

The Problem isn't Attribution; It's Multi-Stage Attacks

David D. Clark¹

MIT Computer Science and Artificial Intelligence
Laboratory
Cambridge, MA 02142
ddc@csail.mit.edu

Susan Landau

Radcliffe Institute for Advanced Study
Harvard University
Cambridge, MA 02138
susan.landau@privacyink.org

ABSTRACT

As a result of increasing spam, DDoS attacks, cybercrime, and data exfiltration from corporate and government sites, there have been multiple calls for an Internet architecture that enables better network attribution at the packet layer. The intent is for a mechanism that links a packet to some packet level personally identifiable information (PLPII). But cyberattacks and cyberexploitations are more different than they are the same. One result of these distinctions is that packet-level attribution is neither as useful nor as necessary as it would appear.

In this paper we discuss why network-level personal attribution is of limited forensic value. We analyze the different types of Internet-based attacks, and observe the role that currently available alternatives to attribution already play in deterrence and prosecution. We focus on the particular character of multi-stage network attacks, in which machine A penetrates and “takes over” machine B, which then does the same to machine C, etc. We consider how these types of attacks might be traced, and observe that any technical contribution can only be contemplated in the larger regulatory context of various legal jurisdictions. Finally we examine the costs of PLPII mechanisms.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]

General Terms

Design, Security, Legal Aspects.

Keywords

Security, botnets, DDoS, spam, multi-stage attacks, attribution, jurisdictional concerns.

1. INTRODUCTION

As a result of increasing spam, DDoS attacks, cybercrime, and data exfiltration from corporate and government sites, there have been multiple calls for an Internet architecture that enables better network attribution at the packet layer. The intent is for a mechanism that links a packet to some packet level personally identifiable information (PLPII). But cyberattacks and cyberexploitations are more different than they are the same, and one result of these distinctions is that packet-level attribution is neither as useful nor as necessary as it would appear.

At the same time, there is more than one sort of attribution

mechanism in the Internet. IP addresses are one sort. They identify machines, not people, and by their design they are visible in the network to any device that handles the packet. At the application layer, the ends of a connection may demand person-level identity information (e.g. your bank and you really want to confirm that the other party is properly identified); this sort of attribution has the feature first that if the connection is encrypted the identity information is private to the end-points, and second that the degree of identity (and thus attribution) that is demanded depends on the requirements of the situation. Banks demand strong identification, while sites that give out sensitive health information usually try hard *not* to gather such information. Packet level personally identifiable information (PLPII) would be a new sort of identity mechanism: one that provides strong identity information independent of application, and which would be visible in the network to third parties. Like a license plate on your packets, it would allow them to be traced back to you. This is the extreme of the accountable Internet.

By analyzing a number of different sorts of attacks, we come to the following conclusions:

(i) The most challenging set of attacks to investigate and deter are “multi-stage” attacks in which computer A penetrates computer B, which is used as a platform for penetrating computer C, which in turn is attacks computer D [4]. Multi-stage attacks within a single jurisdiction may permit the imposition of rules that facilitate technical solutions to attributions. We suggest that such technical solutions form a ripe area for research. But solutions to preventing the attacks of most concern, multi-stage multi-jurisdictional ones, will require not only technical methods, but legal/policy solutions as well. Better attribution techniques will neither solve nor prevent such exploitations.

(ii) From the ability to deflect DDoS attacks so that they are not seriously problematic, to using IP addresses to do partial attribution in investigating network-based criminal activities, there are already multiple solutions to various Internet-based attacks in place today that do not depend on PLPII. In particular, IP addresses are more useful than had been thought as a tool for attribution.

(iii) Redesigning the network to accomplish robust attribution would not solve the most serious network-based cyberattacks and cyberexploitations being experienced today, which are multi-stage and multi-jurisdictional. At the same time, technical solutions enabling personal network-level attribution would not only have the potential to create great harm for privacy, human rights, free

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ReArch, November, 2010, Philadelphia, PA, USA.

