Information Warfare: A Brief Guide to Defense Preparedness

Information systems play an important role in society, so threats to their security should be taken seriously—but there is no need to panic.

Martin C. Libicki

With the cold war over, we Americans could justifiably think that our world is safer. As a nation, we could turn from hostile pursuits and begin to enlist private and public support for other great efforts—such as building a national, nay global, information infrastructure. Such an infrastructure, besides making our lives more productive and pleasant, might also serve as a bridge to global comity.

Ironically, many argue that the infrastructure is itself a venue for a new type of warfare, "information warfare," that may alter today's social order far more effectively than industrial-era contraptions could have. Modern societies rest on the reliability of their information systems. In theory, such systems are under the control of their owners. In reality, however, others can take it away. A dedicated cadre of information warriors (that is, hackers) could get into the system's computers and usurp control, bend them to their purposes and thereby place the information infrastructure, and hence the social order that it supports, at risk. Hitherto, threatening the social order required big armies and usually involved known enemies. Information warriors now argue that a few hackers may suffice. And so the very foundations of our future prosperity—not to mention education and entertainment—would be the very means by which the nation may be undermined.

Are we feeling insecure yet? Many people are. In July of last year, John Deutch, then head of the CIA, told Congress that he ranked information warfare as the second most serious threat to US national security—just below weapons of mass destruction in terrorist hands. That same day, Deputy US Attorney General Jamie Gorelick said she ranked it number one—and called for the equivalent of a Manhattan Project to restore our erstwhile security. That same month, a presidential commission was established to determine the nation's true vulnerability. Such fears have some basis in fact. The Internet suffers around a million successful penetrations every year, and tales of other computer mischief (for example, hackers transferring $10 million from Citibank accounts) are rife.

MARTIN LIBICKI *(http://www.ndu.edu/ndu/inss/libicki.html) is a senior fellow at the National Defense University in Washington, DC. The opinions expressed in this article are his own and do not necessarily reflect the views of the Department of Defense.*

But should we really worry? More specifically, should we worry either as members of society or as scientists? How great is the threat? What is its nature? Are the right steps being taken to eliminate it—or is it more reasonable to manage or work around the threat?

I explore these questions in this essay by taking a brief glance at the origins of information warfare, the nature of the threat to information systems, the particulars of the Internet as a subset of the national information infrastructure and some perspectives on how the threat may be evolving.

Military origins

Is information warfare warfare? To understand much of the current excitement about information warfare, one must understand how information is viewed by those who engage in warfare for a living—the military.

Start with the basic purpose of information: It is to inform decisions and thus action. The better the information (and thus information system), the better the decision. In war, two sets of decisions matter: ours and theirs. Offensive information warfare aims to affect information flowing to the other side so that their decisions are made to our advantage. Defensive information warfare is keeping them from doing that to us.

Because the phenomenon of information is so broad and so pervasive in all human action, a huge tree of activities can be grown from this nutshell definition of information warfare. In ancient times, for instance, information warfare consisted of attacking the enemy's commander, deceiving him as to one's whereabouts and intentions, or manipulating his political organization to achieve this or that end.

Technology's advance complicated the means—but not the ends—of information warfare. Civil War–era balloons and World War I–era aircraft permitted their owners to see how foes were arrayed for battle; those being watched tried to shoot them down. The telegraph enabled troops to be controlled over long distances; marauders would try to tear down the wires. With the invention of radio and radar, techniques of information warfare racheted up further: jamming, counterjamming, interception and broadcast propaganda. The invention and refinement of computers and other digital devices such as space-based sensors have introduced yet more venues for disruption and protection. With many militaries possessing weapons able to hit anything that can be seen and identified as a target,

