From Art to Science: A Vision for the Future of Information Assurance

Susan C. Lee and Donna M. Gregg

As networked information systems become more essential to modern life, the need for information *assurance* (IA)—securing availability, integrity, and confidentiality for our information—becomes increasingly urgent. Capability for IA lags networking capability significantly. While robust science and engineering disciplines underpin the construction of faster, more capable networks, no such foundation exists for concurrently assuring them. This capability gap is a particular problem for APL's customers who depend on their information networks to provide our national security in the face of well-resourced, highly motivated adversaries. To advance the state of the art in technologies for IA, the Laboratory envisions a better understanding of the science that governs networking and assurance, leading to new technology approaches and a rigorous engineering discipline for IA. APL is making contributions to IA technology in the area of secure platforms and is exploring the application of formal methods to the science of IA. The cornerstone of the Laboratory's contribution to achieving the envisioned future, however, will be in IA engineering, especially by determining meaningful IA metrics and the techniques for measuring them.

INTRODUCTION

An explosion of information technology (IT) innovation and investment in the early 1980s has resulted in a society that is highly dependent on computer systems (Fig. 1). The ubiquitous, networked computers in homes and businesses are only a fraction of all the computers that we interact with daily. Computers are embedded in the systems we rely on, from common home appliances to the most vital portions of our critical energy, finance, transportation, and telecommunications infrastructure. Our national security also depends on computers and networks for intelligence, command and control, and weapon guidance. While increased connectivity/access to information has great advantages for both business and military operations, it also offers a tremendous opportunity for societal disruption and information warfare. Such a complex network of interacting computers is susceptible to attacks and breeches from any number of nation states, terrorists, and criminals. The goal of information assurance (IA) is to provide the availability, integrity, and confidentiality of the information resources

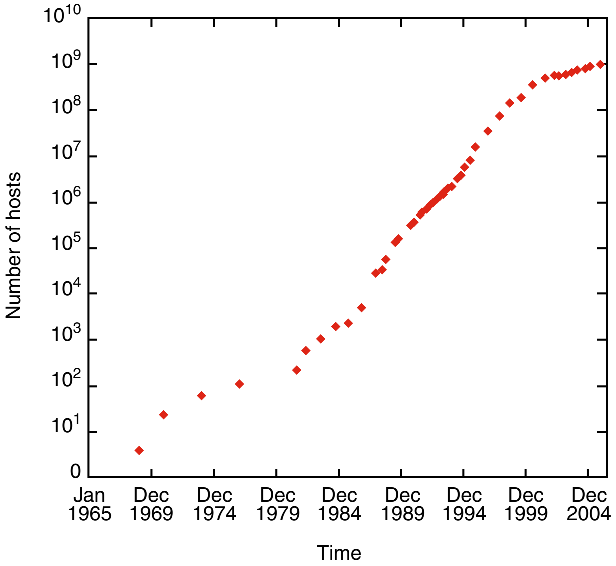

Figure 1. Growth in the number of networked hosts over the past two decades.

needed to reap the benefits of a networked environment that is imperiled by a range of destructive forces.

Despite the importance it assumes in hindsight, IA was not a consideration for the early leaders of the information revolution. The simplifying assumption of the "benign user" allowed them to design software and protocols that only needed to be reliable in the face of inadvertent, random failures. For this reason, the first actors who sought to manipulate networks for malicious purposes found fertile ground. These so-called hackers were initially considered rogue geniuses.[1] As an Information Age lifestyle—online banking and shopping, telecommuting, instant messaging, e-mail, web surfing—became routine, hacker attacks became increasingly frustrating and costly.

IT vendors responded to the hacker threat with research and product lines for IA. Unfortunately, as the daily barrage of spam and frequent denial-of-service (DoS) attacks demonstrate, the IA industry has not prevented malicious individuals from conducting damaging attacks with minimal fear of reprisal. There are two basic reasons for this. First, the fundamental lack of security in the original design of networked systems makes IA, at best, an awkward and cumbersome retrofit. Sometimes, assurance cannot be achieved without inhibiting the very functionality that is desired. Second, the complexity of the systems involved—both the inherent complexity of huge software programs and the emergent complexity of large aggregations of networked systems—defies our ability to prevent the types of flaws and unexpected interactions that are fodder for hacker attacks. Fortunately for the commercial and home user, the hacker threat has not risen to a level that makes the Information Age lifestyle untenable. In an evolutionary metaphor, the ratio of predator (hacker) to prey (legitimate user) is such that individual cyber incidents do not affect the overall viability of the "herd."

