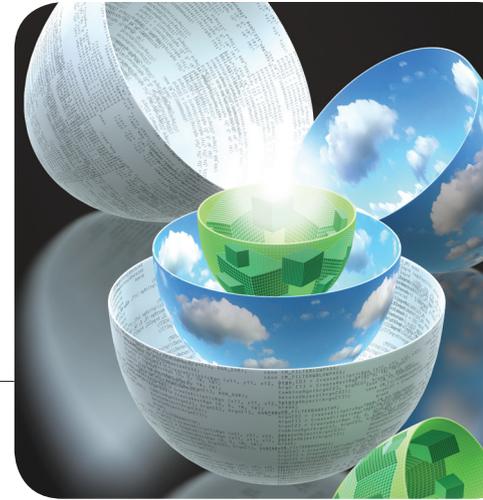


Toward Scalable Trustworthy Computing Using the Human-Physiology-Immunity Metaphor

Achieving scalable trustworthy computing is possible through real-time knowledge-based decisions about cybertrust. This vision is based on the human-physiology-immunity metaphor and the human brain's ability to extract knowledge from data and information.



LEE HIVELY
AND FREDERICK
SHELDON
Oak Ridge
National
Laboratory

ANNA CINZIA
SQUICCIARINI
Pennsylvania
State
University

Recent US federal policy documents have emphasized the importance of cybersecurity for society's welfare (see Figure 1). For example, *Cyber Security: A Crisis of Prioritization* described 10 technologies needed for cybersecurity.¹ The *Federal Plan for Cyber Security and Information Assurance Research and Development* discussed 49 cybersecurity technical topics in eight major R&D areas with corresponding funding priorities.² The Department of Homeland Security's *Roadmap for Cybersecurity Research* listed 11 "hard problems" (eight from the 2005 Infosec Research Council Hard Problem List).³ The *National Cyber Leap Year Summit Co-chairs Report* discussed five cross-cutting solution themes.⁴ (For more on these documents and the cybersecurity problem's scope, see the sidebar.)

Unfortunately, the cybersecurity landscape consists of an ad hoc patchwork of solutions. These solutions have failed to prevent cybercrime or fraud losses, which amount to untold billions of dollars each year. Clearly, we need innovative, effective solutions.

To be effective, cybersecurity solutions must support scalability. To enhance scalability, high-assurance systems should consist of composable components and subsystems, in a system architecture that inherently supports facile composability.³ (Composability is the ability to create systems and applications with predictably satisfactory behavior.) Each component and subsystem should itself be suitably trustworthy, down to the most basic level, thus avoiding development of new methodologies at each successively larger scale. Moreover, scalability should enhance trustworthiness

in areas such as constructive system design, meticulous use of best

practices, error-correcting code to overcome unreliable communications and storage, and encryption to protect insecure communications' integrity and confidentiality. Such techniques are incomplete if they rely on the trustworthiness of developers, users, and administrators. The challenges are, then, to develop

- a sound basis for composability that scales to large, complex, trustworthy systems;
- trustworthiness evaluations of composite systems that are themselves composable and scalable; and
- components, analysis tools, metrics, and testbeds for the solutions we just listed.

We believe that the human body's management of complexity, nonlinearity, and the immune response can act as a metaphor for scalable trustworthy systems. We call this the human-physiology-immunity (HPI) metaphor. Here, we look at examples of human-body functions that hint at achieving scalable, trustworthy solutions, and we outline steps toward enabling this new paradigm.

The Current Scenario: Performance versus Insecurity

An important challenge for cybersecurity is to keep pace with modern systems' evolution. Figure 2 depicts this evolution, revealing that computers' speed

