

Sharon Goldberg --- Research Statement

January 8, 2009

As a **security** researcher, my goal is to develop practical solutions that make computer networks more reliable, predictable, and resilient to attack. To address the grand challenges in security research, I believe that we need to go beyond traditional disciplinary barriers. For instance, **cryptography** is crucial to deal with hacked devices and malicious parties, but cannot deal with the fact that nodes in computer networks are often owned by entities with competing economic goals. **Game theory** is an emerging approach that can be used to reason about the way systems behave in the presence of selfish (rational) behavior, but often requires making strong assumptions on the economic motivations of parties in the network. At the end of the day, solutions must be implemented in extremely resource-constrained network devices, so hardware and software **engineering** concerns are paramount. Finally, ensuring that security research makes a real impact requires an investment in technology transfer activities, including **standardization**, **prototyping** or collaboration with industry. Indeed, my research spans all of these areas.

1. RESEARCH OVERVIEW

When we purchase an item from Amazon.com, traditional cryptography prevents attackers from seeing our credit card numbers or impersonating the Amazon website. But how can we ensure that our request actually arrives at the Amazon.com server, without being dropped or corrupted along the way?

This is exactly the challenge we address when we consider securing Internet's interdomain routing system. The interdomain routing system enables communication between autonomous systems (e.g. Princeton's campus network, AT&T's global backbone network, and China Telecom's residential network). Each autonomous system has different economic goals, each controls a sophisticated network of devices, and each is prone to malfunction, misconfiguration, or attack by malicious parties. Despite these complex relationships, the interdomain routing system was originally designed under the assumption that nodes can trust each other. Furthermore, the system is notoriously resistant to change. This is mainly due to the difficulty of coordinating and convincing autonomous systems to upgrade to one of the new security protocols developed over the last decade. To advance efforts to secure the system, my dissertation research addressed the following two questions:

- Q1. What is the right layer of the system to secure?
- Q2. What are the right tradeoffs between security and efficiency (e.g. storage and computation overhead)?

To address the questions above, I considered separately the effectiveness of securing the following two layers of the interdomain routing system: the routing protocols, used to set up paths through the Internet, and the data-plane mechanisms, used to forward packets along the paths set up by the routing protocols.

- R1. In [1] we used a **game-theoretic analysis** to show that secure routing protocols alone are *not* sufficient to prevent selfish nodes from lying about the paths that data packets take through the network. Thus, nodes may potentially forward their customer's traffic via paths that drop or corrupt packets.
- R2. Thus, I turned my attention to secure data-plane mechanisms. In [2], we used novel **cryptographic and data-streaming approaches** to design lightweight protocols that detect packet loss and corruption on a path through the network, even when some nodes on the path are adversarial.
- R3. I subsequently interned at Cisco Research to **prototype** one of our protocols from [2] on a Cisco software platform. I also collaborated with an IETF working group [6] to **standardize** some of our innovations (see also our security analysis [4]).
- R4. In [3], we took the security guarantees in [2] further, by considering protocols that also enable a sender to locate the adversarial node responsible for dropping or corrupting its packets. We used **cryptographic proof techniques** to argue that *any* such protocol requires unreasonable overhead.

These results suggest that our efforts are better spent on protocols with weaker security guarantees, such as those we considered in [2].

Below, I discuss the details, implications, and open questions that arise from my results on secure routing protocols (Section 2) and data-plane mechanisms (Section 3), and present some future directions (Section 4).

2. SECURE ROUTING PROTOCOLS

How can we prevent a routing event in Pakistan from disrupting communication between Canada and Mexico? More precisely, how can autonomous systems avoid sending packets along costly, inefficient, or insecure paths when better paths through the Internet are available?