For the Laboratory's national security customers, however, the threat is much greater. First, the nation-state adversaries they face are far better resourced and much more highly motivated than the hacker seeking thrills or monetary gain. Second, the stakes are considerably higher. The same DoS attacks that prevent the home users from doing their online shopping for a day could paralyze the military on the first critical day of an operation like Iraqi Freedom. The same confidentiality breach that causes Discover or Visa to issue new credit cards to a dozen customers could alert a dozen terrorists to their imminent capture. Home and commercial users measure their losses in time and money and need only ensure that their gains exceed their losses. Our national security customers measure their losses in lives and, potentially, national sovereignty, and must be able to depend on the immediate and correct operation of their information systems.

Initially, these security concerns prevented the national security community from using commercial IT in their operational systems. The mass market for IT products, however, drives an enormous commercial investment that cannot be matched by government spending. Today, security concerns have been overwhelmed by budgetary constraints and the desire to increase the efficiency and effectiveness of our national security systems by leveraging the vast resources of the commercial sector. Commercial hardware and software systems are now used by the national security community nearly everywhere they meet a need. Unfortunately, there is no mass market for the extremely high level of IA required by national security users. Commercial users are unwilling to pay more for security than is needed to ensure their profits or to sacrifice the interoperability and ease of use that today's state-of-the-art IA demands. Unless the national security community advances IA well beyond the commercial state of the art, it is left with the option of accepting a serious risk or losing a critical information technology.

APL is responding to the IA challenges of its customers, just as it has responded to other national security challenges in the past. To meet *this* challenge, however, the Laboratory must go beyond applied engineering: APL must contribute to establishing the fundamental science and engineering basis for IA to meet the high assurance needs of our customers.

STATE-OF-THE-ART INFORMATION ASSURANCE SCIENCE AND TECHNOLOGY

Like any technical field, IA comprises *technologies* (i.e., hardware/software systems) that perform needed functions, an *engineering discipline* that guides the development of systems, and *fundamental science* that drives both the technology and engineering disciplines. In IA, each has serious shortfalls.

Technology

The technology component of IA is the best developed. Numerous IA devices have been made for attack protection, detection, and mitigation. In an analog to the physical protection of the past, modern IA uses various means of identification and authorization (I&A) to logically bar or admit users to the network. The well-known user ID/password is a form of I&A; however, much more potent technologies, such as public key infrastructure (PKI), which leverages cryptography, and biometrics, which uses physical characteristics of the user for identification, are gaining currency. Firewalls are another form of logical access control; they bar or admit packets, the basic unit of information exchange on a network, based on their characteristics. Scanners have been developed to find and in some cases even correct vulnerabilities, i.e., errors in or unintended consequences of the networked systems, before they allow attacks to occur. Intrusion detection systems (IDS) and virus checkers are among the better-known attack detection devices. Some of these devices have a limited ability to respond to a detected attack by "locking out" suspect users or dropping suspect packets.

The available IA technologies have certain limitations that reduce their effectiveness. First, they are all reactive. For example, vulnerability scanners, IDS, and virus checkers can only work on known attacks. In the usual evolutionary spiral of predator and prey, attackers continually devise new ways of defeating the IA. In this, they are aided by the rapid influx of new IT—new protocols, services, operating systems (OS), and application upgrades—that comes replete with new vulnerabilities ripe for exploitation. Although IA vendors have honed their processes so that new virus signatures can sometimes be distributed hours after the onset of a new virus, users are still completely defenseless against novel attacks.

Second, IA technologies focus on protecting the network, rather than on the purpose the network serves. For example, a firewall can entirely prevent passage of packets destined for a particular application, e.g., a web service, and thus can prevent any web-based attacks from affecting the network it protects. If the network is used to provide web services, however, this so-called protection renders it useless. The reaction capability provided by state-of-the-art IA nearly always results in blocking some connectivity; if the mission the network supports requires that connectivity, then the IA itself is destructive to the mission. IA appliances can actually be co-opted by attackers to inflict a DoS on their victims by tricking them into a reaction that impedes legitimate operations, even though there is no real attack against the system.[2]