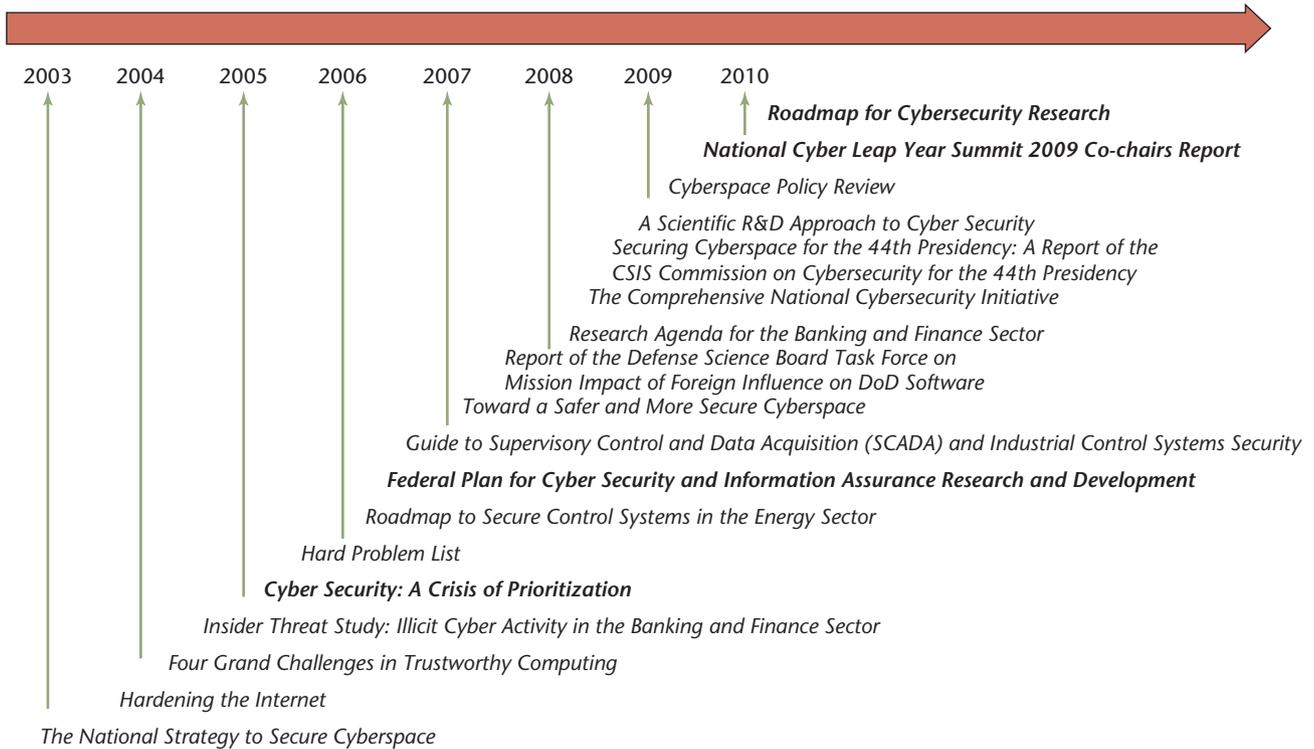


Figure 1. A timeline of selected US federal cybersecurity initiatives, which have all emphasized the importance of cybersecurity for society’s welfare. The bold titles are documents discussed in this article. All the reports are available at www.cyber.st.dhs.gov/documents.html.



Figure 2. Computational performance over time for the world’s top 500 computers. Performance is now limited by parallelization and energy consumption, rather than individual processor speed. (Source: TOP500.org; used with permission.)

has increased by more than 10^4 over the past 15 years. Performance is now limited by programming parallelizability and energy consumption, rather than individual processor speed. Interestingly, the current fastest Tianhe-1A (2.5+ Pflops consuming 4+ megawatts) supercomputer is three times more power efficient than its closest rival (Jaguar, 1.7+ Pflops) and is expect-

ed to be eclipsed in 2012 by two different machines, the Sequoia (20 Pflops consuming 6+ megawatts) by IBM, with world-leading energy efficiency, and a yet-unnamed Cray being built by Lawrence Livermore National Laboratory and Oak Ridge National Laboratory, respectively. Exascale or extreme-scale 1,000+ Pflop machines are predicted in the 2019 time frame.

