Rebecca T. Mercuri

Superscaled Security

Exponential increases in computational speed, memory capacity, and bandwidth impose futuristic security demands and challenges.

Advances in high-performance computing have found their counterpart in new security threats. Yet there is an interesting twist in that computational expansion tends to be relatively predictable, whereas security challenges are typically introduced and mitigated (when possible) in a more chaotic fashion. Few would have surmised, for example, that spam would have exceeded 60% of all email transmissions by the end of 2003, nor that detection software would require sophistication approaching Turing Test intelligence levels. It is useful, therefore, to consider some of the impacts of scaled-up computing on our overall security environment.

Certain rules continue to apply to computational evolution. Moore's Law (which anticipates processing power doubling approximately every 18 months while the equivalent price halves in the same amount of time) has continued to endure. It is likely to surpass even Gordon Moore's own 1997 prediction [5] that it will "run out of gas" around 2017, as new materials and fabrication technologies emerge (including those introduced through the recursive application of improved computers into the manufacturing process). In 2003, Moore told the International Solid-States Circuit Conference "I remember thinking one micron (a milestone the industry passed in 1986) was as far as we could go because of the wavelength of visible light used at the time. Engineers then switched to ultraviolet light."

But as advancements in hardware continue to occur, security appears to be declining (although not necessarily at an equivalent rate). As Paul Kocher of Cryptography Research, Inc. asserted, "Moore's Law, coupled with the business imperative to be more competitive, is driving vendors to build systems of exponentially increasing complexity without making security experts exponentially smarter to compensate." "The current trend is to build systems that conceptually are secure, but in practice and probability the systems' ability to resist the efforts of a creative attacker are less," said Kocher in a pre-USENIX 2002 interview. "The idea is that security should be getting better, however, design flaws are becoming an increasingly catastrophic problem" [2].

Some believe that cryptography can help reverse this trend, and cryptosystems are playing a significant role in providing security assurances. Yet each generation of computers brings with it an obsolescence of some earlier cryptographic methods, usually considerably sooner than has been predicted. This will likely continue to be the case, and may even accelerate as new paradigms for algorithmic attacks, including distributed techniques, evolve. So this implies that if cryptography is used to sign and protect data and

PETER HOEY

software (especially for archival purposes) then a systematic method for updating these wrappers will need to be devised, lest trust in content be undermined.

Cryptography, though, cannot be expected to solve all security problems. PGP encryption guru Phil Zimmerman told London's 2003 Infosec security conference that Moore's Law is a "blind force" for undirected technology escalation. He explained that "the human population does not double every 18 months, but its ability to use computers to keep track of us does," and added, "you can't encrypt your face." He fears the series of initiatives in U.S. homeland security have far-reaching effects on privacy because "it has more inertia and is more insidious. When you put computer technology behind surveillance apparatus, the problem gets worse." We already yield a tremendous amount of personal information to our PDAs and permit tracking of our movements in exchange for continuous incoming telephone service, so the devices we voluntarily adopt may ultimately prove more untrustworthy than the monitoring being imposed. The increasing ability of computers to store and analyze vast quantities of information also increases the likelihood that this data will eventually be used for heretofore unknown and potentially nefarious purposes.

We may have no choice regarding such data dissemination, since global economic forces may be driving us toward total interconnectivity of all humans on the planet. This is ultimately feasible, since in the networking world, Gilder's Law dictates that the total bandwidth of communication systems will triple every 12 months. (By comparison, the number of humans is only expected to double from 6 to 12 billion by 2100.) Metcalfe's Law rates the value of a network as proportional to the number of nodes squared, so the merit of connectivity increases exponentially as units are introduced, while costs tend to remain stable. But this supposition of continually increasing payback is not necessarily correct. Jake Brodsky reminds us of Newton's Third Law ("For every action, there is an equal and opposite reaction"), noting that events are "going on under our noses this very minute: Security holes, hacking, and phreaking. The very tools that make this 'revolution' [in technology] possible are also being used against us." Whether the negatives will eventually outweigh the positives, in terms of adverse impacts on connectivity and usability, is yet to be determined.