Finally, the implementation of existing IA technology uses the same vulnerable computers, software, and networking that it is attempting to protect. For instance, firewalls and IDS are simply computer workstations and are themselves subject to attack. Even when the technology itself is robust, the administration of the system using the technology is often vulnerable. For example, PKI uses a particularly safe and elegant cryptographic scheme to identify legitimate users. The cryptographic algorithm, however, is only a small part of a large network infrastructure needed for issuing, revoking, escrowing, and storing keys as well as making identifications. Although the cryptography is very difficult to attack, the surrounding infrastructure is not.[3]

Engineering

While IA technology is abundant, if limited, engineering tools for IA are both few and seriously deficient. Today, IA engineering is merely a prescriptive process: a set of steps to follow and issues to consider. A major barrier to a more rigorous approach to IA engineering is the lack of useful metrics. There is no agreement on what security is, what "units" it has, or how it can be measured. The metric approaches that have been suggested or are in use are of little value in typical engineering exercises such as requirements specification or design.

Performance metrics for IA appliances, such as the throughput of a firewall or the number of rules it implements, give no indication of how likely they are to protect against an attack. Combinatorial rules for IA devices are also unknown; for example, does overall system security increase if an IDS is added to a firewall? Although this sounds reasonable, APL has shown that the combination of certain security devices can enable an attack that would be impossible with either single device, and thus actually lowers security.[4] Expert opinion is sometimes used to create a relative ranking among risks, vulnerabilities, and mitigation techniques. If these are expressed numerically, they can be used in standard utility functions to score various IA choices; however, having no basis in fundamental theory, the results will change depending on the experts who are consulted.

In an analogy to reliability engineering, probabilistic terms such as "probability of a successful root-level break-in" are sometimes used to levy system-level security requirements. Reliability measures are derived through controlled testing of statistically significant sample populations. Red Team testing of operational systems is envisioned as the analogous process to measure assurance-related probabilities, but obtaining a significant result is problematic. Nearly 4000 Red Team attacks would be needed to establish that the probability of resisting an attack is 0.99 ± 0.01 at a 90% confidence level. An operational network is unlikely to remain static over the period of time needed to carry out these attacks, rendering the earlier results incomparable to the later results. *No* result could preclude the possibility that *every* computer in the network was vulnerable to some new exploit unknown to the Red Team. Finally, even if such statistics could be measured accurately, they give no insight into the mission impact of the 1-in-100 successful attacks.

Without metrics, security requirements cannot be quantified; security cannot be designed to requirements or tested to determine if it meets requirements. Without a means to perform these hallmark engineering activities, IA cannot be considered a true engineering discipline.

Science

Although at least an informal engineering process for IA exists, there is no real "science" of IA. Cryptography and formal methods (FMs) are two areas of true scientific endeavor that are generally accepted as *relevant* to IA. IA technologies that rely on cryptographic algorithms as their basis (and there are many) can provide quantitative, provable security against cryptographic attacks. For example, the likelihood that an encrypted message can be decrypted in a defined period of time can be precisely calculated as a function of the type of cryptographic algorithm, the length of the key, the length of the message, etc. Unfortunately, as noted above, there are many attacks against the system that *surrounds* the cryptographic algorithm, and currently no scientific principles let us calculate if the key can be stolen through a network attack or if the cryptographic algorithm can be bypassed altogether because of OS insecurity.

The field of FMs holds out some hope for providing those principles. FMs allow systems to be defined mathematically so that properties of the system can be formally derived and proven. Many properties of systems related to IA, such as correctness, can be proven for small systems (i.e., software systems with <10,000 lines of code). The major shortfall of FMs is scalability; with a Windows OS at ≈40 million lines of code, correctness cannot be proven with a FMs approach. To extend FMs to larger systems, an understanding of composition rules must be greatly expanded; that is, given the formally defined properties of two small systems, what are the properties of their combination?

Cryptography provides scientific underpinnings only for an IA *tool*, not for security itself. Because of scalability issues, FMs have little to offer in terms of real networks. To overcome the limitations of state-of-the-art technology and to give IA engineering a quantitative basis, the underlying scientific principles impacting IA must be determined.

APL'S FUTURE VISION FOR IA

APL envisions a future in which our customers can confidently leverage the cutting edge in commercial IT to build national security systems that are equal or superior to any that a potential adversary can bring to bear. This confidence must rest on scientific and engineering principles for IA, analogous to the principles that let us employ tanks and bombs with realistic expectations for their efficacy in a given situation.