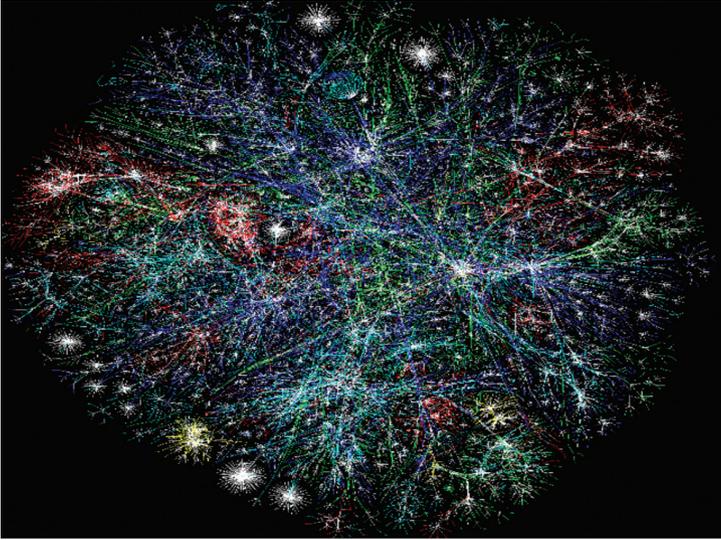


Figure 3. A map of Internet activity over a one-day period (23 Nov. 2008). Red indicates the source is Asia or Pacifica; green indicates Europe, the Middle East, Central Asia, or Africa; blue indicates North America; yellow indicates Latin America or the Caribbean; cyan indicates private networks; and white indicates an unknown source. (Source: Barrett Lyon; used with permission.)

Computational improvements have been accompanied by “ubiquitous insecurity.”⁵ Attacks through the Internet frequently employ malware, involving deliberate infiltration or damage to a computer system. Attacks range from “low and slow” over a day or longer to “fast and focused” attacks at the millisecond level or faster. Such attacks might lie hidden in a sea of normal cyber activity.

In addition, internal attacks can have devastating consequences, including elevated privileges for malware that’s directed by external agents. Insider threats can take many forms, including espionage or cybercrime. Although policy violations can be the result of carelessness or accidents, the core concern is deliberate and intended actions such as malicious exploitation, theft, or destruction of data, or the compromise of networks, communications, or other IT resources.

The greatest challenge is the continuous evolution of attacks. Previous solutions for known threats might not address new attacks, whose effectiveness and disruption are hard to predict. Traditional risk methodologies provide common-sense advice but usually lack specific guidelines for evaluating emerging threats. So, we need better protection from future threats at all levels.

Cybersecurity—a Very Hard Problem

Cybersecurity is a hard, multifaceted problem for the following reasons.

Complexity at All Levels

As Figure 3 shows, the Internet is highly complex and seemingly scale free. All modern computers are themselves networks of systems (for example, CPUs, memory, GPUs, storage, data buses, and I/O devices). All modern software is a complex network of processing functions. The information infrastructure is a complex system of systems of hardware, software, OSs, data, networks, and people. Complex interactions frequently produce emergent, unexpected, and potentially adverse behavior. Failure in such an infrastructure can be so complex that no one can determine the cause, let alone a cure. Scalable trustworthy systems must cope with this complexity.

Immense Amounts of Data

The estimated amount of data globally is 451 exabytes (4.51×10^{21} bytes), or 72 Gbytes for each person on Earth.⁶ Scalable trustworthy systems must be able to process more of this tsunami of data in close to real time for attack characterization, situational awareness, attribution, and appropriate response.

Problems Converting Data to Knowledge

Cybersecurity decisions require converting data into information and hence into knowledge. Analyzing data in the context of other data generates information; processing that information in the context of other information creates knowledge (see Figure 4). Current systems can’t create knowledge; they rely on decisions by humans who can’t respond at computer speeds of milliseconds or less. Moreover, a human can’t detect sparse anomalies in the knowledge-discovery process. Robust cybersecurity requires a new paradigm. Scalable trustworthy systems must process the tsunami of data in nearly real time to enable knowledge-based decisions about cybertrust.

Practical Constraints

Cybersecurity has five practical constraints. The first is protection of private information (which is essential for public acceptance). The second is appropriate handling of imperfect data (errors, incompleteness, inconsistency, and noise). The third is usability and cost effectiveness, including the need to

- scale from the smallest sensor on a chip to the largest high-performance resource,
- allow cross-platform development and interoperability with legacy systems,
- comply with the mandates of law at all levels,
- provide for graceful degradation of safe operation during failure, and
- minimally impact users’ ability to perform real work.