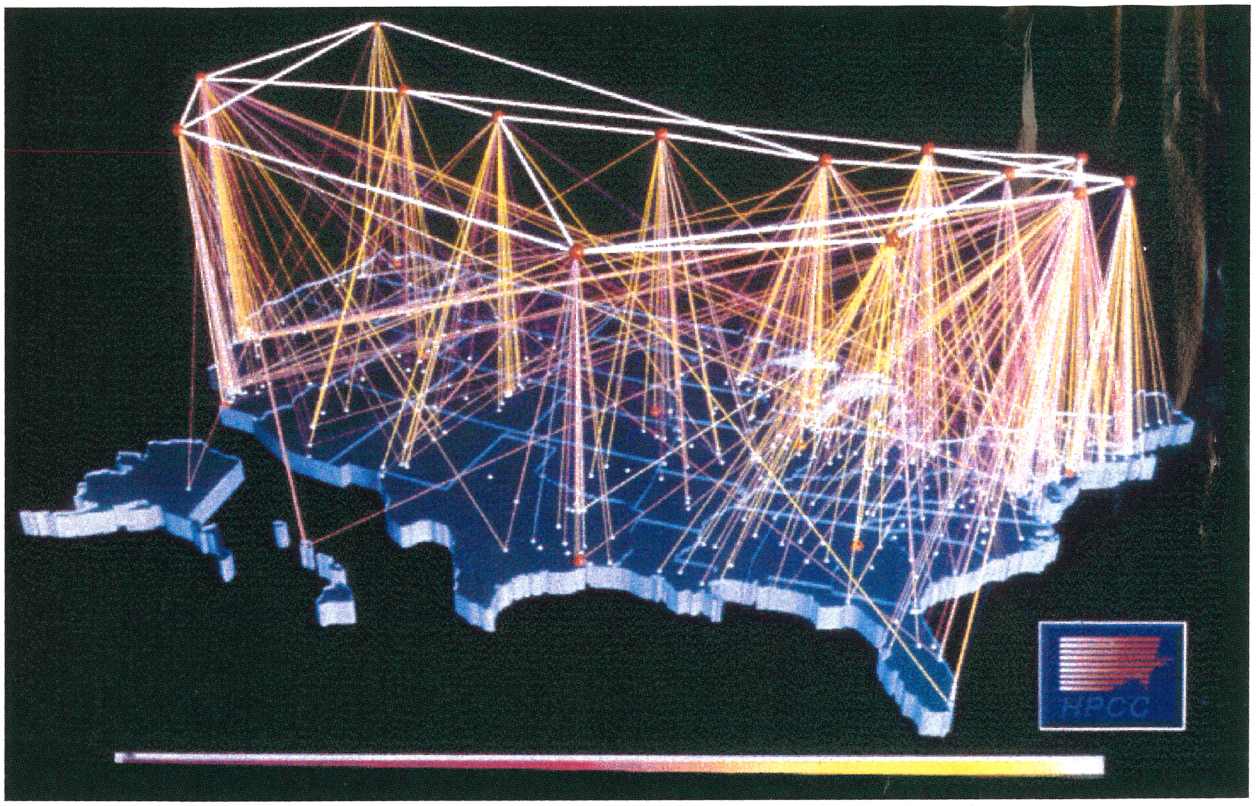

THE US GOVERNMENT'S "BACKBONE" NETWORK, together with selected regional networks. Fiber and wire connections are shown on a plane above the US outline. This large network has both strengths and weaknesses. One strength is its great redundancy. Weaknesses include vulnerability to flooding and easy access to computer administration controls. (Courtesy of the US government's High Performance Computing and Communications program.)

the process of spotting, tracking and identifying things has become central to war—as have stealth, hiding, blending in and shooting at sensors. Information has always been part of warfare. Now it is a very large part. So, therefore, is information warfare.

Ironically, even though information serves to pierce the fog of war, information warfare itself is a hazy enterprise. It requires a knowledge of the other side that is far more intrusive than that required by physical combat. One must know what information feeds decisions; how information is routed in space, time and spectrum; by what rules who gets to see what information; what (and whose) decisions are subject to human override; and many other correlated details. Absent this knowledge, operations are literally shots in the dark. Assessing the effects of operations is trickier. A great deal of data about information systems comes from eavesdropping, a technique that may wilt as encryption spreads. Little physical evidence will tell whether bit flows on the other side have been corrupted or disrupted—much less whether the enemy's ability to make good decisions has been in any way affected. Computer systems can be made to look like a labyrinth of deceptively authentic information.