Technology Vision

Unlike today, future IA design will *not* begin with the placement of individual, specific IA technologies within a network designed to operate only under benign conditions. Instead, it will begin by designing a network that will operate under the hostile conditions that are likely to be found. In much the same way that certain terrain is more defensible than others (e.g., the value of having the "high ground"), some information network designs will prove to be more defensible than others. APL has already shown that by adjusting parameters, networks can be designed to be resistant or impervious to certain types of attacks. For example, increasing the release rate of the SYN packet queue makes the classic "SYN flood" DoS attack much less effective (Fig. 2).[5] Manipulating queue sizes and data refresh rates can also reduce the effectiveness of whole classes of attack. When availability of connectivity is essential, network designs with multiple, redundant paths for communication make a DoS attack less viable. Other potential techniques are the introduction of technology diversity (e.g., different

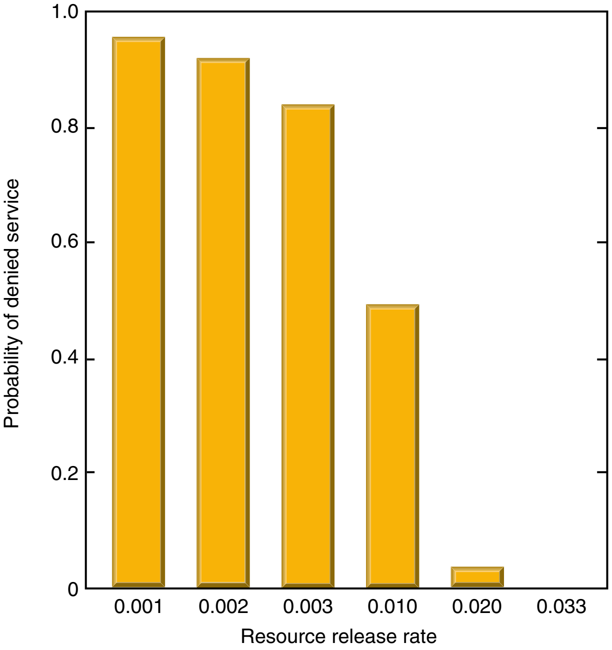

Figure 2. Probability of denied service for TCP connections as a function of resource release rate.

OS, applications, and protocols) to limit the number of systems that can be compromised by a single attack type, out-of-band control for certain essential components, and distribution of control or data to limit single-point failures.

Within a defensible network design, there will be appropriate points at which IA-specific technology can be arrayed to enhance the network's natural resistance to attack, or to fortify unavoidable weak points. Especially in the future when the network design itself will provide assurance, certain non-IA–specific elements of the network (e.g., the network management systems) will also provide key IA functionality. In the end, the security of the whole network rests on the security of the computers that supply the IA functionality. The Laboratory's vision for the future of IA includes the ability to harden individual computers using three principles: attack-independent attack detection, layered monitoring, and hardware support.

The reactive nature of today's attack detection technology (e.g., IDS, virus checkers) is a result of using the characteristics of the *malicious attacks* to differentiate them from normal, desirable system behavior. An attack must be known before its characteristic signature can be determined. To achieve attack-independent results, IA technology must focus on the known, predictable aspects of the *desirable behavior* and use that to detect any attack—previously experienced or novel—that deviates. For example, APL devised a monitor that simply looked for deviations from the specified protocol for establishing a TCP connection and showed that the monitor could detect several common attacks with quite different signatures.[6]

The major drawback of this approach is the difficulty of specifying the correct behavior of complex systems. APL envisions a layered approach to monitoring, where correct behavior is characterized iteratively, building a hierarchy of more and more complicated, but still assured, functionality. An OS, at an estimated 40 million lines of code for Windows XP, is at the top of the hierarchy. Host-based intrusion detection systems (even some called "wrappers" that embody the APL vision of monitoring correct behavior) and virus checkers attempt to prevent network-based compromise of the OS. Since they depend on the OS to provide some of their needed functionality, they can be circumvented without monitoring at a lower layer.

Virtual machine technology, where the base computer hardware is emulated, allows the placement of independent monitoring below the COTS OS. While the COTS OS runs on the virtual machine, a simpler, more secure OS runs on the real computer hardware, ensuring that the higher-layer IA programs are neither tampered with nor bypassed. The secure OS can enforce a policy for correct operation that prevents corruption of the COTS OS, or at the very least, provides a fail-safe functionality. Since this hidden OS does not need to provide all the features of the COTS OS, it can be much simpler and more amenable to validation prior to use. This greater simplicity also allows more complete monitoring against compromise by an even lower layer of intrusion detection.