Copyright 2010 ACM 1-58113-000-0/00/0010...\$10.00.

¹ Clark's effort in this work was funded by the Office of Naval Research under award number N00014-08-1-0898. Any opinions, findings, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the Office of Naval Research.

speech, but also for national-security and law-enforcement use of the Internet².

Our recommendations are that rather than focus future network design on achieving better attribution, better security can be achieved by developing better methods to deter and degrade multi-stage attacks.

We begin in section 2 with a brief discussion of Internet attacks, and in section 3, we consider the role of attribution plays in preventing and investigating these attacks. In section 4, we discuss the potential social costs of strong, network level personal identity mechanisms, and in section 5, we consider potential solutions to the multi-stage attack problem. In section 6, we consider future options for identity in the network.

2. TYPES OF INTERNET-BASED ATTACKS

With each new wave of spam, with each new set of DDoS attacks, whether on Estonia, Georgia, or South Korea, with each large criminal attack, and each new speculation about possible foreign attacks on critical infrastructure, it has become commonplace for there to be a call for “an early-warning system to monitor cyberspace” and to “re-engineer the Internet to make attribution ... more manageable” [12]. But spam differs from DDoS attacks, which in turn differs from cybercrime, which differs from data exfiltration. Understanding those differences clarifies the value that network-level attribution can play in deterring and investigating network-based attacks.

We have found two to categorize cyberattacks that prove useful. *Multi-step attacks* are attacks which occur through several steps that can be widely separated in time. An example of this is a botnet. First the botnet is built by subverting various endhosts. At a later time, the subverted machines are instructed to launch some type of attack. The other category we have found useful to consider is the *multi-stage attack*, and there may be multiple steps in a multi-stage attack. A botnet is an example of such a multi-stage attack, but so is a set of several machines, in which machine A infiltrates machine B to attack machine C. Such attacks are the most technically challenging and complex to deter.

There are many attributes through which one could examine attribution, including structural (what are the different parts of the Internet where attribution would be most useful?), kinds (if users might be identified in some way, what would be the source of the identity?), timing (what are the different roles of attribution before, during, and after an event?). We have found it useful to view attribution along the axis from public to hidden. We start with the most public type of attack: spam.

Most spam leaves a visible trace. The sender’s IP address must be legitimate, and if the spam is offering to sell something, the seller must identify himself. On the other hand, if the spam is only to induce the user to visit a site that may download malware, then we can expect no valid higher-level attribution. Since the final infiltration of the host occurs later only if the web site is visited, spam is a form of multi-step attack. Since the spammer is usually several machines removed from the machine actually sending the spam, it is also a multi-stage attack.

DDoS attacks are designed to be visible. On the other hand, the source of the attack and the preparation is usually carefully hidden. Again, the design of the attack is multi-step (first the botnet is built, and then at an appropriate moment, the subverted machines are instructed to attack) and multi-stage (the bot-master will instruct his bots indirectly to avoid attribution).

“Identity theft,” in which a criminal uses someone else’s personal information to obtain a service, is a criminal activity that has been greatly simplified by the availability of massive amounts of such information through the combination of online databases and the network. The information sought by criminals may be a credit card number, a social security number, a name and address. The full-scale impersonation in which multiple documents are obtained in another person’s name occurs about two million times a year in the United States [15]. Like DDoS attacks, identity theft is a multi-step, multi-stage attack, in which first a system is infiltrated, then personally identifiable information (such as user name and password, or name, address, and credit-card number) is lifted, downloaded perhaps in a multi-stage pattern, and then finally, perhaps, the stolen credentials are used to fool some application-level authentication system. The part of the activity in which the information is surreptitiously lifted, the “theft” in identity theft, is hidden, but the later use of the identity for obtaining goods and services leaves a visible trace.