Hacking as warfare

Computer hacking permits perpetrators to spy on their targets, feed them misleading information and even put them out of business. Hacking has become the icing on the information warfare cake for several reasons. First, computers are becoming ever more pervasive. Second, computer-generated information often flows unimpeded by human attention, so that considerable mischief is possible before anyone notices. Third, hacking requires neither vast resources nor huge cohorts. The more militaries depend on computers, and the more these computers are networked, the more vulnerable they become to such low-cost, almost invisible, attacks.

Hacking can be used strategically. In 1941, Japan attacked Pearl Harbor and, in the course of a few hours, immobilized the bulk of the US Pacific Fleet, thereby permitting the conquest of much of Southeast Asia and Oceania. The possibility of a cyber analogy 60 or more years later cannot be ruled out. The ability to take the US command-and-control system off-line—and keep it off-line (a far harder proposition)—could permit an attacker to wreak worldwide havoc before the US could respond.

In practice, the vulnerability of militaries to hackers is unclear, in that precise knowledge of vulnerabilities is, of course, highly classified. Both attackers and defenders of information systems tend to be paranoid, because gathering information about information warfare is itself information warfare. Yet, systems that handle classified information—for example, nuclear weapons command-and-control, space operations, reconnaissance and surveillance and espionage—are rarely connected to the outside world. Such systems are also internally compartmented, so that a leak in one place does not jeopardize operations elsewhere. And, for good measure, they are based on encrypted files and flows. True, complacency about obvious protection often makes a system open to penetration from the inside—for example, by a cyber equivalent of

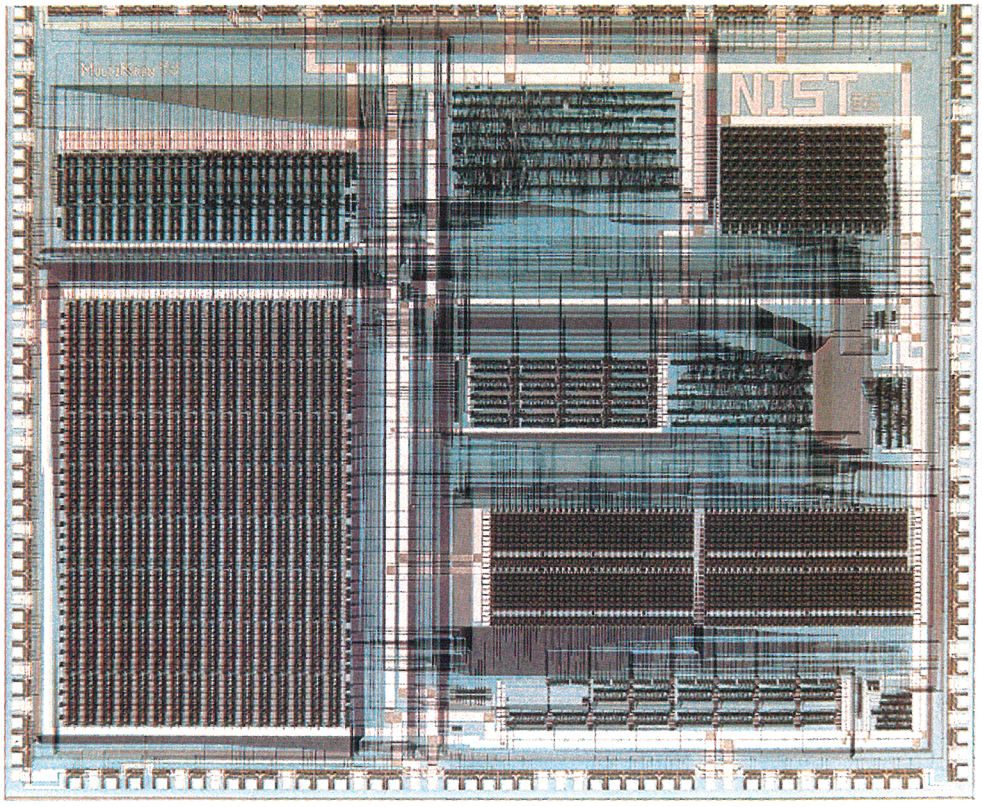

A NEW BATTLEGROUND? Those who see silicon as a venue for conflict have resurrected the vocabulary of the cold war. They talk about deterrence, warnings of attack, minimum essential information infrastructures, defense readiness conditions and so on. (Courtesy of HPCC.)

convicted spy Aldrich Ames. "Air-gapped" (physically isolated) systems are also far less efficient because of their many security precautions. Nevertheless, confounding the US military (much less other militaries, which are less networked or computerized) by breaking into its command, control and intelligence systems is no sure thing.