There is a storehouse of data that we are looking forward to having online, which is currently located in libraries, recordings, and research databases. As David Sarnoff predicted in an article he wrote for the *New York Herald* in 1922, "It is inconceivable that the development of the transmission of intelligence will go forward at a leisurely pace; everything points to a very great acceleration." One wonders whether Sarnoff might have imagined that such acceleration would, within the next decade, make it feasible for a laptop computer to hold the contents of the entire Library of Congress. Certainly it is imperative to ensure that all information will be replicated correctly in such compendia. Imagine an insidious virus that permutes documents such that history eventually reflects that Thomas Jefferson was the first president of the United States. This might not be so damaging, but other transformations could have dire results. Even if data is maintained intact, accessibility to powerful search and logic engines

The ultimate question must be whether or not there will someday be a computational system that can prevent all forms of nefarious attack.

might (justifiably or inappropriately) result in different interpretations of the "truth" than are currently commonplace.

Knowledge acquisition, as we know, is not merely equivalent to information gathering. Knowledge involves analysis and conjecture, whereas information can be obtained in a brute force fashion. For a long time, the sheer power of bulk data provided by the information age has prevailed, but as our chess-playing machines have demonstrated, progress eventually must include an "artificial intelligence" component. As Frederick Friedel, Gary Kasparov's assistant exclaimed, "As Deep Blue goes deeper and deeper, it displays elements of strategic understanding. Somewhere out there, mere tactics are translating into strategy. This is the closest thing I've seen to computer intelligence. It's a weird form of intelligence. But, you can feel it. You can smell it!"

The combination of human reasoning with computer-based data can be extremely powerful. Once the Library of Congress is reduced to the size of a microchip, we might choose to directly implant it into our heads so the contents would be immediately accessible to our thought processes. I recall a conversation with Princeton mathematician John Conway in which we mused about *Matrix*-style cyber-brain enhancements and whether people might eventually be discriminated against in employment if they refused to submit to electronic augmentation. Although I'd prefer to retain my Luddite status, Conway indicated his willingness to enter this Brave New World of embedded micro-technologies. Ethics and security issues abound—for example, recalls for bug fixes might become a bit daunting. Recent vendor speculations have included weaving chips into our clothing or concealing them in medicines, foods, and beverages. The phrase "I've got you under my skin" may take on an eerie meaning in the not-too-distant future.

Back in the world of large computational systems, as these continue to expand, von Neumann architectures will eventually be superceded by or augmented with neural networks, genetic algorithms, and DNA "soups" where even NP-complete problems may someday become solvable. Instead of using detection and eradication, security systems will contain adaptive features—enabling them to encapsulate and incorporate malware into themselves, thus evolving to resist new threats. Of course, viruses will be smarter as well, learning about the systems they infect, beyond basic exploitation of known vulnerabilities. Rather than downloading patches, inoculation software may be injected into networks to counter the effects of currently circulating deleterious code. One might go so far to say that such a mutational process imposed on computational systems will be a necessary aspect of its development, beyond that which humans would be able to provide through programming and intentional design.

Clues to the manner in which architectural expansion may be mitigated are found in a Mead and Conway principle earlier intended for VLSI, as follows: "The task of designing very complex systems involves managing, in some highly structured way, the space and time relationships between the various levels of system building blocks so that the entire system will function as intended when it is finished." Citing that statement, a 1998 NSF workshop report [1] concluded that "the software research literature is replete with design principles, from the highest to the lowest levels. In virtually all cases, the principles have neither been carefully classified for applicability nor validated for practical effectiveness." Clearly more work is necessary. Scalability must be considered. As well, form and functionality in large-scale computing have not kept pace with each other for a number of years, and their disparity introduces the potential for a panoply of adverse consequences.