While a layered approach may allow security to rest on smaller, more focused software that is easier to validate, ultimately all software on a given machine responds in some way to external input, and as long as the possibility for error exists, it remains vulnerable to attack over the network. Ultimately, to make IA mechanisms tamper-proof and non-bypassable, specific hardware support—isolated processing and memory, hardware invocation, and out-of-band communication—is needed. Conceptually, this support is realizable; practically, any change or addition to COTS hardware can become quite costly. Even a small cost per unit becomes large when multiplied by the vast number of commercial systems that are deployed. Further, the rapid obsolescence of commercial systems implies an ongoing and equally rapid development of new COTS IA technology. APL envisions two approaches to enabling practical, hardened systems. First, we have shown that the hardware support already embedded in today's COTS systems (e.g., the underutilized ring architecture) can be leveraged for robust security, and we are exploring the ramifications of the newer generation of computer hardware for IA. Second, when COTS support is insufficient, APL seeks to apply FMs to building an IA hierarchy, described above, so that only a minimum of costly, specialized IA hardware is needed.

Engineering Vision

As with other engineering endeavors, assurance engineering will begin with specifying performance requirements. Rather than network-oriented parameters like "probability of root break-in," APL envisions assurance requirements that are couched in mission-oriented terms that are immediately useful in determining a system's fitness for use. To illustrate, consider the Global Information Grid (GIG) backbone, which has one basic function: to deliver packets from one edge network to another. For the GIG backbone, some appropriate assurance requirements could be to

- Achieve the minimum allowable backbone bandwidth and maximum delay given a distributed DoS (DDoS) attack consisting of a minimum of 10,000 packets per second entering the core from at least 10% of the edge networks
- Restore expected bandwidth and delay within 5 min under the condition that at least 30% of the core routers are compromised and nonresponsive to routing updates

Requirements stated in this fashion have several advantages. First, they are measurable; unlike establishing a

0.0001 probability of root break-in, it is straightforward to plan a test to measure compliance with these requirements (although such a test presents logistical issues!). Second, with such metrics, the cost of assurance can be compared to its benefit *in terms the users can understand*. For example, as less costly means of assurance are contemplated, the restoration time stated in the second hypothetical requirement above might grow to 10 min or 1 hour, an easily understood impact. Finally, these metrics are attack independent. It is unnecessary to know how the DDoS attack or router compromise was achieved; regardless of the attackers' methods, a GIG backbone that meets these requirements will satisfy its assurance requirements.

To analyze the assurance requirements, future IA engineers will perform rigorous, quantitative risk assessments. Based on the well-established premises of fail-safe design and fault analysis, IA engineers will build attack trees (Fig. 3) identifying the steps that attackers can take to achieve a given effect. The effects considered in the risk assessment are those that would result in a degraded mission (e.g., reduction of bandwidth and increase in delay in the examples above). The risk analysis will determine the probability that these effects can be achieved by particular means (e.g., the DDoS attack or widespread router compromise). Given cost or other constraints, IA engineers will be tasked with preventing the adverse effects for all attack "paths" with a probability higher than 50%, for example. As with the future metrics, the attack trees will be specific-attack independent. This seeming paradox is resolved by specifying attack steps in terms of their result (e.g., router compromise) rather than the specific method used to attain the result.

APL has already led GIG risk assessments using attack trees. Even today, this method is useful for finding the common steps that enable realization of many different attacker goals, or that reappear in every path leading to an important goal. It is less useful in determining the quantitative risk because the probability that any given attack step will be attempted and successful is unknown. Today, expert opinion is used to estimate the probabilities; in the future, APL envisions that they will be based on data.

As noted above, the notions of controlled testing of a statistically significant sample to obtain probabilities are problematic when applied to the realm of IA. If they can be obtained at all, these probabilities must be derived from statistical analysis of live networks, rather than the controlled testing used in reliability analysis. In the future, APL envisions a coordinated effort that collects consistent information from all incidents observed in large networks—if not for the Internet, then at least for all DoD networks. Long-term, detailed analysis of vulnerability characteristics and introduction rate, malicious code, incident characteristics and impacts, attacker behavior, and adversary intentions will lead to predictive probabilities for use in risk analysis.

