionably viable—step might lead to more agile policies, but leaves out the performance analysis focus of the computational policy lens and would require a more radical reinvention of existing legal and regulatory mechanisms.

By applying insights from computing to policy designs using the computational policy technique, we can leverage greater knowledge to obtain better outcomes and more scalable

[17] See, for instance, discussion in Ch. 2 on the current industry-by-industry cyber threat information sharing mechanisms and the revealing language of "Presidential Policy Directive 21: Critical Infrastructure Security and Resilience," discussed on p. 52 of this thesis.

[18] O'Reilly, "Open Data and Algorithmic Regulation."

policy solutions to the pressing, growing, “large-n” problems of modern life, including but not limited to cybersecurity.

3.1.4 Cyber Threat Information Sharing in the Computational Policy Frame

This section begins the task of formalizing the cyber threat information sharing problem as a computational policy problem. First, we define a goal of cyber threat information sharing—something which is left unstated in much of the existing analysis, but which is critical if we are to reason about and compare schemes for sharing. Next, we define some of the desirable characteristics of a system that would contribute to this goal or improve its feasibility or scalability.

Defining a Goal

O’Reilly highlights the critical need for “a deep understanding of the desired outcome” and specified performance measurements in order to succeed with algorithmic regulation.[19] While it is tempting to count the very act of sharing—each email sent, phone call made, meeting held, or report written—as success, useful and rigorous analysis requires us to move past this oversimplification: we must at least attempt to evaluate the effect of sharing on actual security posture. With this in mind, and taking O’Reilly’s counsel, we define our goal as *raising the cost to the adversary.* This may mean, for instance, that the adversary has to “burn” infrastructure more often: network defenders maintain “black-lists” of known-malicious addresses, machines, and tools that prevent protected machines from interacting with them, so, once discovered, the infrastructure is significantly less useful.[20] Or, the

[19] O’Reilly, “Open Data and Algorithmic Regulation.”

[20] Internet Protocol (IP) addresses, software tools, and domain names (*e.g.,* `attacker.com`) that an attacker uses to launch attacks or control compromised computers are often blocked. The attacker must then establish new infrastructure, which takes time and money, before attacks can resume. APT1 and APT12 actors changed nearly all the infrastructure identified in Mandiant’s public APT1 report after Mandiant released the report (Mandiant, “APT1: Exposing One of China’s Cyber Espionage Units”) and identified the APT12 attackers who compromised the *New York Times* (Perlroth, “Hackers in China Attacked the

adversary might be required to develop complex new tools to avoid detection, might be prevented from reconnoitering or probing targets with impunity, or might be able to make fewer mistakes in a target network before being discovered. In each case, the adversary has been slowed, forced to rethink targeting decisions and tools, or required to operate with less information; in short, the adversary's life has been made more difficult.

This concept has been made popular by a number of leading individuals and companies in the computer security industry.[21] However, it is not the creation of a shrewd marketer—it is also found in many disciplines of computer science. Most prominently, the standard security model in cryptography—the computational security model—is defined explicitly in terms of the cost (usually measured in time) versus the probability of successful compromise for a hypothetical adversary; designers of cryptographic systems seek to bound the probability of successful compromise for any adversary who spends a fixed amount of resources, or, equivalently, seek to raise the minimum effort required to successfully compromise the system with a given non-negligible probability.[22] In a similar way, reputation systems—systems used to establish some quality about an entity, like trustworthiness—are measured in terms of "adversary cost for manipulating these rankings in a variety of scenarios" and the feasibility of the adversary's response to countermeasures.[23] The proposed metric, cost to the adversary, is clearly not without foundation.

Maximizing adversary cost, in many cases, will also minimize *time to convergence:* the time required for the distributed system to "agree" on a fact, for instance, that a program is malicious. This is a more commonly mentioned metric—"how fast can we stop a new threat?"—but it is subsumed by the fuller notion of raised cost to the adversary.

Times for Last 4 Months"), according to Mandiant, "One Year After the APT1 Report: What Difference Does a Year Make?"

[21] Proponents of the concept include Dmitri Alperovitch, formerly at McAfee Threat Research and now at CrowdStrike (from whom the exact phrase first came to this author's attention) *e.g.*, in Lemos, "Companies See Business In 'Doxing' The Adversary," along with other quoted executives; and government officials, *e.g.*, in Freedberg, "They're Here: Cyber Experts Warn Senate That Adversary Is Already Inside U.S. Networks."

[22] Katz and Lindell, *Introduction to Modern Cryptography*, 48–49.

[23] Marti and Garcia-Molina, "Taxonomy of trust: Categorizing P2P reputation systems"; Vu, Papaioannou, and Aberer, "Impact of Trust Management and Information Sharing to Adversarial Cost in Ranking Systems."

Defining Desirable Qualities

A variety of qualities would contribute to a raised cost to the adversary; we list some here. These are important aspects of the system that make it more realistic and useful, and may make the practical difference between feasibility and infeasibility.

- *Reliability:* The sharing system, if it is to be deployed and used widely as an integral part of network defenses, must not crash or require continuous maintenance.
- *Distributed Survivability:* The system must perform in the face of attack, and should not become a single point of failure in the security system.[24] The "distributed" portion of this quality indicates that the sharing system should not cause the fates of multiple parties to be tied together where they previously weren't. That is, if two entities could previously operate independently, the failure of one should not imperil the other because of the sharing system.[25]
- *Performance under load:* The system's performance should degrade gracefully under load, since the system may be needed the most precisely when it is under the most strain—when its participants are experiencing negative security events.
- *Local autonomy:* For the system to be accepted by disparate partners, not all of whom may trust each other or share a common characterization of threats or level of risk tolerance, it must provide the ability to decide rules and risk tolerance for one's own network. (In other words, it must not too radically upset the current paradigm of local control.) This also is desirable as a matter of practicality, since determining risk management rules that fit all possible networks may not be possible—such a set of rules would leave an unnecessarily large attack surface for some networks and unacceptably

[24] This is based on the definition of "survivability" in Ellison et al., *Survivable Network Systems: An Emerging Discipline*, 2.

[25] This notion is commonly referred to as "fate-sharing". We would seek to limit fate-sharing where it is not necessary.

limit the services that could be offered by others.[26]

- *Resistance to false input:* The sharing system would ideally limit the ability of malicious, ill-informed, or incompetent users to cause other users to take undesirable action.

3.2 A Taxonomy for Cyber Threats

For the purposes of the ensuing analysis, we now define certain classes of attacks. This taxonomy divides attacks only by how specifically they are targeted.[27] This analysis intentionally conflates attackers with their attacks—that is, with the specific steps and tools used to survey, exploit, and maintain access to a target system. While it is true that some attackers use multiple sets of tools and some tools are used by many attackers, this approximation does not imperil the analysis. In practice, distinct tools are usually dealt with separately until the attackers leave enough clues that individual tools are connected to a single group. Even tools that are largely the same often bear identifying marks or are used in idiosyncratic ways that make them distinguishable. The analysis that follows will consider three types of attacks:

- **Generic attacks** are "sprayed" indiscriminately across the Internet in an attempt to infect as many machines as possible, with little regard to the specific identity or characteristics of the target machine. Such attacks may be launched by attackers who are attempting to build up a "botnet" of compromised machines that can be rented to others to send unsolicited email, attempt to overwhelm targeted websites, or infect other machines. They may also be attempts to collect private information, like bank account credentials or payment card numbers, that can be sold or used to steal money.

[26] "Attack surface" is a term that conceptually indicates how many ways an adversary could attack a system. Greater attack surface indicates more vulnerability.

[27] As a scientific experiment manipulates a single variable at a time, this scale attempts to consider only variation in target specificity. It is likely that other characterizations of the threat environment would function better in other analytical situations.

- **Semi-targeted attacks** are somewhat cohesive—an apparent attempt to gather information or disrupt the operations of a certain type of target, such as those from a single industry. Such attacks are more likely to be focused on obtaining a certain type of intelligence or information than on blindly adding to a botnet.

- **Highly targeted attacks** are designed to penetrate a single target, or a small number of highly specific targets, usually to gather information not available elsewhere or to cause unique impact. Such attacks are decidedly *not* crimes of opportunity or convenience, and often use tools generated specifically for the target against which they are used. For example, the targeted attack by Chinese actor "Advanced Persistent Threat 12" on the *New York Times* in 2012–2013 employed 45 custom pieces of malicious software, only one of which was recognized by security vendor Symantec.[28]

[28] This is not an error, despite the similarity of this actor's identifier to the more well-known APT1. Mandiant numbers actors using the prefix "APT"; this is the twelfth. (Perlroth, "Hackers in China Attacked the Times for Last 4 Months"). This is an unusually high number of malicious tools for a compromise of only 53 computers. (Mimoso, "Inside the Targeted Attack on the New York Times").

Chapter 4

Applying Computational Policy

“We need information sharing, in time and at network speed.”

Gen. Keith Alexander, then National Security Agency Director

(July 9, 2012, in: Jackson, “‘Destructive’ cyber attacks ahead, NSA’s Alexander warns”)

The main task of this chapter is to analyze three archetypal information sharing models using the *computational policy* analytical approach developed in the previous chapter, providing new analysis of their performance and scalability.

Before proceeding to the case analyses, we first make some preliminary analyses that will inform the cases that follow. The preliminary subsections are not overly focused on algorithmics, which—far from a criticism of the computational policy approach—actually demonstrates how existing analysis techniques can complement the new one. Following the cases, some general points, a summary, and conclusions are presented.

4.1 Preliminary Analyses

4.1.1 The Effect of Industry Segregation

Most, if not all, of the sharing efforts currently underway are organized around single industries.[1] But, it is not entirely clear that the types of threats facing companies are best addressed with cooperation that respects industry boundaries. This section argues that

[1] These are mostly single “critical infrastructure” industries, as designated by PDD-63, HSPD-7, PPD-21, or the USA PATRIOT Act of 2001. A list of ISACs is found in n. 3 on p. 19.

intra-industry cooperation can be useful to counter some classes of threats in the current threat ecosystem, but that this class of threats is not large.

The fact that most collaboration is accomplished separately per industry does not seem to have been a critical feature of the original design put forward in PDD-63. Notably, PDD-63 refers to a *single* ISAC for the private sector, even while designating multiple sector-specific government lead agencies and personnel.[2] PPD-21, President Obama's 2013 successor to President Clinton's PDD-63 from 1998, explicitly recognizes that the motivation for using sector-specific agencies is partly for economy of effort—to use "existing statutory or regulatory authorities" and "[leverage] existing sector familiarity and relationships"—rather than wholly because such a design will produce improved cybersecurity outcomes.[3] So, is today's intra-industry threat sharing an effective model?