From an economic and national-security vantage point, the most serious type of Internet-based attack occurs with massive data exfiltration of private information from corporate and government sites. As little as twenty years ago, such thefts typically required having a spy in place, perhaps for many years, developing contacts in the targeted industry. Now the probing can happen at a distance. The first public notice of such types of attacks occurred in 2005, when Time magazine reported a 2004 exploitation of four U.S. military sites [14] in which the same security hole was used to access and download multiple classified files in the space of under eight hours. The files were sent first to servers in Korea and Taiwan, and from there to southern China. (Again, a multi-stage pattern.) Since that time, such type of attacks have proliferated, with the targets being civilian industry, military contractors, and government sites. These cyberexploitations have occurred against sites in the U.S., U.K., Australia, Canada, New Zealand, India, Belgium, and Germany [9].

Depending on the site from which data is being exfiltrated, the same technique may be used multiple times (e.g., as was done in the 2004 “Titan Rain” attacks [14]), or the attack may be individually tailored. Such appears to be the case, for example, in the well-publicized attack on Google, in which the first step was a message with a link to website hosting malware that was sent to a Google employee in China. The employee’s machine was compromised after he visited the website, and this compromise enabled intruders to gain access to development sites within Google [11]. At a later point, significant numbers of files were downloaded first to a machine in Taiwan, and then subsequently apparently to a machine in southern China [10]. The pattern was similar to a number of other massive data exfiltrations described earlier in this paper [9]. Attacks of this type are often individually tailored to the site, which makes prevention extremely complicated. They are also multi-step attacks in which the exploiter gains access to the targeted system, then carefully examines the site to determine the files of interest, which are then downloaded rapidly when the time is deemed ripe. Often this

² Nobody needs anonymity more than a spy.

downloading appears to be to intermediate machines—“dead drops”—perhaps in Korea, Taiwan, or Hong Kong, before the files are downloaded further (perhaps to southern China). The multi-stage nature is designed to confound definitive knowledge of the final destination of the files.

This type of cyberexploitation is the most difficult to investigate. The perpetrator may not only have hidden his tracks through the use of a number of intermediate machines, but also through the use of a number of intermediate jurisdictions, the latter of which severely complicates any investigation.

3. WHAT IS THE ROLE FOR ATTRIBUTION IN PREVENTING ATTACKS?

In the non cyber world, attribution is very powerful in deterring attacks and exploitations. The belief that the same is true in the cyber world, namely that an ability to discover the actor behind cyberattack or cyberexploitation will help deter the action from occurring, drives the idea that packet-level personal attribution will secure the network from Internet-based attacks. This is less true than it might appear. Consider now the cyberattacks and cyberexploitations briefly discussed in the previous section in the context of attribution.

An important aspect of dealing with DDoS attacks is preventing such an attack and mitigating it when it occurs. This is not deterrence, but a different approach with a different set of requirements for attribution.

One can take steps ahead of an attack to dilute an attack's strength. For example, machines that are likely to be targets of attack can replicate their content on other servers. (Akamai claims that their replication services make DDoS attacks ineffective [1]). By shutting off traffic from attacking hosts during a DDoS attack, it is possible to dilute the attack's effect. Not all attack traffic has to be shut off; one has only to degrade the efficiency of the DDoS attack in order to be successful in making it ineffective. All these steps depend on knowing the IP addresses of the infected machines. Having PLPII would not be of additional help over having the IP address, since the machines doing the actual attack belong to unwitting owners. ISPs are now experimenting with sending letters or other notices to owners of machines that appear to be infested, based on their ability to map from IP address to billing address. This action is not done in “real time”, and does not depend on PLPII, but rather on the ability of the ISP to use its private data to map IP addresses to users.

DDoS attacks succeed in part because the actual source of the attack, the botmaster, is hidden. Botmasters take great care to be a number of hops removed from the machines doing the attacks, in particular using intermediate services that do not demand or log identity. The approaches are deliberately designed to avoid attribution of any sort, but if any sort of traceback is to be done, it will depend on the IP addresses of the intermediate nodes, not any sort of PLPII associated with them. One must follow a chain of computers, not a chain of unwitting owners, except perhaps to comply with legal requirements for access to traceback information.