The fourth constraint is facilitation of open source software use, parallelism, debugging, and software quality assurance. The fifth is the enabling of multilanguage development for multiple applications. Scalable trustworthy systems must interoperate with legacy systems within constraints that are reasonable and within the context areas we just outlined.

Inadequate Perimeter Defenses

Traditional cybersecurity approaches focus on a layered defense, or defense-in-depth, by erecting physical or cyber walls and fortifications between the layers. This approach is ineffective against malicious insiders as well as malicious outsiders who successfully break in and become indistinguishable from insiders.

Fortification of individual processors on the network doesn't fortify the network. Rather, active, distributed security must be an integral part of novel hardware-software combinations such as

- computers that keep secrets or ignore malware, just as humans can harbor viruses without illness;
- intrinsically secure devices that share provable trust information, confirming their trustworthiness;
- security-hardened hardware that's highly resistant to hacking; and
- systems that determine the trustworthiness of hardware, software, networks, and users (for example, white listing).

Scalable trustworthy systems must provide accountability for all users, software, hardware, and networks.

More and Smarter Attacks

Cyberattacks are growing in number and sophistication. Recent examples include organized nation-state attacks against the Pentagon and other US facilities and against Estonia, Georgia, and Iran;⁷ a rise in identity theft via the Internet; undocumented features in open source applications code (software life-cycle problems); open source flaws (typically on the order of 1 per 10^3 lines of code); the use of botnets and other organized Internet techniques; website and Web application exploits; and the compromise of unsecured data.

So, Are Trustworthy Systems Possible?

One line of reasoning maintains that completely trustworthy systems are impossible. All modern software is complex, as are hardware, networks, and interactions among users. Moreover, flaws, including both malicious and honest mistakes, in complex systems are difficult to detect, analyze, and correct. So, all modern complex systems have vulnerabilities. Updates compound this complexity.

In addition, ubiquitous networking opens a vul-

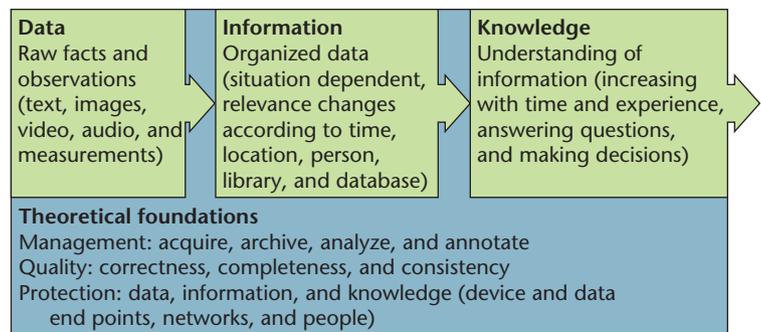


Figure 4. Converting raw data into information (data in the context of other data) and hence into knowledge (information in the context of other information), for understanding and prediction.

nerable computer to Web-based attacks. Most vulnerabilities arise from exploitation of built-in flaws in the security features. For example, network infrastructure enables widespread, distributed attacks, which are readily propagated among networked, homogeneous computing environments. Users frequently use their resources in unanticipated ways. According to the line of reasoning we're describing, the root cause of vulnerabilities is the inherent imperfection of software, hardware, and networks, which by their nature can never be totally secure.

We believe this argument can be refuted simply by viewing examples of complex, living organisms that effectively manage complexity and provide secure, real-time responses.

Why Scalable Trustworthy Systems Might Be Possible

The human brain and immune system provide two compelling models for scalable trustworthy systems.

The Human Brain

The brain exhibits superior speed and insight in processing disparate data for real-time situational understanding and decision-making. For example, a person can read these words and understand the message in real time (approximately 1 sec.) through neuron-based processing with a single-neuron cycle time of approximately 10 milliseconds. This corresponds to approximately 100 neural hops (10^2 processing cycles) per second over approximately 10^{11} brain neurons, for a net processing power of approximately 10^{13} cycles per second.