As in other engineering disciplines, the best architecture and design to meet a set of assurance requirements may not be obvious initially. Modeling and simulation (M&S) is an indispensable tool for analysis of alternatives in many engineering disciplines; the Laboratory envisions M&S of attacks and defenses as an essential tool in the future of IA engineering as well. The size and complexity of networked information systems; the continual revision of the hardware, software, and topology of the systems; and the unpredictable nature of both network traffic and vulnerability discovery make stochastic models suitable for IA engineering. The same statistical data needed to support risk assessment will be needed to perform realistic M&S, along with other statistics, especially those needed to support traffic modeling. Again, in the future, an ongoing activity will be established to collect and publish the needed measurements of live networks.

To make M&S a robust IA engineering tool, two major obstacles must be overcome. First, models must be scalable to large networks, yet retain sufficient fidelity to model the often small-scale features of network attacks. Second, the fidelity of the models must be subject to rigorous validation. Highly accurate, detailed M&S of small networks (i.e., fewer than 500 nodes) can be performed today; data needed to validate them can be collected from live networks or test beds with a modest instrumentation suite. Scaling detailed models to Internet size quickly exceeds the capacity of ordinary computers. Distributed computing can and is being used to attempt to allow this scaling, but

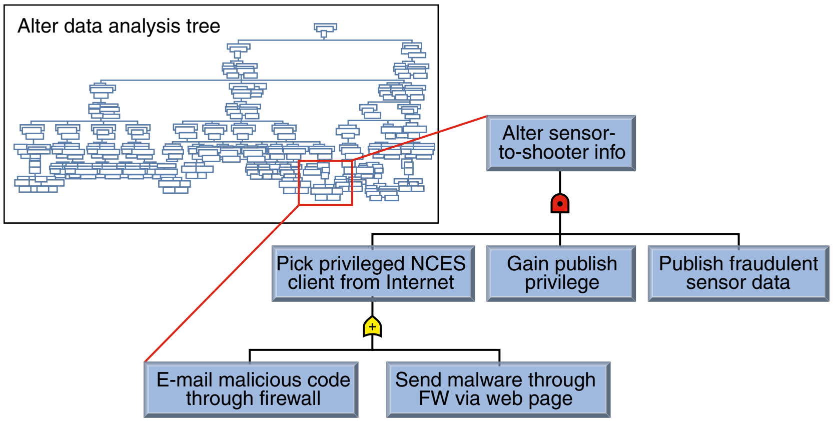
Alter data analysis tree
Alter sensor-to-shooter info
Pick privileged NCES client from Internet
Gain publish privilege
Publish fraudulent sensor data
+
E-mail malicious code through firewall
Send malware through FW via web page

Figure 3. Illustration of a partial attack tree (FW - firewall, NCES = Network-Centric Enterprise System).

very large networks are still beyond reach. In addition, there is insufficient instrumentation to capture validation data. The alternative most often used today is to build coarse models of large networks based on matching the statistics of some short interval of observed data. The accuracy of such models cannot be validated over time or under other network conditions.

APL envisions the future IA engineer using models that have the underpinning of science to guarantee their validity. Much as the science of thermodynamics allows accurate thermal models to be built without modeling the behavior of each and every atom, future network M&S will be built on a foundational understanding of the properties of large networks. The Laboratory envisions advances in science in specific areas to provide this foundation.

Science Vision

In APL's vision for IA, science will provide the underlying theory that can be used to develop both IA technologies with broad impact and engineering tools with true predictive power. This theory will illuminate the properties of assured systems and the fundamental limitations of assurance for real systems.

APL envisions an understanding of system composition as the key to understanding IA. The major factor contributing to the insecurity of information systems today is their overwhelming complexity. These huge systems cannot be proven to be correct, nor can the totality of their behavior be predicted in advance. The most complex system, however, is composed of individual, simple parts. Individually, the parts can be understood, or even formally proven to be correct. Even when the individual parts are correct, however, their interactions may have unexpected results, just as the combination of the firewall and IDS mentioned earlier was shown to have a vulnerability that neither had alone. The ability to derive the behavior of ensembles (combinations of multiple systems), from first principles rather than tests, will enable many parts of APL's IA vision. It will be needed to formally prove correctness of large pieces of software, to compose the hierarchy of host-based IA technologies described in the Engineering Vision, and to create validated, large-scale network models out of individually testable subnet models.