Defending information infrastructures

If military information infrastructures are difficult targets, could an enemy achieve similar effects by going after softer information infrastructures? This question has two parts: Can militaries themselves be so crippled, and can similar strategic effects be achieved by going around militaries and striking at societies?

To begin with, the US military relies on many systems that are poorly protected because they carry no classified information. These systems include logistics (supply, repair and transportation) and databases (finance, medical, R&D, office automation, analytic support, environmental monitoring, open-source intelligence and analysis and so on). Many such systems—for example, at Los Alamos National Laboratory and at the Air Force's Rome Laboratories in Rome, New York, have been broken into electronically with embarrassing results.

The US military also relies on civilian telephone, electricity, gas and water systems—all computer-controlled. For example, 95% of all unclassified defense communications run over the public switched telephone network. Disabling such systems could conceivably hinder or cripple military operations—but how badly and for how long are largely unknown. Armed forces are used to functioning in austere circumstances.

But why attack a military if one can attack a nation by going around it? Thus, of late, an entirely new theory of information warfare has arisen. A widespread and coordinated assault on the national information infrastructure—gaining illicit entry to sensitive domains, rewriting files, leaving trapdoors (means of easy reentry) behind—could eliminate telephone service, turn the power off, misroute transportation systems, corrupt medical devices and scramble financial and legal records. The resulting mess would be so devastating to the nation's economy and well-being that the country would be psychologically incapable of offering effective resistance to any further mischief that the attacking cabal may wish to perpetrate. Because computer intrusions by way of the Internet or phone lines can originate almost anywhere in the world, the perpetrators need never show their faces or leave behind any telltale physical evidence. Indeed, the adroit distribution of computer viruses (malicious computer code that replicates itself to other machines and media and thereafter does damage), worms (like viruses only they congest rather than damage systems), logic bombs (malicious code that activates itself upon some event or signal) and Trojan horses (malicious code buried in otherwise useful software) may allow hackers to plant the seeds of systems destruction and then vanish well before germination is required.

Could it happen?

Will hacking be the leitmotif of 21st-century warfare? Maybe—but there are three major reasons why such a strategy may be harder to pull off than it looks.

First, if it were so easy, why has nothing even close to such mischief already occurred? Of course, one could have asked the same thing just before Pearl Harbor—after all, there is a first time for everything. Yet, the opportunity to assail the information infrastructure did not create itself in 1997; dependence on information systems extends back into the 1960s. Nor does the US have but one enemy

that is waiting for just the right moment to strike. Statistics suggest that if a distribution of incidents includes many small ones and one or more large ones, then it has to include a population of intermediate ones as well.

Second, the national information infrastructure is extremely heterogeneous; it spans many sectors and owners. Some systems are accessible and some are not; some are open and some are proprietary; some are monitored continuously and others are not; some produce archived logs and some do not, and some of those logs are erasable while others are permanent. Some systems are deliberately redundant, some happen to be redundant and some have inexcusable single-point failure modes. Some systems can revert to on-site or even manual control, others cannot.

With intelligence often fuzzy beforehand and inflicted battle damage so hard to assess afterward, how would the big attack be perpetrated? Planting faults in systems in advance risks premature exposure. Planting faults in real time and expecting them all to go off without their having been tested in their environment is a recipe for disappointment even under far more benign circumstances. Yet, orchestrating a crescendo of malfunctions will alert systems administrators to the urgent need for security, which in turn could frustrate the climactic finale. Using the drip-drip-drip of induced systems failures as a strategic tool fails if they are indistinguishable from common accidental failures.

