Towards Better Definitions and Measures of Internet Security (Position Paper)

J. Aspnes and J. Feigenbaum
Yale University
{aspnes,feigenbaum}@cs.yale.edu

M. Mitzenmacher and D. Parkes
Harvard University
{michaelm,parkes}@eecs.harvard.edu

January 30, 2003

1 Introduction

The conventional wisdom is that "the Internet is very insecure." The subtitle of this workshop, namely "deployment obstacles," implies that network owners, operators, and users could have solved pervasive security problems if they had deployed existing security technology. Is there solid evidence that either of these statements is true?

Clearly, there have been some well publicized Internet security problems (*e.g.*, viruses and distributed denial-of-service attacks) during the past five years, and some loss by individuals and businesses is attributable to them. Does this mean that Internet insecurity is really a significant problem? Is it a more serious problem than it was, say, ten years ago, or is there simply more awareness of it now than there was then? What fraction of Internet activity or potential activity is disrupted or prevented because of actual or perceived insecurity? Is this fraction higher or lower than it was ten years ago?

It is our thesis that better models, definitions, metrics, and data would greatly aid the development of *deployable*, effective security technology. We expand upon this thesis in Sections 2-4 below. Some highlights include:

- It might be fruitful to take a *systemic* approach, *i.e.*, to define, measure, and protect the security of the network as a whole, rather than the security of individual hosts and subnetworks.
- Similarly, it might be fruitful to take a *risk-management* approach, in which the goal is to predict, limit, contain, and correct the damage done by security failures, rather than the more traditional complexity-theoretic, logical, and algebraic approaches, in which the goal is to prevent failures or to isolate failed components.
- Asymptotic bounds on computational resources such as time, space, and bandwidth are not sufficient measures of the cost of a security solution; more realistic, comprehensive, and quantitative metrics are needed.
- Security researchers need better data and measurement techniques. In particular, there should be a way to obtain quantitative answers to the questions raised above about the scope of Internet insecurity. Furthermore, we must be able to determine whether, in general and for specific classes of users, things are getting better or worse and whether particular insecurity problems are solved, exacerbated, or unaffected by the deployment of particular solutions.
- Individuals and organizations must be able to compare the cost of deploying a particular security technology with the benefit of the additional security that it creates. Another way to say this is that they need to be able to compare the cost of deploying it with the cost of not deploying it.

2 Willingness to Pay for Universal Access

It is tempting to attribute the lack of adoption of known security technologies to ignorance or foolhardiness, but, in fact, there may be sound economic reasons to prefer a relatively insecure system to a competing secure one. The absence of accurate measures of security vulnerabilities argues for a humble attitude towards the adoption decisions made by individuals and businesses in the field. Although these decisions might not take into account all the costs of reduced security, we should start with the assumption that the people making them are responding rationally to their own assessment of the costs and benefits of adopting particular technologies.

Qualitatively speaking, one of the largest costs of many security technologies is that they stop you from dealing with people you don't know. While it is possible to build large systems that depend on pre-existing relationships between participants, the value of the most successful Internet technologies, including email and the World Wide Web, ultimately depends on their easy, universal accessibility. Users are willing to pay a high price in security for such universal access (*e.g.*, by receiving spam email or by accepting innocuous-looking web pages links might lead to bad places or, more likely, to nowhere at all), and, although there is clear demand for software that can reduce this price locally (*e.g.*, spam filters and web filters), there have been few serious proposals to solve the problem globally by replacing either the SMTP-based email system or the World Wide Web with closed systems usable only by those who have been certified as polite. This may be an inherent property of sufficiently large systems: As the number of participants grows, the strategy of excluding all potential bad participants becomes less and less cost-effective.

Given the choice between security and inclusion, many Internet-based businesses have opted for inclusion. For example, the HTTP protocol is notoriously vulnerable to snooping, and many Internet users are justifiably afraid to send sensitive information like credit-card numbers across unencrypted HTTP connections. Because of this fear, essentially all online retailers now use encrypted HTTPS connections for processing orders. However, many large Internet retailers (Amazon.com is perhaps the most successful example) are willing to process orders across unencrypted HTTP connections opened by users whose obsolete or misconfigured browsers prevent them from using HTTPS. Presumably, the cost of losing a sale is greater than the much more hypothetical cost of exposing such a user's credit-card number. (The cost of fraud is not borne directly by the retailer, but card-issuers who lose enough money have ways to pressure retailers to require encrypted connections.)

In contrast, security technologies that do not alter the relationships among participants in the system seem to be adopted much more easily. The relationship between the Telnet and SSH protocols is closely analogous to the relationship between HTTP and HTTPS; in both cases, the second is a drop-in secure replacement for the first. One difference is that Telnet was almost entirely used to allow remote users to connect to machines on which they already had accounts – a pre-existing relationship; thus, system administrators take much less risk in replacing Telnet by SSH of inadvertently excluding "new customers" who would respond by taking their business elsewhere. Combined with this lower cost was a higher risk of not using security; a server that accepts unencrypted Telnet connections risks the theft of passwords that can then be used to compromise the machine. So, in this case, we might reasonably guess that the difference between the continuing survival of unencrypted HTTP and the near-total disappearance of Telnet is explained by the difference between the costs and benefits of replacing them.