Beyond the augmentation of FM theory and practice to allow formal derivation of the properties of composite systems, the size and observed behavior of today's information systems suggest that the study of complexity (e.g., fractals, chaos, emergent behavior) will be an important contributor to the science underlying IA. Better understanding of the characteristics of complex systems could reveal hard limits on the security that can be expected for networked systems or lead to insights into the design of large, complex pieces of software (e.g., OS) without the unexpected interactions that create vulnerabilities. Most probably, it will provide the basis for accurate large-scale network models that operate faster than real time and can be used in the design of assured networks.

CONTRIBUTING TO THE VISION

The vision presented above has grown out of APL's work in IA over the past 6 years. While the Laboratory's vision for the future of IA is far from a reality, it is one that seems achievable with sufficient effort. APL is contributing to advances in IA technology and engineering and seeks to contribute to the fundamental science of IA.

Technology Contributions

APL has been in the vanguard of the attack-independent intrusion detection technology described above. We have experimented with this approach in creating an IDS for the Army and Navy, and to provide a fail-safe capability to a National Security Agency (NSA)-developed, virtual-machine–based OS.[7] The latter capability took advantage of a standard INTEL processor hardware feature called System Management Mode to make the fail-safe software tamperproof and non-bypassable by a remote user. Currently, APL, in conjunction with Mitre and industry partners, is providing NSA the same type of attack-independent monitor using more advanced hardware support.[8]

Engineering Contributions

Advancing IA as an engineering discipline is the cornerstone of APL's contribution to its future vision for IA. The Laboratory seeks to develop useful IA metrics, to demonstrate the means for measuring them, and to make IA M&S a reliable tool for predicting IA performance.

Taking its long history of weapons development into the information realm, APL has attempted to derive measures of effectiveness (MOEs) for the information defenses it develops from the start. Over a period of about 4 years, APL studied MOEs for DoS attacks, first with internal funds and then for the Defense Advanced Research Projects Agency. To carry out these studies, we were required to both define the metrics and measure them. In all our internal and external IA tasking, APL strives to deliver both the required technology and the rational set of associated MOEs.

The Laboratory has been using M&S of network attacks and defenses to demonstrate its approach to metrics. This activity has led to several years of study on the problem of scalable, high-fidelity models. We are developing an M&S framework that allows information system models of different types and fidelity to be tied together, from actual implementations of IA technologies (highest possible fidelity) to packet-level simulations to abstract mathematical models of large network segments. We continue to investigate new modeling paradigms such as statistical mechanics approaches to representing networks. These new approaches hold out the hope of combining packet flows produced by empirically validated small-network models into models of large networks with mathematically provable validity.

Obtaining accurate probabilities of real events is essential to building any probabilistic model. Understanding the probability of attack, along with attack success given a set of initial conditions, is crucial to quantifying the threat element of APL's IA metrics concept. Large networks like the Internet provide a large "sample" population and typical product mix. Computer Emergency Response Team (CERT) databases record the successes of hackers who are seeking vulnerabilities and exploiting them. Although incomplete and difficult to mine, APL is developing the tools and techniques to use CERT data to derive realistic probabilities to drive our models. For example, our initial study showed that the rate of vulnerability discovery does vary from product to product, but does not decrease significantly, even after years of widespread use (Fig. 4).[9]

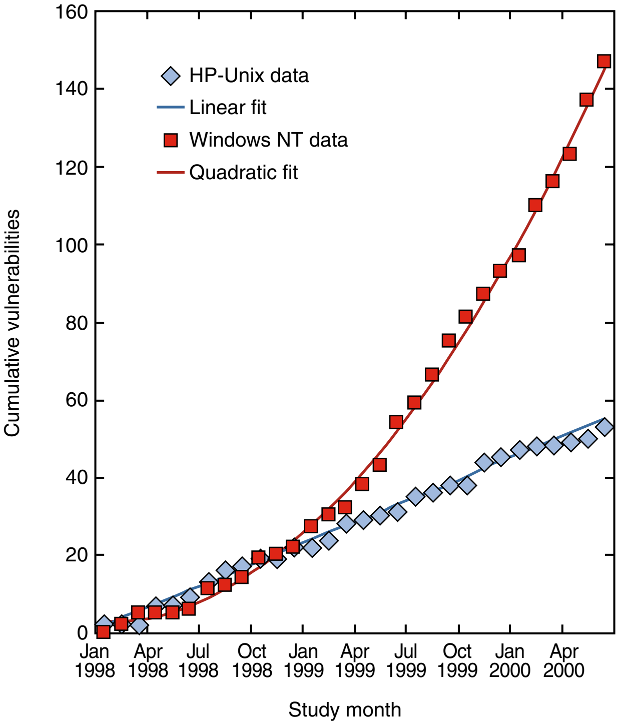

Figure 4. Cumulative number of vulnerabilities found in Windows NT and HP Unix over 2 years.