It is unlikely that single-industry threat sharing cooperation could provide much additional value against generic threats. Such attacks are so widely spread and, frankly, generic, that efforts to counteract them could be undertaken more effectively by cybersecurity companies and distributed over existing channels to all of their customers. The electricity, public transportation, and water ISACs, for instance, should not all have to analyze the threat and issue sector-specific reports on it when the same information is applicable to all companies and even individuals, unless they can provide significant sector-specific value-added information. More strategically, organizations that spend time addressing generic threats incur greater hidden risks: that their reports will come to be viewed as mostly useless, and that the opportunity cost of time spent duplicating outside effort will grow large as highly subtle—and perhaps *actually* sector-specific—threats are ignored. An illustrative example is the Multi-State ISAC (MS-ISAC)'s list of public "Cyber Security Advisories", which appears to consist solely of rewritten versions of vendors' vulnerability announcements; inspection of the ten most recent advisories at the time of writing showed that each one was assigned the

[2] "Presidential Decision Directive 63: Critical Infrastructure Protection," Section IV and Annex A.

[3] "Presidential Policy Directive 21: Critical Infrastructure Security and Resilience," Section "Sector-Specific Agencies".

same risk level—“High”—for all types of governments and businesses, and that the content of each advisory was either directly derivable from the vendor’s announcement or consisted of security advice that was generic enough to apply to almost any vulnerability.[4] The Financial Services ISAC (FS-ISAC) provides a similar “Public Sample” of its “Sector Alerts”; at the time of writing, this sample was five months old and included approximately half financial services-related content: two industry-related member submissions, one industry-specific alert, one general critical infrastructure alert from DHS, and three generic alerts.[5] Since MS-ISAC and FS-ISAC, like many other ISACs, tag information based on the level to which it may be shared, it is likely that they provide more targeted information to members. However, the fact that any nontrivial portion of analysis resources is spent on dealing with low-value, generic threats is troubling.

Sector-specific sharing organizations are best suited to semi-targeted attacks, and specifically to the subset of such attacks that attacks only a single industry. A large number of actors target multiple targets within the same industry, giving ISACs and other industry-specific organizations the opportunity to band together to fight a common infection or adversary.[6] For instance, FS-ISAC reports that it issued a notice about “Hesperbot—An Advanced New Banking Trojan in the Wild” and of a type of fraudulent email that was related to banking transfers.[7] And, FS-ISAC reportedly “validated information sharing” as an important arrow in the network defender’s quiver by facilitating collaboration during

[4] Multi-State Information Sharing & Analysis Center, “MS-ISAC Cyber Security Advisories,” Advisories 2014-038 to 2014-047. Typical advice included instructing users not to click on suspicious attachments, patching systems, running software as a non-privileged user (*i.e.,* not using an administrator’s account), and deploying standard security measures like Microsoft’s Enhanced Mitigation Experience Toolkit (EMET).

[5] Retrieved from the FS-ISAC homepage, `fsisac.com`, on 20 May 2014. Alerts were dated 12/06/2013 to 12/17/2013.

[6] The proportion of white papers and warnings issued by security vendors in response to single-industry attacks can be viewed at, for example, `http://www.symantec.com/security_response/publications/whitepapers.jsp`.

[7] Report title listed on the FS-ISAC homepage, `fsisac.com`, on 20 May 2014. However, it’s unclear whether FS-ISAC provided any additional knowledge beyond the public paper that antivirus firm ESET published on the same topic (Cherepanov and Lipovsky, “Hersperbot—A New, Advanced Banking Trojan in the Wild”). Also, the malware was primarily affecting users in the Czech Republic, Turkey, Portugal, and the United Kingdom, far from the U.S. financial system. The unrelated other alert referred to SWIFT transfers, which are used in the United States.

sustained attacks on the U.S. financial sector by actors associated with Iran in 2013.[8] However, not all actors respect industry boundaries, which might lead industry-based sharing organizations to build up disconnected information "silos". Prominent examples of attackers crossing industry boundaries include some of the most well-known recent groups made public by security firms, including "Shady RAT", which affected 32 "organizational categories" in 14 countries;[9] "APT1", which affected 20 industries;[10] and the "Elderwood Gang", which affected at least 6 types of industries in 10 countries.[11] According to Symantec, the average number of industries affected in email spear-phishing attacks (where an "attack" is linked by attributes like email contents, malicious attachments, and origin) is two, and attackers have reduced the overall volume of attacks in recent years to attempt to evade detection by the increasingly effective and economically-motivated generic attack mitigation industry.[12] These attacks won't be optimally addressed by a sharing system that is fractured along industry lines.

Highly targeted attacks are often refined specifically for a single target or a small group of targets. It is, first, less likely that such an attack will be discovered by commonly used security products, which are best at blocking known, pre-identified threats rather than emerging ones—and which may be the very source of the security alerts we are hoping that companies will share. Then, because highly targeted attack tools and techniques are generally not reused,[13] even if indicators pertaining to the attack are shared within an industry, that information is unlikely to lead to the discovery of other attacks. Sharing organizations therefore have little hope of discovering and mitigating such highly targeted attacks.[14]

[8] Donovan, "FS-ISAC threat information sharing helped thwart DDoS attacks against US banks."

[9] Alperovitch, "Revealed: Operation Shady RAT," 4–5.

[10] Mandiant, "APT1: Exposing One of China's Cyber Espionage Units," 3.

[11] O'Gorman and McDonald, "The Elderwood Project," 5–6.

[12] Symantec, "Internet Security Threat Report Appendix 2014," 62–64, 66. Focused campaigns rose by 68% year-over-year, while "mass-scale" campaigns dropped by 32%.

[13] For example, as noted in the previous chapter on p. 50, attackers used 45 custom pieces of malicious software, only one of which was recognized by security vendor Symantec, to infect 53 machines in an attack on the *New York Times*. (Perlroth, "Hackers in China Attacked the Times for Last 4 Months").

[14] Stuxnet binaries were observed by many virus companies years before it was identified as a targeted attack on Iran. No one noticed! Similar attacks thus may not be discovered until the attacker makes a mistake or loses control of the infection. See Kushner, "The Real Story of Stuxnet." A study considering

It is not clear that the system of ISACs provides appropriate conduits for information sharing across sectors. Government partners might be able to bridge such gaps, for instance through the National Infrastructure Protection Center or other DHS initiatives, but ISACs and private working groups have been disappointed by the quantity and quality of information shared with them by their government partners.[15] The National Council of ISACs does facilitate regular cooperation between ISACs, but its monthly meetings and limited coordination function make its ability to provide near-real-time cross-sector collaboration questionable, though it did recently set up an online portal meant to facilitate cross-sector sharing.[16] Apparently responding to insufficient sharing among ISACs, DHS recently decided to layer another level of sharing on top of the ISACs, establishing the "Critical Infrastructure Information Sharing and Collaboration Program", chartered "to improve sharing among ISACs, information and communications technology service providers, and their respective critical infrastructure owners, operators, and customers".[17] Whether yet another level of sharing is sufficient to solve the problems present in the current design is unclear.

4.1.2 The Makeup of Sharing Organizations

Experienced executives and leaders know that a meeting is a waste of time unless the right people are at the table; in a similar way, even the best-designed sharing organization will be of little use if its participants are not the ones who ought to participate. The strategic success of sharing organizations relies on reaching most, if not all, of the potential targets in their space, so this is a critical issue. The structure and analysis of ISACs and other programs like ECS take as a premise that they can eventually reach a large enough percentage of each industry that the whole industry will be protected—otherwise, the least-protected player in

other exploits is Bilge and Dumitraş, "Before We Knew It: An Empirical Study of Zero-day Attacks in the Real World."

[15] General Accounting Office, "Improving Information Sharing with Infrastructure Sectors," 9; The Rendon Group, *Conficker Working Group: Lessons Learned*, 35, 39, 40.

[16] Donovan, "FS-ISAC threat information sharing helped thwart DDoS attacks against US banks."

[17] U.S. Government Accountability Office, "Cybersecurity: National Strategy, Roles, and Responsibilities Need to Be Better Defined and More Effectively Implemented," 55.

an industry could be used to compromise the rest through existing trust relationships or network interconnections. While individual network operators may not be concerned with this (though some would, due to ongoing business relationships with or dependencies on other companies), those responsible for overall critical infrastructure resilience strategy—that is, those who run programs like ECS and sponsor organizations like the ISACs—definitely must address this issue. The government might also be concerned that the compromise of one critical infrastructure operator could cause significant impact to the nation, even if all others are safe: just one bank or electrical utility could destabilize all interconnected operators. This type of "weakest link" vulnerability was illustrated by the 2014 breach of payment processing systems in Target stores, which was accomplished through the use of credentials stolen from an air-conditioning contractor.[18]

Sharing organizations also overrepresent large, established companies and underrepresent small businesses, non-profits, and other institutions. This imbalance may be due to the fact that large companies typically have a dedicated security staff, while information technology priorities in smaller companies and institutions are generally more focused on providing capabilities, not securing them. A white paper from the ISAC Council, written in 2004 with the consensus of all Council members at that time, noted that many smaller critical infrastructure operators were left out of current ISACs and were unable to participate due to structural or funding constraints.[19] The IT-ISAC, for example, appears rather oddly to contain only large information technology production companies, agricultural companies, IT resellers, and security consultants.[20] The observation that small operators are in fact targets is confirmed by Symantec, whose annual threat report warns that a rapidly increasing portion of attacks are aimed squarely at small businesses; up to 31% of all attacks in 2012

[18] Krebs, "Target Hackers Broke in Via HVAC Company."

[19] ISAC Council, "Reach of the Major ISACs," 8. The white paper makes somewhat specious claims regarding ISACs' ability to reach a large portion of critical infrastructure operators; a number of these claims are based on liaisons with trade associations and other groups that do not truly indicate that any specific company is participating in a sharing effort. This imprecision is an overestimate of ISACs' reach, so it only strengthens the assertion that many operators may be left out.

[20] Based on the list of members available at `http://www.it-isac.org/#!members/c1tsl`.

and 30% in 2013 were directed at them.[21] So, some of the companies most vulnerable to attack may also be those least likely to participate in, and receive the benefits of, sharing organizations.