The preceding discussion is an attempt to account for the separate and often undocumented decisions of millions of system administrators, and so it is necessarily somewhat speculative. But it illustrates the point that the decision to adopt a particular security technology appears to be highly sensitive to cost. In measuring the effectiveness of particular security technologies, therefore, we should start with a baseline assumption that users have already taken into account the costs and benefits that they personally derive from them and that the more sophisticated users are already using all technologies that are personally cost-effective. What remains is the question of predicting security externalities – the costs to other users of a single user's decision to adopt a particular technology or not. Here we need good models, definitions, and metrics that can distinguish between failures that disrupt the activities of individual users and patterns of failures that threaten to bring down the entire system.

Economic analysis of security-technology adoption is currently an active area of study; see, for example, [AND1, VAR1] for further discussion.

3 Towards Systemic Notions of Security

Much of established cryptography and security research focuses on providing worst-case guarantees, typically in a stylized model of the network, and from the viewpoint of a single component of the network. Some beautiful results have been produced, *e.g.*, those on zero-knowledge proof systems and secure, multiparty function evaluation, but their usefulness in solving real-world network-security problems has yet to be demonstrated.

We believe that the security of computer networks is best viewed at the system level rather than at the component level, and that it is better to provide quantitative measures of security with respect to a realistic model of user behavior rather than absolute guarantees of security with respect to a stylized model of behavior. Useful security metrics should allow a comparison between the cost to deploy security measures and the benefit from system-wide security improvements.

A systemic, global view, of security is important for a number of reasons. First, networks often contain redundancy, which can make them robust to failures of individual nodes. In October 2002, there was a Distributed Denial of Service attack on the thirteen root DNS servers on the Internet that forced eight of the servers offline [LEW1]. However, the attack on the DNS system did not have a noticeable affect on the performance of the Internet because of redundancy, both because most DNS data is stored locally and because of the remaining five operational root servers. The Internet was designed to be a robust and distributed network, and any measure of the security of the Internet must account for the layers of redundancy within its design. Second, Internet protocols are adaptive (*e.g.*, one can often route around network elements that are under denial-of-service attack), and massive, creative effort can be mobilized in response to serious security breeches (*e.g.*, security patches or new anti-virus software can be installed quickly on massive numbers of compromised machines). Both of these properties were in evidence during the recent "Slammer Worm" attack on Microsoft's SQL Server [SLA1]. Third, computer networks may only be as secure as their weakest components (*e.g.*, the fact a master attacker can exploit vulnerabilities in a large number of slaves is what makes distributed denial-of-service attacks work).

Economics provides a natural framework within which to define metrics for systemic security. Consider network A and network $A+$, which is identical to A, except that additional security technologies have been adopted (possibly only in a subnetwork of A). From an economic point of view, it is natural to measure the value of this additional security in terms of the expected increase in the *total* utility of users of the network. At one extreme, it might be that A was already "secure enough," and the additional effort made by the adopters of the security technology produced no economic value. At another extreme, $A+$ may remain so insecure that the total utility of the system is negligible, and again the additional security has no economic value. Thus, a useful systemic metric of security should capture the aggregate opportunity cost to users of the system. Similar economic tools have been used by the networking community in its comparison of best-effort-only service and reservation-capable service (see, *e.g.*, [BS98]), and their analysis may carry over to our setting. Market-based methods have also been suggested as a method to expose the value of incremental security in a system [SCH1].

Given quantitative metrics, decisions about the deployment of new security technology can be addressed not only in terms of absolute security for individual components but also in terms of a systemic cost-benefit analysis and return on investment. For further discussion, see Anderson's argument that, because there are many possible attacks, it can be better to develop methods to contain an attack and construct a patch than it is to anticipate and fix all attacks before they occur [AND1].

An economic approach to the development of metrics for network security also suggests adopting a *rational*, but *non-adversarial*, model of attack. We advocate an *algorithmic mechanism design* approach to security, in which security measures are designed to maximize the value of incremental security as measured with respect to an economically-rational model of attack. One should identify points in the network where security fixes are most cost effective and aim to position incentives so that optimal system-wide security decisions are made by the individual actors within a network. For an introduction to algorithmic mechanism design and its application to network protocols, see, *e.g.*, [FS02, NR01].

In addition to improved metrics for security, security researchers need to be able to determine whether security is getting better or worse by measuring actual network events and activities. This will require application-specific models to capture users' utility through measurable network properties, such as dropped packets and server down time, as well as network-level models to account for the ability of specific classes of users to switch to alternative hosts and subnetworks and route around problems.

4 Example: Virus Propagation and Immunization

A natural problem domain in which to study security externalities is computer viruses. To determine reasonable and cost-effective solutions to the virus problem, it is important to clarify the underlying goal. From the point of view of the individual user, the goal is obviously to avoid becoming victim of a computer virus. At the network level, however, it seems clear that complete prevention of all viruses is an unrealistic goal. By analogy with physical viruses, a more reasonable goal appears to be preventing a widespread outbreak that cripples the whole network. Even this goal, however, is rather ill-defined and must be formalized in future work.