There is a different trajectory in spam cases. Spam has two purposes: induce the user to visit an infected site or induce the user to purchase goods. While the former is typically a set-up for a multi-stage attack, the latter leaves a track for investigators,

namely the merchant from whom the user has purchased goods. Attribution of the actual machine sending the spam is less valuable to investigators than is “following the money” and discovering who is offering the get-rich-fast deals and low prices on medications.

In criminal cases, the eventual goal is to *prove* the case against the individual involved in the commission of the crime. Attribution must identify the individual responsible for the criminal activity, not the machine, and the evidence found must be of *forensic quality* [4]. Superficially, this seems to call for very robust personal-level identity and attribution mechanisms. However, this assumption may be simplistic. As a practical matter, evidence found through on-line forensics is generally much less convincing for juries than actual physical evidence [Landau, personal communication with senior member of the FBI, December 14, 2009]. However, because criminal activities using the network almost always occur for the purpose of making money, there will be a money trail. (This is true even of child pornography.) Following that trail is the most useful route for law enforcement. During the course of the investigation, evidence that is not of “forensic” quality, including IP addresses, is useful in leading investigators to potential suspects.

Consider identity theft, a multi-stage activity that first involves stealing the identity data, then exploiting that data directly by stealing money from accounts, using stolen credit-card numbers to purchase items, etc. or for building false identities. At some point money must be taken as part of the identity theft (the exception to this is if identity theft is used in a spying scheme); at that point, the investigation has hard evidence to use. An illustrative example is the break-in conducted by Russian and Estonian criminals against RBS WorldPay, an Atlanta-based credit-card processing company. With the aid of an insider, the group changed account information to allow high withdrawal limits. But while the initial online activity was theft was of identity information, the goal and final theft was of tangibles (money). More importantly, it was the tracking of fraudulent ATM activities in Tallinn Estonia that was crucial for the initial arrests and cracking of the case [16].

A case with a similar outcome occurred with the theft of patient records from the online pharmacy Express Scripts [8]. Here the purpose of the data exfiltration was blackmail. But as with RBS, while the initial activity was online data theft, the final outcome involved paying the thieves. (In fact, their criminal activities did not succeed). There is no reason to believe that PLPII would have been an important tool for law-enforcement investigators. The idea of following the money is important in almost all investigations of on-line criminal behavior.

Of course, following the money does not always work. One place where it fails is when the data exfiltration is done by the government of a foreign nation. This may involve theft of military or government data for national-security purposes, it may be the theft of political information for diplomatic advantage, it may be the theft of corporate trade secrets or technical work for business advantage. This is not a new sort of problem. Thefts of military and political information has occurred as long as there have been governments, while the thefts of business and technical material from other nations' industries is over a century old. What is new is the ease with which such thefts can occur.

PLPII would generally be of no use here. If the attacker is an insider and has credentials issued by the target, that sort of identity information is not carried at the packet level, but at the application level, and usually is intentionally hidden from observation inside the network. If the attacker is an outsider, any credentials provided, PLPII or application-level, would almost certainly be falsified. Further, assuming that this is a multi-stage attack, any identity information, whether PLPII or application level, would just be that of the unwitting owner of the “last hop” machine. Multi-stage attacks incorporate a form of identity theft.

Thus we conclude that while in such multi-stage attacks, it is very difficult to determine who the attacker actually is, PLPII won't help the situation. What is desired is to trace the attack back to a larger entity—a company, a government or agency, and the like. IP addresses, which cannot be forged in attacks of these sorts, are more likely to be useful than PLPII. This raises such questions as whether it should be easy or hard to map IP addresses to jurisdictions.