This would be a bad enough outcome if small critical infrastructure operators were simply unaffected by sharing. However, effective sharing among the larger, better-prepared operators could cause the actors launching semi-targeted attacks at an industry to concentrate on the low-hanging fruit: the least prepared companies.[22] The designers of sharing organizations must carefully consider whether they are simply removing the target from their participants' backs and placing it on less prepared targets; if the goal is to improve the overall security posture of critical industries, such solutions are insufficient.

We now proceed to case analysis. The first information sharing design is the Information Sharing and Analysis Center (ISAC), the long-standing model encouraged for critical infrastructure sectors by Presidential Decision Directive 63 (PDD-63). The critical feature we analyze is the 24/7 "watch floor", so this analysis could apply to analogous models as well.

The next is an generic version of social network-based sharing, patterned after the industry-based responses to certain attacks like Aurora and the Facebook "watering-hole".[23] Similarly, it would apply to other sharing organizations that principally promote social trust to later enable sharing.

The final design is a distributed system in which a greater volume of security-relevant signals are shared more widely. This sort of design is becoming more common in actual practice but remains rare in the analysis of cyber threat information sharing. The exact specifics of the design do not materially affect the analysis, so this model stands in for many of the new and promising efforts in this space that we briefly note later.

[21] Symantec, "Internet Security Threat Report 2013," 4, 16; Symantec, "Internet Security Threat Report 2014," 30.

[22] This is an example of the "diversion effect" discussed in Rosenzweig, "Cybersecurity and Public Goods: The Public/Private 'Partnership'," 9.

[23] This term is explained in subsection 2.2.1.

Each case is presented in narrative form along with a more formal algorithmic description written in *pseudocode*. Pseudocode is simply a notation allowing easier specification of algorithms; it is intended to be read like English.[24] Each actor's behavior is encapsulated in a function that is named for that actor. The design is then analyzed for performance using the qualities from the previous chapter. Particularly relevant qualities are highlighted for each case.

4.2 Watch Floor Sharing

The "watch floor" model, employed by Information Sharing and Analysis Centers (ISACs) and government agencies, involves a staff of individuals whose job is to coordinate any reaction to observed activity in their area of responsibility.[25] ISACs, in particular, are generally expected to operate watch floors 24 hours a day, 7 days a week, to respond to physical and cyber threat incidents within their industries.[26] (Other activities carried out by ISACs, like awareness efforts and conferences, do not have immediate effects during time-critical attack events and have less direct impacts on cost to the adversary, so we exclude them from this analysis.)

Following the computational policy approach, we identify a procedure for each actor to follow. In this case, the actors are the individual reporting institutions and the watch floor. A simplified version of the watch-floor sharing process is presented in Algorithm 1. In brief, the process for sharing an incident begins with an ISAC member. The member's network security staff would first identify an attempted or actual intrusion, denial of service attack, or other security threat. If they decided to report the incident to the ISAC—available

[24] One construct that may be unclear is the *loop*. To loop is to simply repeat a task indefinitely.

[25] These may represent a form of institutional isomorphism (DiMaggio and Powell, "The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields") from the many 24/7 operations centers in the Intelligence Community and in law enforcement. These include, as examples, the National Security Operations Center (NSOC) at NSA and the National Counterterrorism Center's watch floor (Office of the Director of National Intelligence, "National Intelligence: A Consumer's Guide," 24, 43).

[26] U.S. General Accounting Office, "Critical Infrastructure Protection: Improving Information Sharing with Infrastructure Sectors."

survey data suggests this does not happen reliably[27]—the security staff would spend time gathering data, ensuring its correctness, and completing other tasks to prepare a report for submission. The watch floor staff would then evaluate, anonymize, and disseminate the report if it were considered likely to affect other ISAC members.

Algorithm 1 Idealized Algorithm for Watch Floor Sharing

```
function Member
    loop
        look for threat
        if threat is found then
            investigate the threat further
            determine indicators of compromise
            if threat is worth sharing then
                write a report
                upload report and supporting materials to ISAC
            end if
        end if
    end loop
end function
function Watch Floor
    while there are pending reports from members do
        judge whether report is worth sharing
        if it is worth sharing then
            ensure that anonymity is protected
            disseminate report to members
        end if
    end while
end function
```

We now present analysis of the complexity and performance of this algorithm.

4.2.1 Runtime

Each step of this process is expensive in terms of time. The staff of each member must determine whether a specific action is malicious—a task that grows more difficult with an

[27] A 2013 survey found that only 22% of companies even participate in an ISAC (PricewaterhouseCoopers LLP, *Key findings from the 2013 US State of Cybercrime Survey*). In addition, off-the-record interviews and other sources (*e.g.,* those discussed in Ch. 2) indicate that even ISAC members may be reluctant to report particularly sensitive incidents.

increasing amount of events on the network. With network traffic and the number of legitimate services growing considerably each year, network security teams are playing a catch-up game. Further, when adversaries' tactics are only useful as long as they remain undiscovered, delays can mean further intrusions; a series of slow, human-enabled analysis steps at each level of the sharing path may thus reduce effectiveness.

First, we note that computation is distributed: complex local decisions—judgments about whether a network has been compromised—are made by each member. This is not intrinsically a bad idea, but applying approximations earlier in a complex computation can cause accuracy to diminish; for example, if a student rounds numbers during each step of a scientific calculation, the result they obtain will likely be nontrivially wrong. It also introduces a significant delay—at least a constant factor, potentially greater as a function of the amount of network traffic the local staff member is responsible for. In this case, distribution could lead to poor performance against the very type of threat that sharing organizations seem best-tuned to confront: those that attack multiple targets of the same type. Consider that an intruder in one's network is neither guaranteed to be discovered nor guaranteed to be shared outside the member organization; denote the probability of discovery $p_{discover}$ and the probability of sharing a discovered event p_{share}. The optimal case would be discovering and sharing all relevant security events that occur in multiple members' networks; with fully shared data, this would happen with some probability approximately $p_{discover}$ since, once found in one network, identifying the same event in a perfectly shared dataset would be trivial. However, in the federated structure of the ISAC model, the probability of discovering a shared event depends on multiple organizations not only both detecting but also *both reporting* the issue. This occurs with probability approximately $1 - (1 - p)^k - kp(1 - p)^{k-1}$ where $p = p_{discover}p_{share}$ and k members are affected.[28] We can slightly relax this bound by noting that the watch floor could theoretically share every report it receives, meaning that the probability of sharing this knowledge is closer to $p_{discover}p_{share}$, but the ISAC might not

[28] This is the binomial distribution for any number of discoveries > 2.

share every single event that occurs for fear of overwhelming their members with so much noise that the most important signals are lost. Distribution may thus significantly reduce the likelihood of threat discovery.

The evaluation step may also be slow—with any appreciable number of reports coming into a watch floor, determining whether an attack is also occurring on another member's network could, in the worst case, take up to the time required to compare all pairs of incoming reports. Process optimizations and searchable indicators could reduce this time from the worst-case $O(n^2)$.[29]

Effectiveness Against Different Types of Attacks

Analysis of how effective this type of industry-based watch-floor strategy might be against the categories of attacks presented in section 3.2 was presented in subsection 4.1.1. With broader sharing (across industry lines) the performance argument above would dominate and the issues identified with industry segregation would be insignificant.

4.2.2 Performance Under Load

The presence of a central sharing authority means that there will always be a limited-capacity bottleneck through which all new reports, updates, questions or clarifications must flow; in crisis—when such information may be most valuable—this situation will become even worse as the capacity of those at the center of the sharing network is saturated. This is a classic weakness of centrally-managed services.[30] Direct communications, avoiding this bottleneck, would come at the expense of anonymity, or, even with an anonymous message system, could potentially overwhelm the reporter with replies from the rest of the ISAC members.

[29] This upper bound comes from the number of possible pairs of n reports: $n(n-1)/2$.

[30] Denning, "Computing is a Natural Science," Table, p. 16.

4.2.3 Other Qualities

The other qualities are easier to explain and less critical to performance. The *reliability* of the watch-floor model is mostly assured by non-technical means; the actual watch floor facility's reliability is dependent on the same infrastructure as other critical facilities. In the same way that a data center or critical military installation can be kept running, the watch floor can—reliability is an external quality, not intrinsic to the algorithmic design.

Distributed survivability and *local autonomy* can both be achieved by simply ignoring the watch floor. No two participants in a watch floor sharing model are required to be in constant contact, and neither is a participant required to be constantly connected to the watch floor; operations can continue without any other party. In a similar way, the only influence the watch floor has over its participating members is through reports and collaboration. There is no way for a watch floor organization to force a change to local security policies, except perhaps under penalty of ejection from the group.

This model also achieves *resistance to false input* through mostly non-technical means. While any respectable watch floor would use adequate cryptographic and security measures to authenticate users and data sent by them, the correctness of the actual contents of reports depends largely on the trustworthiness of the sharing party. (This point suggests that smaller groups of participants might be better able to resist misinformation; large programs like InfraGard could more easily suffer from compromise in this regard.)[31] This quality also relies on non-technical vetting of participants before they are admitted to the group. Presumably this is a standard process for most watch floor organizations, since failing to perform due diligence before sharing potentially sensitive information would be a glaring failure of judgment.

[31] To give InfraGard credit where it is due, participants are informed on the InfraGard membership application form that a background check will be conducted on them.

4.2.4 Analytical Limitations

There are some limitations to this analysis. It appears that some ISACs operate only "on-call" services after hours, raising questions about their ability to truly respond in real time to incident reports.[32] These working hours would interact poorly with those of attackers, who are known to increase attacks when they expect people to be overwhelmed with email or computer tasks, or to be less diligent about security (for instance, during holidays or on Monday mornings).[33] In addition, for example, the workday in China Standard Time corresponds to times when few U.S. workers are in the office, and some of the most active attacker groups are known to operate during these hours.[34] Nevertheless, for simplicity, the above analysis takes ISACs at their word, assuming they can actually perform their duties around the clock.