Honestly announcing routes in routing protocol messages. In recent years, practitioners' efforts to solve this problem have focused on securing the interdomain routing protocol (i.e., the Border Gateway Protocol (BGP)). In BGP, nodes discover paths through the Internet via announcement from their neighbors. Thus, in [1], we considered the following security requirement: ensuring that the paths announced in the routing protocol messages match the paths that packets actually take through the Internet (i.e., in the data plane). If this requirement was satisfied, a node could rely on routing protocol messages to choose efficient routes for its packets, or to avoid sending its traffic via intermediate nodes that are known to drop or corrupt packets.

Because this requirement is quite strong, we investigated whether it can be met in a weaker (but still realistic) 'threat model' where nodes are assumed to be rational economic actors. (This should be contrasted with the more common approach in the security literature, where misbehaving parties are assumed to be arbitrarily malicious). Assuming that players are rational allows us to use game-theoretic tools to reason about when nodes have incentive to lie in routing protocol messages. Specifically, we used tools from distributed algorithmic mechanism design (DAMD) to look for conditions under which we could *prove* that nodes have no incentive to send routing protocol messages that lie about the routes that they actually use for their packets. Our approach differs from earlier attempts within the DAMD framework, because, for the first time, we considered nodes that are interested in attracting incoming traffic in order to increase revenues.

Our analysis yielded some surprising results. We showed that even if we assume that nodes are rational, and even if they all use the strongest known secure routing protocol (Secure BGP [5]), some nodes may still benefit from lying in their routing messages. (While we did prove that nodes have no incentive to lie in some special cases, to do this, we had to make unrealistic assumptions on the routing policies at each node.) Our work [1] also includes examples of situations where nodes can reap financial benefits by lying, and I am currently using empirical approaches to explore the impact these lies can have on the network.

3. SECURE DATA-PLANE MECHANISMS

Our work in [1] suggests that nodes should not rely on secure routing protocols to choose good paths through a network. Thus, we focus instead on securing the mechanisms responsible for packet delivery. Here, the question of security vs. efficiency is paramount, because these mechanisms need to keep up with the extremely high packet-processing speeds in the core of the Internet.

Detecting the adversary. How can a node *detect* if the traffic it sends on a path is being discarded or corrupted? In [2], we designed secure and lightweight measurement protocols for this task. We assumed the strongest possible adversarial model (where adversarial nodes on the path are fully aware of the details of the protocol, and can, add, drop, modify and delay packets at will) and designed highly efficient protocols for high-speed routers on multi-Gbit/sec links. For instance, our "secure sketch" protocol can monitor up to a billion packets without marking or modifying existing packets, and requires only two control messages and 200-600 bytes of storage at the sender and receiver only.

Prototyping at Cisco. I spent the summer of 2008 at Cisco Research, prototyping our secure sketch protocol on a Cisco software platform, working with product engineers to analyze use cases, and preparing a full technical specification of the protocol.

Standardization efforts. Related to our work in [2], we reviewed an IETF charter on packet sampling [6] that proposed solutions related to protocols we developed in [2]. Unfortunately, constructions proposed in the charter contained a number of security flaws. We addressed these in [4] by developing a formal

cryptographic model of security for packet sampling, presenting a number of secure approaches, and explicitly breaking many of the constructions proposed in the charter. After presenting our work to the IETF working group, the charter was revised according to our comments.

Localizing the adversary. While it is useful to enable senders to *detect* packet loss and corruption on path, it is even more useful to be able to *localize* the adversarial node responsible for tampering with packets. Unfortunately, all known solutions to this problem (see [3]) require high overhead at *every* node on the path (i.e., keys, cryptographic operations, storage). In [3] we studied this problem by developing a formal cryptographic model that allowed us to find flaws in previous work, to design new secure protocols, and also to apply sophisticated proof techniques (reductions, black box separations) to argue that the large overhead used by these protocols at *every* node is both inherent and unavoidable.

Secure Internet measurement. The success of our work on measurement protocols in [2] (and our negative results on localizing adversarial nodes in [3]) motivate more work on cryptographic protocols that *detect* metrics other than packet loss and corruption. I am very interested in developing protocols that securely monitor latency, delay variation (jitter), and packet reordering on an Internet path.