4. THE DANGERS OF FULL ATTRIBUTION

An IP address is a form of identifier; how much of a binding it provides to an individual depends upon the ISP and the particularities of its billing system. A coffeehouse offering wireless service will likely have no ability to tie activity to particular users, while an ISP serving an individual with a hardwired system will be able to do so. The latter is what enables organizations pursuing potential copyright infringement to track down the violators within willing jurisdictions. IP addresses, depending on how they are managed and the rules of the jurisdiction within which they are located, can be mapped to a variety of higher-level information—the barriers are not technical but legal. For example, traceback may need to proceed rapidly, but the law moves slowly. This is a problem, but not a technical one. Within a regime that allows the gathering and use of IP-level connectivity information (with protections that might or might not be personally comforting) mechanisms could be imagined that could do rapid traceback.

At the same time, attribution is not always the desired goal. As noted above, some applications (e.g. a web site offering sensitive health information) may desire to demonstrate that they do not gather or have the means to gather personal attribution. As well, there are identity systems that provide pseudonymity, whether for providing comments on blogs or for disabling linkages between different online transactions.

PLPII destroys the ability to do Internet activity anonymously. In that context, we draw attention to a 1969 U.S. Supreme Court decision that stated, “If the First Amendment means anything, it means that a State has no business telling a man, sitting alone in his own house, what books he may read or what films he may watch. Our whole constitutional heritage rebels at the thought of giving government the power to control men's minds.” [13] Given the increasing prevalence for content to appear on the Internet --- and sometimes only on the Internet --- it would seem that mandated PLPII would run counter to this principle of “the right to read anonymously” [5].

In order for such “anti-attribution” behaviors to be fully untraceable, there needs to be anonymity not only at the application layer, but also at the IP layer [17]. One way to do so is

through Tor³, which is used by activists and dissidents in the U.S. and abroad, whistleblowers, and journalists and investigative reporters. Tor was developed by the U.S. Naval Research Laboratory, and is also used by law-enforcement officers to hide their government addresses as they investigate chat rooms and online criminal activities, and by military personnel seeking to hide their affiliation when communicating from insecure sites. These government investigators benefit by a broad use of the technology since broad usage provides cover for *all* the users of the system [6]. Any robust scheme for PLPII would be deadly to these sorts of circumstances, since this information would not be under the control of the end-nodes but would be revealed to third parties (e.g. repressive governments) observing the network.

There is a perhaps intrinsic tension between the desire for attribution and deterrence on the one hand, and the need to provide for core elements of civil society—free, voluntary and private association—on the other hand. Resolving this tension might well be the holy grail of security research, always beyond our reach. But a balance must be our goal.

5. HANDLING MULTI-STAGE ATTACKS

Multi-stage operations in and of themselves are not evidence of nefarious activity; they often occur as part of legitimate business operations in which certain aspects are done by different systems (e.g., federated identity management or composed Web 2.0 applications). The problem in multi-stage computations occurs when intermediate machines have been infiltrated and corrupted in such a way that the machine is no longer following the instructions intended by its owner.

There are several ways that multi-stage attacks can be deterred or degraded. One can make it harder to penetrate and keep control of intermediate machines, and one can try to trace back attacks when they occur. Suppose, as before, we are in the situation where computer A has penetrated and controls B, which in turn has penetrated and controls C, which has penetrated and is controlling D. If one is to determine that machine C is really being controlled by machine B or even A, one has to be able to study what C is actually doing. Investigating and preventing such a systemic failure requires the ability to either to examine C, its communications, or both; studying C's communications may mean simply studying the transactional information rather than full content. Tools to facilitate traceback of multi-stage attacks could be imagined. For example, every user could be encouraged to “protect its reputation” by installing some sort of enhanced home router that logs all incoming and outgoing packets. One can imagine technical approaches to analysis of this stored traffic, e.g. the research on stepping stones (see, for example, [2]).