4.3 Social-Network Sharing

The private, social trust-based sharing structures discussed in Chapter 2 offer a very different communication paradigm from the ISACs. Social groups, including the Bay Area CSO Council, provide social contact and build relationships that can be used later during a company's response to a discovered incident. Technical groups formed from previous work experiences, like those that spawned the Conficker Working Group, are similarly based on trust, but more professionally and transitively—members accepted other trusted members' assertions about third parties, and most trust was from previously working together or from

[32] For instance, the Research and Education Networking ISAC indicates that its "normal hours of operation are 0800 to 1700 U.S. Eastern" and its goal is to return a "human-provided acknowledgement within four normal business hours" for "normal priority reports" (Research and Education Networking Information Sharing and Analysis Center, "REN-ISAC CSIRT RFC2350 Description," 2.11, 4.1). Similar descriptions in the standard RFC2350 format could not be located for other U.S. ISACs.

[33] Symantec, "Internet Security Threat Report Appendix 2014," 66.

[34] For instance, APT1 operators were observed to observe the Chinese New Year holiday and begin their work at approximately 8a.m. China Standard Time each day, which is 5p.m. Pacific and 8p.m. Eastern Daylight Time. (Mandiant, "One Year After the APT1 Report: What Difference Does a Year Make?," 18; FireEye Labs, "The PLA and the 8:00am-5:00pm Work Day: FireEye Confirms DOJ's Findings on APT1 Intrustion Activity").

trusting work done by a given institution.[35]

Ad hoc trust-based sharing has in many cases filled the gap where organizations failed to facilitate sharing. Their communications tend to be very candid— communication occurs via phone call or other personal contact[36] or via a private broadcast email list, in which every member can see all communications.[37] Such sharing relies less on formal structures and processes, but can still be modeled algorithmically. Our simple model is presented in Algorithm 2.

Algorithm 2 Idealized Algorithm for Social-Network Sharing

```
function Participant
    loop
        look for threat
        if threat is found then
            if colleagues' contact information is missing then
                gather contact information
            end if
            while more information is being discovered do
                call or email colleagues from similar organizations
                investigate the threat further
            end while
        end if
    end loop
end function
```

We now present analysis of the complexity and performance of this simpler algorithm.

4.3.1 Runtime and Information Spread

We model this problem as an information spread problem through a friendship graph to analyze its performance. A friendship graph is a graph with people as nodes and friendships or trust relationships as edges, possibly weighted to indicate relationship strength. Then,

[35] See the Chapter 2 for more detailed discussion of these two examples.

[36] For example, in the Facebook "watering-hole" attack (Facebook Security Team, "Protecting People On Facebook").

[37] For example, in Conficker (The Rendon Group, *Conficker Working Group: Lessons Learned*). This model broke down somewhat as membership grew past the level that personal or limited-social-distance trust could accommodate.

information spread is easily modeled with probabilistic contagion models or information spread models, which suggest that the rate of information diffusion will begin high and monotonically decrease, eventually converging to zero.[38] The number of other entities which will be informed through such mechanisms is thus limited. In the most favorable cases, the information will reach an entire connected component of the graph—all friends of friends of friends, and so on—but if ties are weak the diffusion may stop earlier. In addition, prior expectations between these participants might preclude any further transitive sharing, making the first step the only one.

However, despite the limits of total information diffusion, such a process is likely to provide much faster information diffusion in the initial stages of an incident response. Again, Conficker and the Facebook watering hole responses are instructive, since they indicate that sharing unfinished information can be useful for security. This sharing model effectively short-circuits longer paths in the information spread graph (from the originator to the sharing center to the other entities), suggesting that the two mechanisms could be used in parallel to improve (that is, raise) overall cost to the adversary. Further strengthening this point is the high likelihood that personal trust relationships would form within existing technical or business communities in the real world, suggesting that saturating a community in the friendship graph is equivalent to informing those who work on similar problems and are thus likely to be similarly targeted. In addition, the reduced time spent evaluating information and the reduced barriers to direct information sharing together with the faster network diffusion suggest that relevant parties are more likely to be informed sooner.

The tasks required before sharing can begin—for example, gathering necessary contact information—may be very expensive in terms of time. Security professionals may prefer different communication methods during an incident; for instance, they may use external email accounts, phone calls, or encrypted mail to avoid tipping off an attacker that defenders are preparing to take action. By design, these contact methods may be difficult for to

[38] Leskovec, "Network Effects and Cascading Behavior."

determine without personal contact. Tasks like these, while expensive, can take place long before an incident begins. Social sharing indeed provides opportunities for such tasks to be completed: forming a personal relationship with one's security colleagues almost certainly involves the exchange of preferred contact information, for instance.

4.3.2 Effectiveness Against Different Attack Types

As we explored at length above, industry-based groups are best suited to respond to semi-targeted attacks. In this case, social groups may be more or less effective depending on the mix of industries and institution types represented in the social network. Groups that remain within industry boundaries may perform no better than an industry ISAC in terms of eventual information diffusion. However, social networks often involve different types of connections, and a variety of distinct communities can be "merged" through critical people.[39]

Highly targeted attacks and generic attacks are not the strong point for this type of arrangement, but social contact with others responsible for the security of diverse organizations may lead to better security knowledge in general. Although sharing across communities may be limited if the originator of a tip asks receivers to keep it secret, it is likely that certain details could be shared without imperiling anonymity. Such arrangements thus could perform well against the group of semi-targeted attacks that cross industry boundaries.

4.3.3 Other Qualities

Just as a sharing participant could simply ignore alerts from a watch floor, they could ignore calls or emails from an unwanted sharing participant. They similarly could run their network without any new information from partners, and are under no obligation to set policy based on their social network's opinions. So, this solution has inherently good distributed survivability and local autonomy. As in the watch floor model, false input is prevented by

[39] Yang and Leskovec, "Overlapping Community Detection at Scale: A Nonnegative Matrix Factorization Approach," Section 1.

social trust and norms.

It is likely that, under load, performance will degrade since the social networks built by these arrangements only include a small number of members per participating institution. For instance, the one participant might be the executive responsible for security, which yields benefits—they have the power to act on shared information and to, themselves, share—but also downsides—they may be busy or unavailable during an incident. So, this system is not strongly reliable under some common circumstances.

4.4 Distributed, High-Volume Data Sharing

The previous sharing designs rely heavily on human-mediated steps. While it is not unheard of in computer system design to assign some particularly difficult tasks to humans, it is somewhat surprising to computer scientists who are accustomed to mostly using computers when solving computer security problems. In recent years, the security and policy communities have begun to consider information sharing paradigms that call for high-volume, automated sharing of security data. This is alluded to in many ways, including calls for initiatives like Continuous Diagnostics and Monitoring, which is not necessarily a sharing scheme, but a machine-enabled analysis technique;[40] machine-readable indicators and other automated tools to enable information sharing;[41] and automated sharing programs like the Enhanced Cybersecurity Services program, which facilitates automatic signature sharing, at least unidirectionally from government to industry.[42] Promising initiatives, such as the FS-ISAC "Cyber Intelligence Repository" and the Georgia Tech Research Institute "Titan"

[40] For instance, p. 8 or the DE.CM ("Detect: Security Continuous Monitoring") series of recommendations of National Institute of Standards and Technology, "Framework for Improving Critical Infrastructure Cybersecurity, Version 1.0."

[41] Such efforts include The MITRE Corporation's TAXII (Connolly, Davidson, and Schmidt, "The Trusted Automated eXchange of Indicator Information (TAXII$^{\text{TM}}$)"). As noted in subsection 2.1.1, Presidential Policy Directive 21 encourages the development of information systems, interoperable formats, and redundant systems ("Presidential Policy Directive 21: Critical Infrastructure Security and Resilience," "Three Strategic Imperatives" #2).

[42] These programs are described in detail in subsection 2.1.2.

Project, have adopted this mindset, and implementation is ongoing at the time of writing.[43]

Actors are individual members of the network security teams at participating institutions.[44] Actors communicate minimally evaluated information about hosts, programs, and other information to sharing partners. They are encouraged to use automatic means to generate these low-confidence data points, which often take the form of simply *facts* rather than complex *judgments*; for example, "IP address 1.2.3.4 sent us a malicious document" rather than "we have determined 1.2.3.4 to be controlled by foreign actor x attempting to exfiltrate y type of data." Communications could involve a central hub for analytic support or to enforce certain privacy or anonymity protections. Alternatively, communications could be over a peer-to-peer structure that allows for distributed storage of data and efficient lookups.[45] Another model might follow the design used by the Internet's routing fabric—the way that computers decide what steps to take to reach other computers—is based on distributed protocols that provide eventual approximate convergence of global state.[46]

More confident, actionable information would be derived computationally from either expert judgments from a trusted set of users (similar to antivirus definitions today) or fully automatically by the derivation of a risk measure based on the suspicious inputs found in the shared data. In addition, the collected corpus of network traffic data, appropriately anonymized, could be shared with academic researchers to develop new algorithms for the identification of malicious activity—real-world data is difficult to obtain, constraining the development of new security techniques.[47]

[43] Financial Services - Information Sharing and Analysis Center, "Cyber Intelligence Sharing Vision"; Georgia Tech Research Institute, "The Titan Project."

[44] If desired, the sharing role may be limited to a subset of the security staff per institution without imperiling the following analysis.

[45] The broadcast method could be, if using a distributed storage structure, completed in $O(\log n)$ bandwidth using methods found in the systems literature. For example, a protocol like that of El-Ansary et al., "Efficient broadcast in structured P2P networks" is being used commercially by BitTorrent, Inc.'s Sync technology, `http://www.bittorrent.com/sync/`.

[46] Among these protocols is BGP: Rekhter, Li, and Hares, *RFC4271: A Border Gateway Protocol 4 (BGP-4)*. Neighboring routers exchange information about what sections of the Internet they can provide access to, and eventually each one determines how it will direct traffic based on this peer-to-peer input.

[47] Sommer and Paxson, "Outside the Closed World: On Using Machine Learning for Network Intrusion Detection," subsection III.E.

Our model of this sharing design is presented in Algorithm 3.

4.4.1 Runtime and Effectiveness

This analysis was spurred by, and is a reaction to, the pervasive idea that the way to solve cyber threat sharing is the same as to staff more 24-7 watch floors, create more task forces, train more "cyber professionals", and otherwise put more humans "in the loop".[48] There is a certain idea, accepted without proof, that humans can effectively curate—and then share—the vast amounts of data involved in network management and security, and even that sharing cannot be effective without it.[49]

Computer scientists tend to disagree with such impulses, searching instead for ways for computers to assist humans in making difficult decisions. Illustratively, current research efforts in computer networking recognize that humans are incapable of understanding the state of complex systems—even single corporate networks—and design simple abstractions that hide the complexity of implementation so that administrators can have a fighting chance to get it right.[50] In contrast, sharing organizations call upon multiple different humans to interpret data in isolated networks, properly merge that information together with tips from other network operators, and form coherent judgments about what security measures to take and with whom to share information.