Third, the dismal state of the infrastructure's security today is, ironically, another reason to believe that hacking is not the future of war. Passwords are easy to guess; they are often common words and are too seldom changed. And they are easy to steal, from network sniffers or by tunneling from trusted computers. Events are inconsistently logged, if at all. Operating systems too often let users examine and alter files they have no business seeing. Firewalls often leak from the inside. Not only have we muddled through, but there are vast grounds for improvement.

Securing systems

Systems can be made far more secure should a deteriorating cyber environment make security imperative. Where the consequences of induced failure are sufficiently serious—for example, at a nuclear power plant—a system can be disconnected from global networks such as the Internet or the international phone system (for example, via modem). Less drastically, computers can be shielded from outside penetration if they respond only to a limited number of message types; consider, for example, the difficulty of inserting a virus into a bank's computer by hitting keys on an automated teller machine (ATM). Media that cannot be written to—a read-only memory (ROM) chip or a compact-disc ROM, for example—are impervious to attack by viruses; media that cannot be erased can log transactions so that hackers cannot go back and hide their tracks. Digital signatures, if adequately managed, are much better authentication devices than are passwords; the world never has to see the private key that authenticates both user and message (and the mathematics of digital signatures also means that users cannot use their dog's nickname as a private key). By associating every change in a database, program or other file with a specific user, digital signatures inhibit corruption by insiders.

Not every system need be secured at great cost. But critical systems—those that control key functions such as electric power, hold important data such as bank records, are needed in real time or whose owners have real enemies—should be secured.

The Internet, at any rate, is not the entire national information infrastructure, or even its most critical part. Much of what makes it so enjoyable—its open accessibility, constant evolution, lack of central control, focus on the user rather than on any specific mission—tends to detract from its security. UNIX, its dominant operating system, gives too many user-initiated processes the power to read and write files that users ordinarily should not work with. Critical systems can use Internet technologies such as packet switching but be isolated from the Internet itself.

True, the Internet is gaining both "bitshare" (portion of all electronic communication) and "mindshare" (portion of our attention). But to juxtapose tomorrow's Internet and today's security levels and so predict disaster is misleading—analogous to combining airline accident rates of the 1950s with passenger loadings of the 1990s and looking for a constant stream of funerals. Forty years ago, Boeing recognized that as long as four out of five Americans were afraid to fly, its business would be limited. It thus developed a single-point failure philosophy that ensured that no one aircraft failure would by itself be fatal. As accident rates dropped, loadings soared. Simi-

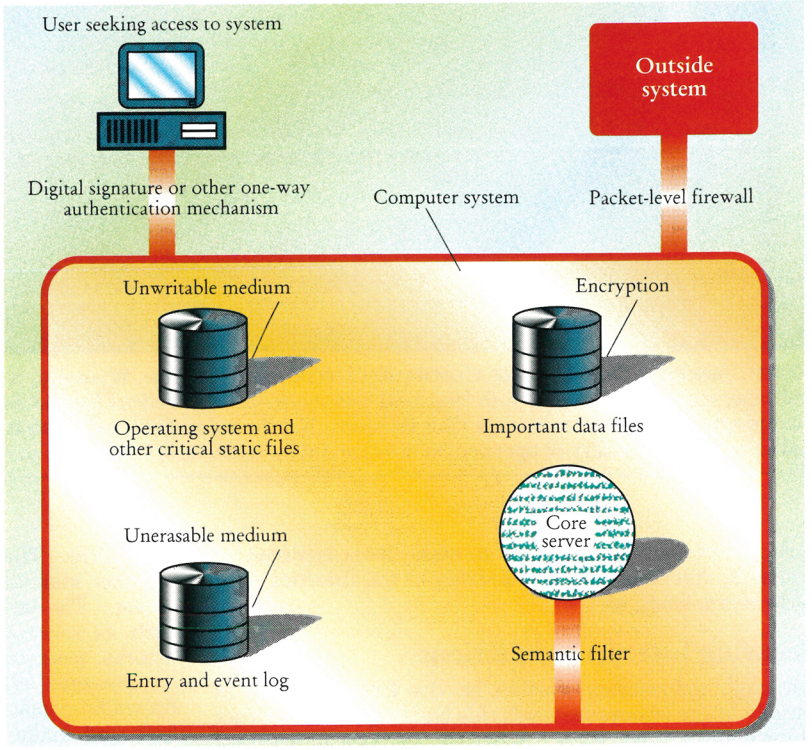

SECURITY POINTS. Systems must be secured in multiple ways and places—at their borders through access devices and firewalls, internally through semantic filters and the use of unerasable and unwritable media.