Fast hash functions. Furthermore, the secure measurement protocols in [2] require hash functions with very *weak* cryptographic properties to be computed at extremely high speeds. Thus, I am very interested in designing new hash functions that are optimized for this unusual tradeoff between security and efficiency.

4. FUTURE DIRECTIONS

As academics, we are uniquely positioned to provide an unbiased analysis of the complex security issues that arise in computer networks. Furthermore, we have the opportunity to draw on rigorous techniques to deal with these issues, and to reach out to various stakeholders (industry, standards bodies) to ensure our solutions have impact. My dissertation research has taken preliminary steps to deal with some of the security issues that arise in interdomain routing, and has opened a number of fascinating new questions related to the design of *network architectures* that withstand selfish or adversarial behavior. Answering these questions will require techniques that bridge between traditional disciplines (e.g., cryptography, game theory), for instance:

Deployment of security protocols. Deploying new network security protocols is a Catch-22 problem. On the one hand, the security properties of many protocols (e.g. Secure BGP [5] and the localization protocols in [3]) only take effect after they have been adopted by a large number of nodes. On the other hand, nodes will only take on a costly upgrade if the security benefits of the new protocol have already taken effect. Inspired by recent work on the adoption of new technologies by economically-motivated nodes in social networks [7], I intend to develop analytic models to yield insights on strategies for deploying network security protocols.

Intersection of incentives and cryptography. In [3] we prove that *any* localization protocol can, at best, localize a link *adjacent* to the adversary, but cannot distinguish which of the two nodes bordering that link are responsible for tampering with traffic. Thus, an economically-motivated node might prefer to deviate from the localization protocol to ensure that that he is not blamed when his neighbor corrupts packets. However, the model we used in [1] does *not* capture the notion of rational players; like most works in the cryptographic literature, we rely on the assumption that nodes on the data path are either adversarial (arbitrarily malicious) or honest (always follow the protocol without deviation). While this issue has recently been addressed in the cryptographic literature (e.g., for secret sharing, Byzantine agreement, etc.), more work is required before we fully understand how to deal with both rational and malicious players. I am very interested in increasing our understanding of this issue, beginning with the question of rational actors in secure localization protocols [3].

Reputation in networks. The networking literature sometimes (informally) claims that nodes in a network are unlikely to misbehave because they value their long-term reputations. I am fascinated by the idea of formally reasoning about this argument, and using it to inform the design of network security protocols. We made a preliminary attempt to do this in [1] with our notion of “loop verification”, (that formalizes the idea that a node will be honest if it fears getting caught lying), but I intend to contribute more to this area, potentially by adapting or generalizing results from the economics literature on repeated games.

5. REFERENCES

- [1] **S. Goldberg**, S. Halevi, A. Jaggard, V. Ramachandran, R. Wright. “Rationality and Traffic Attraction: Incentives for Honest Path Announcements in BGP”. *Proc. ACM SIGCOMM 2008*..
- [2] **S. Goldberg**, D. Xiao, E. Tromer, B. Barak, J. Rexford. “Path-Quality Monitoring in the Presence of Adversaries”. *Proc. ACM SIGMETRICS 2008*.
- [3] B. Barak, **S. Goldberg**, D. Xiao. “Protocols and Lower Bounds for Failure Localization in the Internet”. *Proc. LACR EUROCRYPT 2008*.
- [4] **S. Goldberg**, J. Rexford. “Security Vulnerabilities and Solutions for Packet Sampling”. Invited Paper, *IEEE Sarnoff Symposium 2007*.
- [5] S. Kent, C. Lynn, and K. Seo. “Secure Border Gateway Protocol (S-BGP)”. *J. Selected Areas in Communications* 18(4):582-592, Apr. 2000.
- [6] IETF Working Group on Packet Sampling (PSAMP), <http://www.ietf.org/html.charters/psamp-charter.html>
- [7] D. Kempe, J. Kleinberg, E. Tardos. “Maximizing the Spread of Influence through a Social Network.” *Proc. 9th ACM SIGKDD*, 2003.