Further, computer network exploitation and attack can escalate very quickly: an attacker who has breached a network can rapidly shift from expanding access and establishing

[48] Putting something "in the loop", a common phrase in computer systems design, indicates the addition of that thing into the operation of the process. "Putting a human in the loop" further refers to placing a step involving a human into an otherwise automated process, perhaps for quality control or to reserve certain decision-making authority to humans and not their computer assistants.

[49] The scale of this data is truly huge. Networking company Cisco, for instance, receives 20 terabytes of data each day for analysis, and this is only a subset of the data generated by its security analysis centers. Most personal computers today could store less than 10% of the data Cisco receives each day. (Woodie, "A Peek Inside Cisco's Hadoop Security Machine"). Machines on the Internet generate thousands of log alerts per day just from the "background noise" of scanners and generic attackers (based on this author's experience).

[50] They often don't get it right, even if they're experienced, due to the extraordinary difficulty of understanding interactions between all devices on large networks. See Reitblatt et al., "Abstractions for Network Update."

Algorithm 3 Idealized Algorithm for High-Volume Data Sharing

```
function Generator
    loop
        repeat
            record security-relevant information
        until timer expires
        send information to aggregator
        reset timer
    end loop
end function
function Consumer
    loop
        if new data is available then
            update local copy of information
            generate new rules and judgments
            if potential threat is found then
                alert security staff to issue
                while staff are working do
                    send information about flagged indicators to aggregator
                    use other sharing mechanisms to inform partners if serious
                end while
            end if
        else
            wait
        end if
    end loop
end function
function Aggregator
    loop
        repeat
            wait for information from generators
        until new information sent by a generator
        update summary statistics
        optionally perform analytics to identify potential threats
    end loop
end function
```

a foothold to denying service and causing damage.[51] Data is by definition exfiltrated at network speed, not human speed, and the compounded delay of detection, analysis, report writing, and sharing could allow an attacker to succeed before they are even noticed.[52]

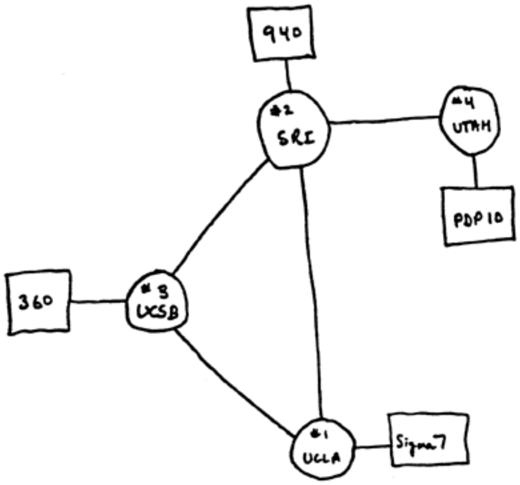

Figure 4.1: Computer History Museum, *1969 4-node ARPANET diagram*

There are two main advantages to this method, both centered around sharing more and sharing information that is individually less vetted. Analysis can proceed with more inputs, instead of the filtered, potentially incomplete, possibly non-machine-readable results shared from a member organization to an ISAC or similar sharing organization.[53] As we began discussing above, computing with more relevant data is likely to produce better results than even the most clever algorithms working on incomplete data. Such success is seen in the "unreasonable effectiveness of data" identified by Google researchers, which demonstrated that adding more data, even if it is less curated, can often yield better results faster than they could be obtained otherwise.[54]

[51] Junio, *Away from the Keyboard? The Challenges of Command and Control over Cyber Operations*.

[52] While humans are ultimately controlling the malicious software, the point is that security decisions need to be made at machine speed; humans can't, for instance, review all outgoing network traffic before allowing it to pass. Such a system would break nearly all network protocols.

[53] Non-machine-readable information requires a human to translate it into technical rules or signatures that can be used in security products. This is a nontrivial hassle that impedes effectiveness and timeliness.

[54] Halevy, Norvig, and Pereira, "The unreasonable effectiveness of data."

4.4.2 Reliability and Other Qualities

As we've observed in the other cases, the sharing system is not a single point of failure for most computer network defense architectures: while it is a good addition, it is not so critical that defensive activity stops if it fails. So, this system performs similarly to the others in terms of *reliability*, but in this case reliability depends on the system's architecture more than on the presence of a staff in a watch floor or the ability to reach a colleague. As in the other solutions, *distributed survivability* is a side-effect of this reliability.

4.4.3 Performance Under Load

The other qualities do exhibit some differences from the previous two cases, however. Under load, this system will likely scale much better—while both this and a central watch floor share an a workload complexity that grows at least linearly with the number of incoming reports, computers provide significantly smaller "constant factors" than human analysts. So, under extreme load, the computers will run a bit hotter and perhaps slightly slower, but their performance will degrade gracefully, depending on machines' capacity rather than humans' capacity to deal with crisis events. (Humans also have needs that do not affect computers, like the requirement for periodic rest and for times of reduced stress between demanding events.)

4.4.4 Local Autonomy and Robustness

An automated system still allows local decision making authority. In fact, it may even provide a more expressive way of specifying local policy, since the watch floor and social network models of sharing tend to involve only indicators that are known, or at least strongly suspected, to be malicious. In contrast, a local policy could specify more of a risk tolerance, quarantining traffic from computers that have been observed probing other networks but allowing through other computers or types of traffic.

Robustness against false input is a bit of a harder challenge, and will depend on the specifics of the sharing system's design. However, it would be reasonable to enforce a similar type of pre-membership investigation or qualification process, after which a member is given access to contribute and receive security event data. Technical means could also be employed to deemphasize reports from sources that are frequently questioned by others, or some similar scheme from the reputation systems literature.[55]

4.5 Conclusions

A summary of the results of the three cases just analyzed is presented in Table 4.1.

At their core, today's security problems are truly failures to *scale*. In the days when many of the protocols we use today were developed—the late 1980s—computer networks were literally small enough to draw on a (very small) sheet of paper (see Figure 4.1). Administrators at least had the hope of knowing the real-life identity behind each computer connecting to their network, and could conceivably assure themselves of their own systems' security posture. All of the computers in the entire network used to be listed in a single file, managed manually by a single person and periodically redistributed.[56] Today's networks comprise many more computers per security professional, and even simple lunch-time web browsing can cause connections to hundreds of other networks, some of which may never have been observed by the administrator before.

A broader implication of this structural assumption is that cyber threat information sharing should rely on a small number of high-confidence data points—for example, *known* bad addresses or programs—versus a high number of low-confidence data points, which would be analyzed. The above analysis suggests that sharing more minimally-filtered data could, in fact, significantly improve network security.

More generally, allowing computers rather than people to make security judgments

[55] This literature was briefly discussed in subsection 3.1.4 when dealing with security models.

[56] Leiner et al., "A brief history of the Internet."

is vastly more scalable; it is significantly easier to add more computers to a compute cluster than it is to hire more highly trained computer security professionals into government or ISAC bureaucracies.[57]

[57] Evans and Reeder, "A Human Capital Crisis in Cybersecurity: Technical Proficiency Matters," vi; Office of Inspector General. U.S. Department of Homeland Security, "DHS' Efforts to Coordinate the Activities of Federal Cyber Operations Centers," 10.

[58] Remember, this is higher volume, lower confidence/individual value data shared automatically.

[59] But not to the level of a human's degraded performance.

Table 4.1: Summary of Sharing Scheme Performance

Design Quality	ISAC/ Watch Floor	Private Association	Distributed Sensors
Convergence Time	Limited by human analysis steps at multiple stages	Limited by human analysis at each sharing step	Machine speed
Probability of Information Spread	Good to members	Good within social circle	Nearly certain to all[58]
Performance versus Generic Attack	Probably redundant	Probably redundant	Possibly redundant, but may help discover attacks
Performance versus Semi-Targeted Attack	Good strategically and operationally	Good strategically and operationally	Good operationally
Performance versus Highly Targeted Attack	Poor	Possibly higher cost to adversary	Likely higher cost to adversary
Reliability	Depends on central hub	Depends on personal communication methods	Depends on system architecture
Distributed Survivability	Good	Good	Good
Performance Under Load	Poor—central human analysis bottleneck	Degraded—security personnel busy responding to incident being shared	Degraded technically[59]
Local Autonomy	Good—can ignore reports	Good—can ignore calls or emails	Good—can still set independent security policies based on aggregated data
Resistance to False Input	Assured by institutional trust and organization vetting	Assured by personal trust	Needs to be assured by contributor vetting or by technically de-emphasizing suspect data

Chapter 5

Conclusions

“There are two possible outcomes in cybersecurity for the United States. We can continue to pursue outdated strategies and spend our time describing the problem until there is some crisis. . . . Alternatively, we can take action on measurably effective policies.”

CSIS Commission on Cybersecurity for the 44th Presidency
(January 31, 2011: Lewis et al., “Cybersecurity Two Years Later,” 15)

“The government is using a ‘Ford sedan’ policymaking system to manage the cyberspace ‘Porsche’ system.”

Professor Harvey Rishikof
(Paraphrased in Rosenzweig, “Cybersecurity and Public Goods: The Public/Private ‘Partnership’,” 10)

Sharing is a *tool*, not a goal itself. In fact, it is a *very useful* tool for computer security, as suggested by a survey of other types of information sharing and cooperation, and even some evidence from existing security collaborations. However, existing cyber threat information sharing organizations have inherited unhelpful design assumptions from the 1990s. Sharing structures like the ISACs misunderstand the cyber threat environment, improperly assuming that crisis events will be the most pressing threats. This idea, while perhaps appropriate in physical security or law enforcement, simply does not apply in computer network defense; the “slow drip” of a wide spectrum of espionage, data theft, and attack preparations are not best addressed with such a model. While the size of the Internet, the speed of communications links, the amount of data stored on computer systems, and the processing power available at all levels of the economy have expanded by more than an order of magnitude since the beginning of this policy history, the organizations and structures developed to fight cyber threats have not adapted.