Science Contributions

As APL capability and experience in IA grow, we recognize increasingly that the lack of scientific underpinnings seriously hampers technology and engineering efforts. To address this lack, deep expertise will be required in key areas such as FM and complexity theory, which may have little application at APL outside our focus on IA engineering. APL plans internal investment to explore this space and is embarking on pilot studies of FM applications to metrics definition and measurement, partnering with the University of Texas to supply additional FM expertise.

CONCLUSION

The importance of IA to APL's customers and the difficulty of IA challenges make it a natural fit with the Laboratory's mission and image. Although IA is a relatively new domain for us, significant progress has been made in establishing a presence in the community, a body of experience and technology to draw on, and a vision to shape our contributions to its future. APL is confident that by pursuing its vision, in concert with its sponsors and through partnerships with leaders in the APL community, we will make critical contributions to the critical challenges facing our national security customers in IA.

REFERENCES

[1] Denning, D. E., "Concerning Hackers Who Break into Computer Systems," in *Proc. 13th Nat. Computer Security Conf.*, Washington, DC, pp. 653–664 (Oct 1990).

[2] Kewley, D. L., and Lowry, J., "Observations on the Effects of Defense in Depth on Adversary Behavior in Cyber Warfare," in *Proc. 2001 IEEE Workshop on Information Assurance and Security*, U.S. Military Academy, West Point, NY (5–6 Jun 2001).

[3] Cohen, F., "50 Ways to Defeat Your PKI and Other Cryptosystems," in *Proc. 23rd Nat. Information Systems Security Conf.*, Baltimore, MD (27 Jun 2000).

[4] Blackert, W. J., Castner, A. K., Gregg, D. M., Hom, R. L., Jokerst, R. M., and Kyle, E. M., *Distributed Denial of Service Defense/Attack Tradeoff Analysis*, VS-03-073, JHU/APL, Laurel, MD (Sep 2003).

[5] Gregg, D. M., Blackert, W. J., Furnage, D. C., and Heinbuch, D. J., *Denial of Service Attack Assessment Report*, VS-01-071, JHU/APL, Laurel, MD (Jul 2001).

[6] Lee, S. C., and Heinbuch, D. V., "Building a True Anomaly Detector for Intrusion Detection," in *Proc. IEEE Military Communications Conf. (MILCOM2000)*, available on CD, Los Angeles, CA (Oct 2000).

[7] *N-FORCE Daemon Prototype Technical Description*, VS-03-021, JHU/APL, Laurel, MD (Jul 2003).

[8] *N-FORCE Attestation Study Contract Summary Report—Final*, VS-05-059, JHU/APL, Laurel, MD (Sep 2005).

[9] Lee, S. C., and Davis, L. B., "Quantitative Security Data From Statistical Analysis of Vulnerability and Incident Reports," in *Proc. Third Annual Int. Systems Security Engineering Assoc. Conf.*, available on CD, Orlando, FL (13–14 Mar 2002).

THE AUTHORS

Susan C. Lee is an APL Principal Professional Staff member and serves as Chief Scientist for both the Laboratory's Applied Information Sciences Department (AISD) and the Infocentric Operations (IO) Business Area. As Chief Scientist, Ms. Lee is responsible for defining and guiding the IO technology program, including concept development, internal research and development, new business development, and sponsored research. Since 1999 her focus has been in the information assurance (IA) area where her activities have included the invention of a novel intrusion detection system and working with the National Security Agency to define the IA architecture for DoD's Global Information Grid. **Donna M. Gregg** is Head of the Information Systems Branch in the AISD and serves as the focus area lead for IA in the IO Business Area. As the IA focus area lead, Ms. Gregg is responsible for establishing strategy and guiding business area investments in IA research, most recently focusing those investments on establishing a secure computer platform and techniques for transforming IA into an engineering discipline. Ms. Gregg has worked in IA since 1998, at first using modeling and simulation to quantify the security of an information system. She was the Supervisor of the Laboratory's first technical group focused on IA. The team can be contacted through Ms. Lee at susan.lee@jhuapl.edu.

Susan C. Lee

Donna M. Gregg