If the use of machine C is not a simple pass-through of packets, other techniques might be used to try to diagnose the source of the attack that has corrupted the machine, and (if it is part of a bot-net) how it is being controlled. See, for example, [3]. One could imagine, for example, aggregating logs from a large number of infested machines, a similar number of clean machines, and looking for differences in their patterns of communication. Such a technique could give a measure of privacy protection while yielding hints about what is going on.

³ See www.torproject.org.

Schemes for traceback raise issues of cross-jurisdictional rights and responsibilities, but they also raise issues of personal privacy. An interesting research question concerns the possibility of doing attack detection and resolution across jurisdictions and across ISPs in ways that give some degree of anonymity to intermediate actors.

If a user has not acted to capture this information, and it appears that his machine is infested with malware, regulation could permit (or require) that his serving ISP log his traffic. In other words, regulation could define that the penalty for failure to self-protect is loss of privacy, which is a punishment that might fit the crime, as opposed to more draconian alternatives such as a total disconnection or quarantine. ISPs seem to be able to detect that a machine is part of a bot-net; their dilemma is that they do not know what they can or should do with that information. This is both a technical and policy issue, having to do with privacy and the rights of the various actors.

Of course, examining this information would require either the permission of C's owner or legal authority from the jurisdiction in which C resides. Thus we come to a clear and very important point: **the multi-jurisdictional multi-stage attack problem is a problem that involves both legal and policy tools as well as technical ones; it cannot be solved by technical means alone.** In particular, this means that the problem of cyberexploitations carried out against U.S. industry and government sites is unlikely to be amenable to technical solutions.

This analysis has been framed in the context of the current Internet. Our first future-looking conclusion is the negative one that a strong attribution mechanism along the lines of PLPII should *not* be an objective. But another future-looking objective might be to design a different sort of network that makes multi-stage attacks harder to realize or easier to trace. For example, some future network proposals, such as PSIRP⁴ and Named Data Networking [7], use information dissemination rather than inter-computer communication as their basic service level. We have not attempted to analyze these proposals from the perspective of multi-stage attacks, but we urge that this sort of analysis should be a part of the process of design and evaluation of new architecture proposals.

6. THE ROLE OF ATTRIBUTION

Our conclusion is not that attribution has no role to play in the Internet, but rather that a public, personally identifiable packet-level mechanism is neither appropriate nor particularly needed or helpful. As per the previous section, in many cases it would be counterproductive and destructive.

With respect to the current Internet, its entrenched nature suggests that a radical idea such as PLPII is not likely to happen. However, in a future architecture, one could imagine a proposal to make person identity management not an application-level service to be used as needed, but a "network-level" mechanism built into the core of the architecture as a mandatory tool. To us, this would be a bad idea.

First, we have argued that "forensic quality" PLPII is not actually desired by law enforcement. They want information of a quality that can guide an investigation, not that they bring into court.

Second, the fundamental question about PLPII would be "who issues the credentials". At the application layer, identity management is a decision private to the parties. When a user contacts a bank, the bank decides what certificate to offer, and (perhaps in the future) what certificates to accept. Or it might issue its own user certificates. Those are application-specific decisions.

But if every user were required to have and present PLPII in their packets, who would issue and vouch for that information? The home country of the user? If so, we might conclude that the resulting PLPII was essentially useless. A credit card company? If so, we disenfranchise a major part of the world that does not have access to a credit card.

These are not technical questions, but social and policy questions. Technologists should seek satisfactory answers to such questions to guide the design of any scheme for PLPII. We do not see satisfactory answers.

In contrast to mandatory deployment of PLPII we should consider other innovations that would make appropriate sorts of attribution easier to accomplish. For example, we alluded above to the idea of allocating addresses to countries so that addresses could more easily and robustly be linked to a jurisdiction. Careful thought would be required to consider whether such a change would be in the best interest of a majority of the actors on the Internet.

We also suggested that regulation could hold owners of intermediate machines in a multi-stage attack responsible to some degree for the resulting harm of the attack. We suggested a specific approach, in which poor system maintenance would result in a loss of privacy.