By defining the goals of cyber threat information sharing in terms of a new analysis technique—*computational policy*—we can better analyze the strengths of each type of design by applying the insights of decades of computer systems design. With this analysis frame, we can better focus our efforts: should we try to reduce the constant factors inherent in an asymptotically expensive process, or might our energy be more effectively spent reducing the asymptotic complexity of the process? Is a given design even likely to succeed? These questions can be more readily and rigorously answered by drawing on the cumulative experience of those who have built some of the most astonishingly complex yet functional systems in the history of humankind: computer systems. And, such constraints will only grow in importance as technology continues to progress while, even today, we struggle to train and support a cadre of network security experts who can protect existing systems.[1]

The application of this analysis technique suggests that a repertoire of sharing strategies for the variety of threats faced by businesses and government agencies, including a much greater focus on day-to-day sharing of many more minimally-evaluated indicators, would contribute to an improved national computer network defense posture. The critical problem is not that there are too many impediments to sharing; it is that there are too few reasons to even try to overcome the few impediments that remain. As a society, we ask every cyber threat information sharing organization to solve the whole problem of computer security, and receive in return a constellation of agencies and private organizations that are each trying to take on a problem that is too big for them to solve. If instead we more optimally match sharing strategies to the specific goals they are meant to achieve, each organization can concentrate on the goals it is uniquely positioned to address, and overall benefits can increase to the point that companies and agencies will be considered naïve for not getting involved.

Computational policy is, to this author's knowledge, a novel analysis technique for

[1] Evans and Reeder, "A Human Capital Crisis in Cybersecurity: Technical Proficiency Matters," v.

policy prescriptions. Its full potential remains to be explored in applications against more "wicked problems" of policy in the Internet Age, in computer security, but even—especially—in fields outside the traditional boundaries of "the artificial".[2] But, for now, it provides explanatory power when put to the task of analyzing the predominant models of cyber threat information sharing, providing reasons for some of the observed deficiencies in existing organization and suggesting that the innovative, automated efforts being undertaken by well-regarded organizations and computer security researchers are built on a solid theoretical foundation.

[2] Rosenzweig, "Cybersecurity and Public Goods: The Public/Private 'Partnership'," n. 47; Denning, "Computing is a Natural Science," 13.

Appendix A

List of Acronyms

APT	Advanced Persistent Threat
ARPA	Advanced Research Projects Agency
CERT	Computer Emergency Response Team
CFTB	California Franchise Tax Board
CIO	Chief Information Officer
CISAC	Center for International Security and Cooperation (Stanford)
CISO	Chief Information Security Officer
CSIS	Center for Strategic and International Studies
CSO	Chief Security Officer
DC3	Defense Cyber Crime Center
DECS	DIB Enhanced Cybersecurity Services
DHS	Department of Homeland Security
DIB	Defense Industrial Base
DNS	Domain Name System
DoD	Department of Defense
ECS	Enhanced Cybersecurity Services
EO	Executive Order
FBI	Federal Bureau of Investigation
FS-ISAC	Financial Services Information Sharing and Analysis Center
GAO	Government Accountability Office
HSPD	Homeland Security Presidential Directive
IG	Inspector General
IP	Internet Protocol
IS	Information System
IT	Information Technology
IT-ISAC	Information Technology Information Sharing and Analysis Center
ISAC	Information Sharing & Analysis Center
JCSP	Joint Cybersecurity Services Program
MS-ISAC	Multi-State Information Sharing and Analysis Center
NCFTA	National Cyber-Forensics and Training Alliance
NIST	National Institute of Standards and Technology

NSA/CSS	National Security Agency/Central Security Service
NTOC	NSA/CSS Threat Operations Center
OIG	Office of the Inspector General
ONG-ISAC	Oil and Natural Gas Information Sharing and Analysis Center
PDD	Presidential Decision Directive
PPD	Presidential Policy Directive
RAT	Remote Access Tool
SSA	Sector-Specific Agency (PDD-63)

References

Alexander, Keith. "Securing our Government Networks." Speech to the Armed Forces Communications and Electronics Association Cybersecurity Symposium, 2009. http://www.afceadc.org/events/special-events/past-conferences-symposia/cybersecurity-symposium-fy09/presentations/alexander.pdf.

Alperovitch, Dmitri. "Revealed: Operation Shady RAT" (August 2011). http://www.mcafee.com/us/resources/white-papers/wp-operation-shady-rat.pdf.

El-Ansary, Sameh, Luc Onana Alima, Per Brand, and Seif Haridi. "Efficient broadcast in structured P2P networks." In *Peer-to-Peer Systems II*, 304–314. Springer, 2003.

Axelrod, Robert M. *The evolution of cooperation.* New York: Basic Books, 1984. ISBN: 0465021220.

Bankrate.com. "ChexSystems" (October 2008). http://www.bankrate.com/finance/checking/chexsystems.aspx.

Bilge, Leyla, and Tudor Dumitraş. "Before We Knew It: An Empirical Study of Zero-day Attacks in the Real World." In *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, 833–844. CCS '12. Raleigh, North Carolina, USA: ACM, 2012. http://doi.acm.org/10.1145/2382196.2382284.

Booz Allen Hamilton. *Commercial/Civil Cyber Community Snapshot.* http://www.whitehouse.gov/files/documents/cyber/Comm-Civil_CyberSnapshotPoster.pdf.

———. *National Cybersecurity Center Policy Capture.* http://www.whitehouse.gov/files/documents/cyber/CybersecurityCentersGraphic.pdf.

Brearly, Harry Chase. "A Symbol of Safety: The Origins of Underwriters' Laboratories." In *Reputation: Studies in the Voluntary Elicitation of Good Conduct*, edited by Daniel B. Klein, 75–84. Edited from the 1923 version as noted in source. Ann Arbor, MI: U. of Michigan P., 1997.

Brenner, Joel. "Cyber Threat Information and the Antitrust Canard." *Lawfare* (April 2014). http://www.lawfareblog.com/2014/04/cyber-threat-information-and-the-antitrust-canard/.

Cherepanov, Anton, and Robert Lipovsky. "Hersperbot—A New, Advanced Banking Trojan in the Wild." http://www.eset.com/us/resources/white-papers/Hesperbot_Whitepaper.pdf.

Chex Systems, Inc. "Consumer Assistance" (2009). https://www.consumerdebit.com/consumerinfo/us/en/index.htm.

Computer History Museum. *1969 4-node ARPANET diagram.* In "Exhibit: Internet History", 2004.

Connolly, Julie, Mark Davidson, and Charles Schmidt. "The Trusted Automated eXchange of Indicator Information (TAXII^{TM})" (May 2014). http://taxii.mitre.org/about/documents/Introduction_to_TAXII_White_Paper_May_2014.pdf.

Cormen, T. H., C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms.* 3rd. The MIT Press, 2009. ISBN: 978-0-262-03384-8.

Coviello, Arthur W., Jr. "Written Testimony to the U.S. Senate Committee on Commerce, Science, & Transportation" (July 2013). http://www.emc.com/collateral/corporation/coviello-congressional-testimony-2013.pdf.

Defense Cyber Crime Center. "DIB Enhanced Cybersecurity Services (DECS)" (February 2013). https://www.dc3.mil/data/uploads/dcise-pdf-dib-enhanced-cybersecurity-services-procedures_updated-feb-26-2013.pdf.

Denning, Peter J. "Computing is a Natural Science." *Commun. ACM* (New York, NY, USA) 50, no. 7 (July 2007): 13–18. ISSN: 0001-0782. doi:10.1145/1272516.1272529. http://doi.acm.org/10.1145/1272516.1272529.

Department of Justice and Federal Trade Commission. "Antitrust Policy Statement on Sharing of Cybersecurity Information" (April 2014). http://www.justice.gov/atr/public/guidelines/305027.pdf.

Dhillon, Gurpreet, and James Backhouse. "Current directions in IS security research: towards socio-organizational perspectives." *Information Systems Journal* 11, no. 2 (April 2001): 127–153. ISSN: 1350-1917. doi:10.1046/j.1365-2575.2001.00099.x. http://doi.wiley.com/10.1046/j.1365-2575.2001.00099.x.

Diakopoulos, Nicholas. *Algorithmic Accountability Reporting: On the Investigation of Black Boxes.* http://towcenter.org/wp-content/uploads/2014/02/78524_Tow-Center-Report-WEB-1.pdf, December 2013.

DiMaggio, Paul J., and Walter W. Powell. "The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields." *American Sociological Review* 48, no. 2 (1983): 147–160. ISSN: 00031224. http://www.jstor.org/stable/2095101.

Donovan, Fred. "FS-ISAC threat information sharing helped thwart DDoS attacks against US banks." *FierceITSecurity* (November 2013). http://www.fierceitsecurity.com/story/fs-isac-threat-information-sharing-helped-thwart-ddos-attacks-against-us-ba/2013-11-14.

Ellison, R. J., D. A. Fisher, R. C. Linger, H. F. Lipson, T. Longstaff, and N. R. Mead. *Survivable Network Systems: An Emerging Discipline.* Technical report CMU/SEI-97-TR-013. PA: Carnegie-Mellon Software Engineering Institute, 1999.

Evans, Karen, and Franklin Reeder. "A Human Capital Crisis in Cybersecurity: Technical Proficiency Matters" (November 2010). http://csis.org/files/publication/101111_Evans_HumanCapital_Web.pdf.

"Executive Order No. 13636: Improving Critical Infrastructure Cybersecurity." In *Federal Register, Vol. 78, No. 33*, 11739–11744. 2013.

Facebook Security Team. "Protecting People On Facebook." https://www.facebook.com/note.php?note_id=10151249208250766.

Federal Bureau of Investigation. "InfraGard: A Partnership That Works" (March 2010). http://www.fbi.gov/news/stories/2010/march/infragard_030810.

———. "National Crime Information Center." Accessed May 17, 2014. http://www.fbi.gov/about-us/cjis/ncic.

———. "The NCFTA: Combining Forces to Fight Cyber Crime" (September 2011). http://www.fbi.gov/news/stories/2011/september/cyber_091611.

———. "When Off-Line is Better: Another Way to Search Crime Records" (January 2010). http://www.fbi.gov/news/stories/2010/january/ncic_010410.

Fedorowicz, Jane, Janis L. Gogan, and Mary J. Culnan. "Barriers to Interorganizational Information Sharing in e-Government: A Stakeholder Analysis." *The Information Society* 26, no. 5 (September 2010): 315–329. doi:10.1080/01972243.2010.511556.

Financial Services - Information Sharing and Analysis Center. "Cyber Intelligence Sharing Vision" (November 2013). https://www.fsisac.com/sites/default/files/Avalanche%20One-Sheet%2020Nov2013.pdf.