NORTH AMERICAN AIR DEFENSE COMMAND CENTER near Cheyenne, Wyoming, is a computer-laden operation. The military and other national institutions with large field operations—AT&T, Federal Express and Union Pacific, for example—rely increasingly on networks to manage their activities. But the vulnerability of a system is more than just a function of its hardware and software. The ability of operators to detect and react to anomalies, as well as to recover from disasters, is equally critical.

larly, people will be reluctant to stake their business or wallet on the security and reliability of the Internet until facts merit doing so.

Securing Internet connections

Although information warfare on the Internet today is largely at the annoyance level, this hardly guarantees that things will not get worse. What can be done to ward off an escalation?

Market forces may solve some problems. Networks are easier to flood (by, for example, overloading links or nodes with pointless traffic) because the Internet lacks a per-message or per-byte billing structure. If free Internet service deteriorates, entrepreneurs are likely to step in with service that is more guaranteed. Even within the free Internet, a consensus to tamp down could create pressure on independent service providers to identify recurrent troublemakers or trouble spots.

The Internet's structure can be improved. Achieving locktight security for every computer on the Internet is absurdly difficult, but some high-performance routers are particularly critical and should be made able to withstand any feasible attack. Today, a hacker could also theoretically get in the way of the Internet's domain name service, which translates domain names (for example, libickim@ndu.edu) into domain addresses (for example, 198.76.89.37). Proposals are being circulated that call for cryptographic methods to ensure that addresses that purport to come from domain name servers actually do originate there. The newest version of the Internet Protocol provides standard facilities for encryption and authentication. The latter should reduce many security threats such as source-address spoofing, source-related routing attacks, password sniffing and connection hijacking.

Systems administrators can do their part. Putting firewalls in place is one method, even if they sometimes leak and make systems harder to use. The Computer Emergency Response Team, based at Carnegie Mellon University, issues periodic advisories on this or that security hole; careful attention to such warnings and prompt fixes are helpful. Password administration can be improved in simple ways such as by disabling default passwords, requiring frequent changes in extant ones and requiring alphanumeric rather than alphabetic encoding. Cryptographic methods can offer very high levels of authentication and privacy—but key management is not trivial, and, at least initially, will eat up network time.

As for users, good security practices include the judicious choice of passwords, a certain wariness about downloading free software (including "applets," the small Java programs) and installation of devices to keep others from using a computer in its owner's absence. Backing up important files and programs (unless supplied by the network) can accelerate recovery in the event of disaster, of which malice is but one cause.

Unfortunately, sufficient attention to security may also require a certain fastidiousness about absorbing the newest innovations. Until recently, one could not disable

one's computer by opening e-mail. But with word-processing macros embedded in text, opening e-mail can now unleash a virus in a network or a hard disk. Web browsers can also download running code (something Java was explicitly invented to do), some of it possibly malign. Distributing objects (structured data packages with procedures that operate on them) over global networks without a good way to authenticate them leads to similar risks.

Finding the right metaphor

Is the scientific community heir to the information warfare threat? Some science is classified in nature or purpose, and its products clearly must be protected. In an age of industrial espionage, the fruits of commercial research and development may also need to be protected from exposure or corruption. Yet, by and large, scientific information is meant to be shared among all nations, regardless of their politics. Why steal what one can have for free? What use is corruption in a system that is publicly self-correcting?