Making it harder or impossible to forge source IP addresses would do no harm and some good: it would eliminate a few classes of DDoS attacks. But it would not improve the situation with multi-stage attacks: in such attacks the IP addresses cannot usefully be forged today, since preparing the attack usually require interaction at the level of a TCP connection, which implies two-way packet exchange.

In conclusion, to handle the issues of criminal activities, cyberexploitation, etc., we should focus on other approaches, such as making multi-stage attacks more difficult and costly. And rather than issuing calls for better attribution on the network, we should be designing applications that do a better job of integrating identity and attribution when and only when it is actually needed for the purpose at hand.

REFERENCES

- [1] Akamai Security Capabilities, whitepaper, http://www.akamai.com/d/whitepapers/Akamai_Security_Capabilities.pdf?campaign_id=AANA-65TPAC, visited April 20, 2010.
- [2] Blum, Avrim, Dawn Song, Shobha Venkataraman, 2004. "Detection of interactive stepping stones: Algorithms and confidence bounds". Conference of Recent Advance in Intrusion Detection (RAID), (Sophia Antipolis.)

⁴ See www.psirp.org.

- [3] B. Carrier and C. Shields, 2004. The Session Token Protocol for Forensics and Traceback, *ACM Transactions on Information and System Security*, 7(3); 333-362, August, 2004
- [4] Clark, David D. and Susan Landau, 2010. "Untangling Attribution," *Proceedings of a Workshop on Deterring CyberAttacks: Informing Strategies and Developing Options for U.S. Policy*, National Research Council. September 2010
- [5] Cohen, Julie, 1997. "A Right to Read Anonymously: A Closer Look at 'Copyright Management' in Cyberspace," 12 *Berkeley Tech. L.J.* 161
- [6] Dingedine, Roger and Nick Mathewson, 2004. "Anonymity Loves Company: Usability and the Network Effect," Workshop on Usable Privacy and Security Software,
- [7] V. Jacobson et al, Networking Named Content, CoNEXT 2009, Rome, December, 2009.
- [8] Kirk, Jeremy, 2008. "Data Thieves Blackmail U.S. Drug Company," *techworld.com* (November 8, 2008), <http://news.techworld.com/security/106670/data-thieves-blackmail-us-drug-company/>.
- [9] Krekel, Bryan, 2009. "Capability of the People's Republic of China to Conduct Cyber Warfare and Computer Network Exploitation," Prepared for the U.S.-China Economic and Security Review Commission (2009).
- [10] Markoff, John, 2010. "After Google's Stand on China, U.S. Treads Lightly," *New York Times*, January 14, 2010.
- [11] Markoff, John, 2010. "Cyberattack on Google Said to Hit Password System," *New York Times*, April 19, 2010.
- [12] McConnell, Mike, 2010. "Mike McConnell on How to Win the Cyberwar We're Losing," *Washington Post*, February 28, 2010.
- [13] *Stanley v. Georgia*, 394 U.S. 557 (1969).
- [14] Thornborough, Nathan, 2005. "Inside the Chinese Hack Attack," *Time* (August 25, 2005).
- [15] Synovate, *Federal Trade Commission --- 2006 Identity Theft Survey Report*, November 2007, 4.
- [16] United States Department of Justice, Office of Public Affairs, *Alleged International Hacking Ring Caught in \$9 Million Fraud (November 9, 2009)*, and United States District Court, Northeastern District of Georgia, Atlanta Division, United States v. Viktor Pleschuk, Sergei Tsurikov, Hacker 3, Oleg Covelin, Igor Grudijev, Ronald Tsoi, Evelin Tsoi, and Mikhail Jevgenov, Defendants, Criminal Indictment 1-09-CR-492 (November 10, 2009).
- [17] Xie, Yinglian, Fang Yu and Martin Abadi, 2009. "De-anonymizing the Internet Using Unreliable Ids," *Proceedings of the ACM SIGCOMM 2009 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications* (August 2009), 75-86.