Financial Services Information Sharing and Analysis Center. "FAQs — FS-ISAC." https://www.fsisac.com/faqs.

FireEye Labs. "The PLA and the 8:00am-5:00pm Work Day: FireEye Confirms DOJ's Findings on APT1 Intrustion Activity." *FireEye Blog* (May 2014). http://www.fireeye.com/blog/technical/2014/05/the-pla-and-the-800am-500pm-work-day-fireeye-confirms-dojs-findings-on-apt1-intrusion-activity.html.

Fischer, Eric A, Edward C Liu, John Rollins, and Catherine A Theohary. *The 2013 Cybersecurity Executive Order: Overview and Considerations for Congress.* Technical report. Congressional Research Service, 2013.

Fossum, Lori. *Cyber Conflict Bibliography.* George Washington University Law School, 2013.

Freedberg, Sydney J., Jr. "They're Here: Cyber Experts Warn Senate That Adversary Is Already Inside U.S. Networks." *Breaking Defense* (March 2012). http://breakingdefense.com/2012/03/they-re-here-cyber-experts-warn-senate-that-adversary-is-alread/.

Fung, Brian. "Washington is making is easier for businesses to swap notes on hackers." *The Washington Post: "The Switch"* (April 2014). http://www.washingtonpost.com/blogs/the-switch/wp/2014/04/10/washington-is-making-it-easier-for-businesses-to-swap-notes-on-hackers/.

Gal-Or, Esther, and Anindya Ghose. "The Economic Incentives for Sharing Security Information." *Information Systems Research* 16, no. 2 (June 2005): 186–208. ISSN: 1047-7047. doi:10.1287/isre.1050.0053. http://pubsonline.informs.org/doi/abs/10.1287/isre.1050.0053.

General Accounting Office. "Challenges for Critical Infrastructure Protection," no. GAO-03-233 (February 2003).

———. "Improving Information Sharing with Infrastructure Sectors," no. GAO-04-780 (July 2004).

Georgia Tech Research Institute. "The Titan Project." http://gtri.gatech.edu/ctisl/titan.

Gjelten, Tom. "Cyber Briefings 'Scare the Bejeezus' Out Of CEOs." *NPR* (May 2012). http://www.npr.org/2012/05/09/152296621/cyber-briefings-scare-the-bejeezus-out-of-ceos.

Gordon, Lawrence A, Martin P Loeb, and William Lucyshyn. "Sharing information on computer systems security: An economic analysis." *Journal of Accounting and Public Policy* 22, no. 6 (November 2003): 461–485.

Gow, Brad. "Data security breaches: More reach and frequency requires more diligence." *Zurich* (n.d.). http://www.zurich.com/NR/rdonlyres/C4FC10D0-2156-42F8-84E7-63C3BF69B6B6/0/Tech_Cold2_DataBreach.pdf.

Graham, Robert. "How we know the 60 Minutes NSA interview was crap." *Errata Security* (December 2013). http://blog.erratasec.com/2013/12/how-we-know-60-minutes-nsa-interview.html.

Halevy, Alon, Peter Norvig, and Fernando Pereira. "The unreasonable effectiveness of data." *Intelligent Systems, IEEE* 24, no. 2 (2009): 8–12.

Hancock, Jay. "Actuaries In Denver Will Get First Peek At Obamacare's Full Cost." *NPR shots* (March 2014). http://www.npr.org/blogs/health/2014/03/07/287255108/actuaries-in-denver-will-get-first-peek-at-obamacares-full-cost.

Hausken, Kjell. "Information sharing among firms and cyber attacks." *Journal of Accounting and Public Policy* 26, no. 6 (November 2007): 639–688. doi:10.1016/j.jaccpubpol.2007.10.001.

Higgins, Kelly Jackson. "'Operation Aurora' Changing the Role of the CISO." *Dark Reading* (March 2010). http://www.darkreading.com/attacks-breaches/operation-aurora-changing-the-role-of-th/223900131.

Hodge, Nathan, and Ian Sherr. "Lockheed Martin Hit By Security Breach." *Wall Street Journal* (May 2011). http://tinyurl.com/wsj-lockheed.

"Homeland Security Presidential Directive 7: Critical Infrastructure Identification, Prioritization, and Protection." December 2003. https://www.dhs.gov/homeland-security-presidential-directive-7.

IID. "ActiveTrust." Accessed May 17, 2014. http://internetidentity.com/threat-intelligence/.

InfraGard National Capital Region Members Alliance. "Washington, D.C. – National Capital Region Members Alliance Chapter Website" (2014). Accessed May 17, 2014. https://www.infragard.org/%25252Bck3h5h66wLRAMW1ejxKMGAKIO%25252Bwnz20Sbm%25252F2JH1LSE%25253D!.

Internet Systems Consortium. "Internet Domain Survey, January 2014." http://ftp.isc.org/www/survey/reports/2014/01/.

ISAC Council. "Reach of the Major ISACs" (January 2004). http://isaccouncil.org/images/Reach_of_the_Major_ISACs_013104.pdf.

Jackson, William. "'Destructive' cyber attacks ahead, NSA's Alexander warns." *Government Computer News* (July 2012). http://gcn.com/articles/2012/07/10/alexander-destructive-cyber-attacks-on-way.aspx.

Junio, Tim. *Away from the Keyboard? The Challenges of Command and Control over Cyber Operations.* (Lecture, Stanford University, Center for International Security and Cooperation, 7 March 2013.)

Kaspersky, Eugene. "The contemporary antivirus industry and its problems" (November 2005). http://www.securelist.com/en/analysis/174405517/The_contemporary_antivirus_industry_and_its_problems.

Katz, Jonathan, and Yehuda Lindell. *Introduction to Modern Cryptography.* Chapman & Hall/CRC Cryptography and Network Security Series. Chapman & Hall/CRC, 2007. ISBN: 1584885513.

Klein, Daniel B. "Credit-Information Reporting: Why Free Spech is Vital to Social Accountability and Consumer Opportunity." *The Independent Review* 3, no. 3 (Winter 2001): 325–344. ISSN: 1086-1653.

Krebs, Brian. "Target Hackers Broke in Via HVAC Company." *Krebs on Security* (February 2014). http://krebsonsecurity.com/2014/02/target-hackers-broke-in-via-hvac-company/.

Kushner, David. "The Real Story of Stuxnet." *IEEE Spectrum* (February 2013). http://spectrum.ieee.org/telecom/security/the-real-story-of-stuxnet.

Lee, Ronald D., and Lauren J. Schlanger. "Department of Defense Expands Defense Industrial Base Voluntary Cybersecurity Information Sharing Activities" (May 2012). http://www.arnoldporter.com/resources/documents/Advisory%20Department_of_Defense_Expands_Defense_Industrial_Base_Voluntary_Cybersecurity_Information_Sharing_Activities.pdf.

Legro, Jeffrey. *Cooperation under fire: Anglo-German restraint during World War II.* Ithaca: Cornell U. P., 1995. ISBN: 0801429382.

Leiner, Barry M, Vinton G Cerf, David D Clark, Robert E Kahn, Leonard Kleinrock, Daniel C Lynch, Jon Postel, Larry G Roberts, and Stephen Wolff. "A brief history of the Internet." *ACM SIGCOMM Computer Communication Review* 39, no. 5 (2009): 22–31.

Lemos, Robert. "Companies See Business In 'Doxing' The Adversary." *Dark Reading* (May 2012). http://www.darkreading.com/threat-intelligence/companies-see-business-in-doxing-the-adv/240001335.

Leskovec, Jure. "Network Effects and Cascading Behavior." CS224W Lecture, Stanford University (October 2013). http://www.stanford.edu/class/cs224w/slides/07-cascading.pdf.

Lewis, James A., James R. Langevin, Michael T. McCaul, Scott Charney, and Harry Raduege. "Cybersecurity Two Years Later" (January 2011). http://csis.org/files/publication/110128_Lewis_CybersecurityTwoYearsLater_Web.pdf.

Lim, Dawn. "Pentagon Cyber-Threat Sharing Program Lost Participants." *NextGov* (October 2012). http://www.nextgov.com/cybersecurity/2012/10/pentagon-cyber-threat-sharing-program-lost-participants/59028/.

Liu, Charles Zhechao, Humayun Zafar, and Yoris A. Au. "Rethinking FS-ISAC: An IT Security Information Sharing Model for the Financial Services Sector." *Communications of the Association for Information Systems* 34, no. 2 (2014).

Mandiant. "APT1: Exposing One of China's Cyber Espionage Units" (February 2013). http://intelreport.mandiant.com/Mandiant_APT1_Report.pdf.

———. *M-Trends: Beyond the Breach (2014 Threat Report).* 2014. https://dl.mandiant.com/EE/library/WP_M-Trends2014_140409.pdf.

———. "One Year After the APT1 Report: What Difference Does a Year Make?" In *M-Trends: Beyond the Breach (2014 Threat Report)*, 17–21. 2014. https://dl.mandiant.com/EE/library/WP_M-Trends2014_140409.pdf.

Marti, Sergio, and Hector Garcia-Molina. "Taxonomy of trust: Categorizing P2P reputation systems." *Computer Networks* 50, no. 4 (2006): 472–484.

Matsumura, Toshihiro, and Marc Ryser. "Revelation of Private Information about Unpaid Notes in the Trade Credit Bill System in Japan." *The Journal of Legal Studies* 24, no. 1 (January 1995): 165–187.

Mimoso, Michael. “Adequate Attack Data and Threat Information Sharing No Longer a Luxury.” *Kaspersky Labs Threatpost* (2012). http://threatpost.com/adequate-attack-data-and-threat-information-sharing-no-longer-luxury-111512/77221.

———. “Defenders Still Chasing Adequate Threat Intelligence Sharing.” *Kaspersky Labs Threatpost* (2012). http://threatpost.com/defenders-still-chasing-adequate-threat-intelligence-sharing/102904.

———. “Inside the Targeted Attack on the New York Times.” *Kaspersky Labs Threatpost* (January 2013). http://threatpost.com/inside-targeted-attack-new-york-times-013113.

Moteff, John D. “Critical Infrastructures: Background and Early Implementation of PDD-63,” no. RL30153 (June 2001).

Mueller, Robert S. “Working Together to Defeat Cyber Threats.” Speech to the RSA Cyber Security Conference, San Francisco, CA, February 2013. http://www.fbi.gov/news/speeches/working-together-to-defeat-cyber-threats.