This disparity includes the relationship of our concept of trust to the necessity that computer systems maintain fundamentally deterministic behavior patterns by disallowing the use of self-modifying code. So if we relinquish control and allow evolutionary processes to occur, it becomes less clear how to assess trustworthiness. Ultimately, we might want to permit systems to adjudicate among themselves whether or not they are to be trusted, since it may not be possible for scientists and engineers to appropriately assess a breed of computers with constantly changing logical infrastructures. But then how can we trust them to make the right decisions about each other?

Perhaps this can be understood by examining some concepts from complexity science. As Jim Pinto explained to the Chaos in Manufacturing Conference [6], critical complexity occurs when processing power becomes intelligence, and when connected intelligence becomes self-organizing. Emergent behavior results from organized complexity, such as can be found in mutations, selection, and evolution. William Roetzheim defines complexity science as "the study of emergent behavior exhibited by interacting systems operating at the threshold of stability and chaos."

It would appear that such a threshold occurs whenever security breaches reach the point of threatening the stability of computational systems. For example, Gilder's Law has allowed faster bandwidth to overcome the adverse impacts of spam. As well, Moore's Law has provided increased computational speeds to allow background virus detection. As attacks have become more sophisticated, new levels of hardware exponentiation have arrived, pushing the instability threshold back. But this will soon be insufficient, since attacks that include emergent behavior will likely only be able to be thwarted through intelligence and mutation. Self-modifying processes could therefore become the panacea rather than the plague, if we can determine how to deploy them effectively in our security toolkits.

So, the ultimate question must be whether or not there will someday be a computational system that can prevent all forms of nefarious attack. To simplify this issue, let us restrict our consideration to breaches initiated by humans. Ray Kurzweil refers to the time when computers exceed human intelligence as "The Age of Spiritual Machines" [4] and believes we will be entering this era within the next two decades. He conjectures that by the end of this century, self-replicating hardware, decentralization, and redundancy will remove security concerns, for all practical purposes.

On the other hand, Seton Hall's philosopher/physicist Stanley Jaki, using Gödel's incompleteness theorem, concludes that "the fact that the mind cannot derive a formal proof of the consistency of a formal system from the system itself is actually the very proof that human reasoning, if it is to exist at all, must resort in the last analysis to informal, self-reflecting intuitive steps as well. This is precisely what a machine, being necessarily a purely formal system, cannot do, and this is why Gödel's theorem distinguishes in effect between self-conscious beings and inanimate objects" [3]. The extrapolation from this must be that self-modifying code, without the ability to perform introspection, cannot surpass human intelligence, hence systems will remain vulnerable.

Which answer is correct? With any luck, we shall all live long enough to see how things turn out. If not, then perhaps future readers, happening across this article in their Library of Congress brain chips, will hopefully have retained the ability to either laugh, or shake their heads and sigh.

References

1. Basili, V.R. et. al. NSF Workshop on a Software Research Program for the 21st Century, Oct. 15–16, 1998, published in *ACM SIGSOFT 24*, 3 (May 1999).
2. Hyman, G. The Dark Side of Moore's Law. (Aug. 7, 2002); siliconvalley.internet.com/news/article.php/1442041.
3. Jaki, S.L. *Brain, Mind and Computers, 3rd ed.* Regenery Gateway, 1989.
4. Kurzweil, R. *The Age of Spiritual Machines*. Viking Penguin, 1999.
5. Moore, G. An update on Moore's Law. *Intel Developer Forum*, (Sept. 30, 1997); www.intel.com/pressroom/archive/speeches/gem93097.htm.
6. Pinto, J. Symbiotic life in the 21st century. In *Proceedings of the Chaos in Manufacturing Conference* (May 4, 2000); www.calculemus.org/MathUniversalis/NS/09/symbiotic.html.

Rebecca T. Mercuri (mercuri@acm.org) is a research fellow at the John F. Kennedy School of Government, Harvard University.