These questions suggest the limited relevance of information warfare for science—an enterprise that, at least in pure form, faces few organized enemies. But is hacking actually warfare? The information infrastructure is a strategic national asset, and thus assaults on its integrity are putatively a national security concern. To call something warfare, though, is to presuppose motive and organization. Advocates have exhumed the vocabulary of the cold war: deterrence, indications and warning of attack, minimum essential information infrastructures, DEFCONs (defense readiness conditions) for information warfare alerts and extended reconstitution. "Warfare" suggests the need for a government agency to protect the infrastructure from its enemies. Yet, the US cannot build a national firewall, cannot force others to build firewalls, cannot regulate a private system's security features or inspect its source code. In cyberspace, there is no such thing as forced entry; if someone has entered your system, it is because you have created a path inward, or more typically, have failed to close those paths that lie in the software.

Perhaps hacking is more akin to disease and pollution that are present in the environment and that systems must actively filter out. Systems security may resemble systems safety—the price one pays because important complex systems carry risks. Safety engineering (for example, protecting airliners from geese) is prosaic but essential. A disease- or pollution-based metaphor has the advantages of integrating rogues and idiots together as sources of problems and draws attention to all unwanted flows in the system: from malicious computer intrusion to self-serving or false information, unwanted missives and similar wastes of resources. At the very least, such a metaphor puts the onus back on the system designers and administrators.

As computers get smarter

An information system girded with firewalls and gates, broken vertically into compartments and horizontally by access privileges, where suspicion is the norm and nothing can be trusted, will probably reduce the risk of information warfare as we know it today to negligible levels. Yet, a security architecture of increased sophistication may become increasingly intrusive and somehow antithetical to the purposes for which science in general and physics in particular are pursued. It is no accident that the World Wide Web was invented to enable particle physicists to share knowledge.

Worse, as systems (or systems of systems) become more complex, the possibility arises that even with a perfect understanding of components, one cannot predict all fault modes. Although today's hacking requires active intruders, much of the risk and tedium associated with malicious hacking can be transferred to "bots"—pieces of code that wander the Internet looking for a computer to roost in. A single hacker unleashing a flood of junk to clog a system can be traced and stopped, perhaps even caught. Similar effects, however, may be generated by implanting a virus in thousands of systems. Turning on one or two (for example, by including a key word in otherwise innocent text) can simulate a chain letter on steroids, as each infected system turns on another, with thousands of sites flooding the system, well before their owners know what is going on.

This scenario suggests a completely different approach to information security. Consider how complex humans are compared to computers and how they, like computers, network themselves to exchange knowledge. Yet, for the most part, it is difficult for humans to pass disabling bitstreams among themselves. Few word combinations can cause a rational person to seize up, spin cycles endlessly or flood friends and acquaintances with pointless communications. Why? Humans accept inputs at the semantic rather than syntactic level; messages are converted into specific meanings and then processed as such. Information and commands are screened for harmful nonsense.

Yet, as any observer of human affairs can attest, cognition does not prevent certain messages from achieving widespread and deleterious effects. Such an idea, which zoologist Richard Dawkins calls a "meme," grips the imagination and spreads; it may induce "extraordinarily popular delusions and the madness of crowds." (Indeed, "information warfare" is itself such a meme.) This is where education, notably classical education, comes in. People are trained to externalize and objectivize information and view it as something separate from the self—to examine information critically and thereby avoid the twin hazards of excessive naïveté or total cynicism. Rhetoric deals with the proper inference of intentions from statements. Logic deals with generating correct implication therefrom. Language is taught as a precision tool so that information can be passed and understood with clarity. Etiquette governs what people say (and how intrusively) and inculcates propriety so that people keep private information private. Thus can humans interact with others, even strangers, to form functioning organizations and other relationships.

Building such sophistication into silicon will not be easy, either technically or socially. To exchange information, or even to agree on what constitutes a legitimate piece of information, requires conformance to shared norms at several levels: of meaning, of intention, of context, of behavior. Without shared norms—standards as it were—there is no meaning but only programmed reactions. As long as that is true, it will be difficult to defend information systems from a potentially polluted environment except through mechanisms antithetical to the scientific ethos.

The irony is that one cannot expect to make computers consistently more powerful, and expect them to remain safe, if they are not endowed with the power not only to input and output bytes but also to understand them. But, having taught them to listen and speak, we must simultaneously teach them proper manners. And that, ultimately, is what effective defense against information warfare is about. ■