Multi-State Information Sharing & Analysis Center. “MS-ISAC Cyber Security Advisories.” 2014. http://msisac.cisecurity.org/advisories/.

Nakashima, Ellen. “Google to enlist NSA to help it ward off cyberattacks.” *The Washington Post* (February 2010). http://www.washingtonpost.com/wp-dyn/content/article/2010/02/03/AR2010020304057.html.

National Council of ISACs. “Member ISACs.” Accessed May 15, 2014. http://isaccouncil.org/memberisacs.html.

National Institute of Standards and Technology. “Framework for Improving Critical Infrastructure Cybersecurity, Version 1.0” (February 2014). http://www.nist.gov/cyberframework/upload/cybersecurity-framework-021214-final.pdf.

Nelson, Rick “Ozzie”, and Rob Wise. “Homeland Security at a Crossroads: Evolving DHS to Meet the Next Generation of Threats.” *Center for Strategic and International Studies Commentary* (February 2013). https://csis.org/print/41645.

Nooteboom, Bart. *Inter-Firm Collaboration, Learning & Networks: An integrated approach.* London: Routledge, 2004. ISBN: 041532954X.

North Texas InfraGard Members Alliance. “InfraGard – North Texas Chapter Events Home Page” (2014). Accessed May 17, 2014. http://www.ntinfragard.org/.

Office of Inspector General. U.S. Department of Homeland Security. “DHS’ Efforts to Coordinate the Activities of Federal Cyber Operations Centers,” no. OIG-14-02 (2013). http://www.oig.dhs.gov/assets/Mgmt/2014/OIG_14-02_Oct13.pdf.

Office of the Director of National Intelligence. “National Intelligence: A Consumer’s Guide” (April 2009). http://www.dtic.mil/get-tr-doc/pdf?Location=U2&doc=GetTRDoc.pdf&AD=ADA501547.

Office of the Secretary of Defense. "Cyber security information sharing." In *Federal Register, Title 32, Vol. 2, Sec. 236*, 2:546–547. 236. July 2012.

O'Gorman, Gavin, and Geoff McDonald. "The Elderwood Project" (September 2012).

O'Reilly, Tim. "Open Data and Algorithmic Regulation." Chap. 22 in *Beyond Transparency.* Code for America, 2013.

Pennsylvania Department of Transportation. "Driver License Compact Fact Sheet" (July 2011). `http://www.dmv.state.pa.us/pdotforms/fact_sheets/fs-dlc.pdf`.

Perlroth, Nicole. "Hackers in China Attacked the Times for Last 4 Months." *New York Times* (January 2013). `http://www.nytimes.com/2013/01/31/technology/chinese-hackers-infiltrate-new-york-times-computers.html0`.

"Presidential Decision Directive 63: Critical Infrastructure Protection." May 1998. `http://www.fas.org/irp/offdocs/pdd/pdd-63.pdf`.

"Presidential Policy Directive 21: Critical Infrastructure Security and Resilience." February 2013. `http://www.whitehouse.gov/the-press-office/2013/02/12/presidential-policy-directive-critical-infrastructure-security-and-resil`.

President's Commission on Critical Infrastructure Protection. *Critical Foundations: Protecting America's Infrastructures.* October 1997.

PricewaterhouseCoopers LLP. *Key findings from the 2013 US State of Cybercrime Survey.* Technical report. June 2013. `https://www.pwc.com/en_US/us/increasing-it-effectiveness/publications/assets/us-state-of-cybercrime.pdf`.

Rajkumar, Ragunathan (Raj), Insup Lee, Lui Sha, and John Stankovic. "Cyber-physical Systems: The Next Computing Revolution." In *Proceedings of the 47th Design Automation Conference*, 731–736. DAC '10. Anaheim, California: ACM, 2010. `http://doi.acm.org/10.1145/1837274.1837461`.

Reed, John. "DoD-DHS' info sharing program on cyber threats isn't shrinking (updated)." *Foreign Policy: The Complex* (October 2012). `http://complex.foreignpolicy.com/posts/2012/10/09/dod_dhs_cyber_threat_info_sharing_program_isnt_shrinking`.

Reitblatt, Mark, Nate Foster, Jennifer Rexford, Cole Schlesinger, and David Walker. "Abstractions for Network Update." In *ACM SIGCOMM.* 2012.

Rekhter, Y., T. Li, and S. Hares. *RFC4271: A Border Gateway Protocol 4 (BGP-4)*, January 2006.

Research and Education Networking Information Sharing and Analysis Center. "REN-ISAC CSIRT RFC2350 Description." July 2013. `http://www.ren-isac.net/csirt/`.

R.L.G. "Hackers strike at a foe." *The Economist: Shumpeter Blog* (July 2011). `http://www.economist.com/blogs/schumpeter/2011/07/security-breach-booz-allen-hamilton`.

Rosenzweig, Paul. “Cybersecurity and Public Goods: The Public/Private ‘Partnership’.” In *Emerging Threats in National Security and Law*, edited by Peter Berkowitz. 2011. http://www.emergingthreatessays.com.

———. “Public-Private Partnerships for Cybersecurity Information Sharing.” *Lawfare* (September 2012). http://www.lawfareblog.com/2012/09/public-private-partnerships-for-cybersecurity-information-sharing/.

Rumsey, Matthew. “Troubling, Broad FOIA Exemptions Not Limited to CISPA” (2013). Accessed November 18, 2013. http://sunlightfoundation.com/blog/2013/04/23/troubling-broad-foia-exemptions-not-limited-to-cispa/.

San Francisco Bay Area InfraGard Chapter. “SF Bay InfraGard Meetings” (2014). Accessed May 17, 2014. https://www.sfbay-infragard.org/MEETINGS.htm.

Shou, Darren, and Tudo Dumitraş. “Worldwide Intelligence Network Environment (WINE): Symantec’s Data Sharing Environment.” Presentation, 2011 CAIDA Workshop on Network Telescopes. http://www.caida.org/workshops/telescope/slides/telescope1103_wine.pdf.

Sipser, Michael. “Space Complexity.” In *Introduction to the Theory of Computation*, 2nd ed. Course Technology, 2006.

Sommer, Robin, and Vern Paxson. “Outside the Closed World: On Using Machine Learning for Network Intrusion Detection.” In *Proceedings of the 2010 IEEE Symposium on Security and Privacy*, 305–316. SP ’10. Washington, DC, USA: IEEE Computer Society, 2010. ISBN: 978-0-7695-4035-1. doi:10.1109/SP.2010.25. http://dx.doi.org/10.1109/SP.2010.25.

Symantec. “Internet Security Threat Report 2013.” 18 (April 2013). http://www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_v18_2012_21291018.en-us.pdf.

———. “Internet Security Threat Report 2014.” 19 (April 2014). http://www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_v19_21291018.en-us.pdf.

———. “Internet Security Threat Report Appendix 2014.” 19 (April 2014). http://www.symantec.com/content/en/us/enterprise/other_resources/b-istr_appendices_v19_221284438.en-us.pdf.

The Rendon Group. *Conficker Working Group: Lessons Learned.* (Commissioned by the Department of Homeland Security.) June 2010.

Timm, Trevor. “CISPA, “National Security,” and the NSA’s Ability to Read Your Emails” (2012). Accessed November 18, 2013. https://www.eff.org/deeplinks/2012/04/cispa-national-security-and-nsa-ability-read-your-emails.

U.S. Department of Commerce. “Incentives To Adopt Improved Cybersecurity Practices.” In *Federal Register, Vol. 78, No. 60*, 78:18954–18955. 60. 2013.

U.S. Department of Homeland Security. *Information Sharing: A Vital Resource for a Shared National Mission to Protect Critical Infrastructure.* Accessed November 18, 2013. http://www.dhs.gov/information-sharing-vital-resource-shared-national-mission-protect-critical-infrastructure.

———. *Privacy Impact Assessment for the Enhanced Cybersecurity Services (ECS).* Technical report 703. Washington, DC: U.S. Department of Homeland Security, 2013.

U.S. General Accounting Office. "Critical Infrastructure Protection: Improving Information Sharing with Infrastructure Sectors," no. GAO-04-780 (2004).

U.S. Government Accountability Office. "Cybersecurity: National Strategy, Roles, and Responsibilities Need to Be Better Defined and More Effectively Implemented," no. GAO-13-187 (2013).

———. "Information Technology: Federal Laws, Regulations, and Mandatory Standards for Securing Private Sector Information Technology Systems and Data in Critical Infrastructure Sectors," no. GAO-08-1075R (2008).

Vu, Le-Hung, Thanasis G. Papaioannou, and Karl Aberer. "Impact of Trust Management and Information Sharing to Adversarial Cost in Ranking Systems." In *Trust Management IV*, edited by Masakatsu Nishigaki, Audun Jøsang, Yuko Murayama, and Stephen Marsh, 321:108–124. IFIP Advances in Information and Communication Technology. Springer Berlin Heidelberg, 2010. http://dx.doi.org/10.1007/978-3-642-13446-3_8.

Waldspurger, Carl A. "Memory Resource Management in VMware ESX Server." In *Proceedings of the 5th Symposium on Operating Systems Design and implementationCopyright Restrictions Prevent ACM from Being Able to Make the PDFs for This Conference Available for Downloading*, 181–194. OSDI '02. Boston, Massachusetts: ACM, 2002. ISBN: 978-1-4503-0111-4. doi:10.1145/1060289.1060307. http://doi.acm.org/10.1145/1060289.1060307.

White House. *National Strategy for Information Sharing and Safeguarding.* Washington, 2012. http://www.whitehouse.gov/sites/default/files/docs/2012sharingstrategy_1.pdf.

Wing, Jeannette M. "Computational Thinking." *Commun. ACM* (New York, NY, USA) 49, no. 3 (March 2006): 33–35. ISSN: 0001-0782. doi:10.1145/1118178.1118215. http://doi.acm.org/10.1145/1118178.1118215.

Woodie, Alex. "A Peek Inside Cisco's Hadoop Security Machine." *Datanami* (February 2014). http://www.datanami.com/2014/02/21/a_peek_inside_cisco_s_hadoop_security_machine/.

Yang, Jaewon, and Jure Leskovec. "Overlapping Community Detection at Scale: A Nonnegative Matrix Factorization Approach." In *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining*, 587–596. WSDM '13. Rome, Italy: ACM, 2013. doi:10.1145/2433396.2433471.