There is preliminary work on applying epidemiological models to the computer virus problem [KCW1, KW1]. In the most widely studied models, nodes are either healthy or infected; neighbors of infected nodes can become infected over time according to some fixed rate. Nodes are also cured at some rate, but they can become re-infected. More recent work has examined these models in the context of Internet-like graphs (sometimes referred to as *scale-free graphs* or *power-law graphs*) [PV1, PV2, ZTG1]. This recent work has demonstrated that even these simple models are hard to analyze fully in the context of the Internet. In many other network models (*e.g.*, random graphs or lattices), the virus will die out quickly if the ratio of the spreading rate to the cure rate is sufficiently low; in Internet-like graphs, however, viruses tend to survive in some small fraction of the population, posing the risk of returning in a subsequent epidemic burst. Early stages of propagation are exacerbated by the highly connected nodes of the Internet, which have the potential to impact many neighbors quickly. As one might expect, however, immunizing these high degree nodes can greatly reduce the effect of an outbreak.

One possible obstacle to effectively combatting viruses is that the benefit obtained by using anti-virus software does not accrue only to the user of the software but rather to all users of the network. Hence, those that employ anti-virus software incur costs (*e.g.*, the cost of the software itself) but may not be adequately reimbursed for the benefits they provide to the network. As an extreme example, Williamson has demonstrated that virus outbreaks can be significantly curtailed if machines are limited in the number of connections that can be made over short periods of time [W1]. A user who employs software to curtail connections does obtain some benefits; for example, in the case that a virus that attempts to spread rapidly is placed on the machine, it may be detected and removed more rapidly. However, there are some associated costs, including the cost of the software and the restriction on connections in the course of normal machine use. The main benefits of using such software accrues to the network as a whole, because such software limits the rate at which viruses can spread.

How to provide incentives, particularly for highly connected nodes, to use anti-virus software offers a novel direction for research in distributed algorithmic mechanism design.

References

[AND1] R. Anderson, “Why Information Security is Hard – An Economic Perspective,” http://www.cl.cam.ac.uk/users/rja14/econ.pdf

[AND2] R. Brady, R. Anderson, and R. Ball, “Murphy’s Law, the Fitness of Evolving Species, and the Limits of Software reliability,” http://citeseer.nj.nec.com/brady99murphys.html

[BS98] L. Breslau and S. Shenker, “Best-Effort versus Reservation: A Simple Comparative Analysis,” in *Proceedings of the 1998 ACM Sigcomm Conference*, pages 3–16.

[FS02] J. Feigenbaum and S. Shenker, “Distributed Algorithmic Mechanism Design: Recent Results and Future Directions,” in *Proceedings of the 2002 ACM Workshop on Discrete Algorithms and Methods in Mobile Computing and Communications*, pages 1–13.

[KCW1] J. O. Kephart, D. M. Chess, and S. R. White, “Computers and Epidemiology,” *IEEE Spectrum*, May 1993.

[KW1] J. O. Kephart and S. R. White, “Directed-graph Epidemiological Models of Computer Viruses,” in *Proceedings of the 1991 IEEE Computer Society Symposium on Reserach in Security and Privacy*, pages 343–359.

[LEW1] J. Lewis, “Assessing the Risks of Cyber Terrorism, Cyber War and Other Cyber Threats,” http://www.csis.org/tech/0211_lewis.pdf

[NR01] N. Nisan and A. Ronen, “Algorithmic Mechanism Design,” *Games and Economic Behavior*, 2001.

[PV1] R. Pastor-Satorras and A. Vespignani, “Epidemic Spreading in Scale-Free Networks,” http://complex.upc.es/ romu/Papers/virus.pdf

[PV2] R. Pastor-Satorras and A. Vespignani, “Immunization of complex networks,” http://complex.upc.es/ romu/Papers/immuno.pdf

[SCH1] S. Schechter, “Quantitatively Differentiating System Security,” http://www.sims.berkeley.edu/resources/affiliates/workshops/econsecurity/econws/31.pdf

[SLA1] http://www.microsoft.com/security/slammer.asp

[VAR1] H. Varian, “Managing Online Security Risks,” Economic Science Column, *The New York Times*, June 1, 2000, http://www.nytimes.com/library/financial/columns/060100econ-scene.html

[VAR2] H. Varian, “System reliability and Free riding,” http://www.sims.berkeley.edu/resources/affiliates/workshops/econsecurity/econws/49.pdf

[W1] M. Williamson, “Throttling Viruses: Restricting propagation to defeat malicious mobile code,” http://www.hpl.hp.com/techreports/2002/HPL-2002-172.pdf

[ZTG1] C. Zou, D. Towsely, and W. Gong, “Email Virus Propagation Modeling and Analysis,” http://tennis.ecs.umass.edu/~czou/research/emailvirus.pdf