Modern high-performance computers run at more than 10^{15} operations per second, or 100-fold more (1-million-fold expected by 2019) processing power than the brain, yet they can't perform "intelligent" real-time processing of the same data. Accordingly, we view the human brain's capacity for intelligent,

The Cybersecurity Problem's Scope

Table A. Some recent US federal cybersecurity priorities.*

Federal effort to characterize a problem			Solution themes
Cybersecurity priorities, President's Information Technology Advisory Committee (2005)	Some top cybersecurity and information assurance R&D priorities, US Nat'l Science and Technology Council (2006)	Hard-problem list (ver. 2), US Dept. of Homeland Security (2009)	National Cyber Leap Year Summit (2009)
Authentication (3)	Authentication, authorization, trust management, access control, and privilege management (4)	Scalable trustworthy systems (including system architecture and the requisite development methodology) (4)	Hardware-enabled trust (knowing when you've been had)
Secure software engineering (2)	Large-scale cyber situational awareness; automated attack detection, warning, and response (3)	Enterprise-level security metrics (measures of overall system trustworthiness) (3)	—
Holistic system security (2)	Insider threat detection, mitigation, forensics, traceback, and attribution (4)	System evaluation life cycle (including approaches for sufficient assurance) (2)	Cybereconomics (crime and fraud don't pay)
Monitoring and detection (3)	Secure Domain Name System and routing, and protocols and process control systems (3)	Combating insider threats (3)	—
Secure fundamental protocols (2)	Domain-specific security (for example, wireless and RFID) (2)	Combating malware and botnets (3)	Moving-target defense (attacks work once if at all)
Mitigation and recovery (1)	Detection of vulnerabilities and malicious code; metrics and software testing and assessment (3)	Global-scale identity management (3)	—
Cyberforensics (3)	Secure OSs, software engineering, and information provenance (3)	Survivability of time-critical systems (4)	Digital provenance (basing trust decisions on verified assertions)
Modeling and testbeds (3)	Cybersecurity, information assurance R&D testbeds, IT systems, Internet modeling, simulation, and visualization (3)	Situational understanding and attack attribution (2)	—
Metrics, benchmarks, and best practices (3)	Trusted computing base architectures and composable, scalable, and secure systems (3)	Provenance (relating to information, systems, and hardware) (4)	Nature-inspired cyberhealth (moving from forensics to real-time diagnosis)
Nontechnology issues (2)	Inherently secure, high-assurance, and provably secure systems and architectures (3)	Privacy-aware security (3)	—
—	Trust in the Internet and privacy (3)	Usable security (3)	—

* Progress in a solution theme area will support advances in the related problem areas; numbers in parentheses indicate the number of solution themes that the problem area covers.

real-time, knowledge-based decisions as a basis for envisioning scalable, secure situation awareness.

Jeffrey Hawkins and Sandra Blakeslee's 2004 book *On Intelligence* focused on the brain's neocortex, which has approximately 10^{10} neurons and 10^{14} connections.⁸ The neocortex's key features are

- an irreducible representation for each item in memory;
- auto-associativity among items (for example, recalling one line of a song leads to the remainder), because a memory is recall of a time-serial sequence of stored items;
- hierarchical processing (for example, combining the

Table A maps the problem space to the solution space through analysis of the four documents we mentioned in the main article’s introduction. A long-term vision for scalable trustworthy systems requires solutions for all the problems listed in the table.

When analyzing these documents, we considered this question: if a particular priority or hard problem is resolved, what solution theme will be addressed? For example, we found that attempts to improve or deploy stronger authentication (Table A, column 1, row 1) would be quintessential toward progressing three themes in the table:

- hardware-enabled trust,
- moving-target defense, and
- digital provenance.

However, such efforts would likely play only a minor role in the other two themes—cybereconomics and nature-inspired cyberhealth. More details about this analysis appear elsewhere.¹

In the table, a number in parentheses indicates the number of solution themes that a problem area affects. The larger the number, the stronger that problem area’s cross-cutting nature. Quite a bit of overlap occurs, indicating that the priorities and problems haven’t significantly changed over time. The federal government’s recent efforts to simplify all these into five themes will undoubtedly help focus public and private research in these areas.

It seems reasonable to conclude that federal funding and policy will emphasize these themes. One measure of progress would be an increase in the number of indictments, convictions, and extraditions from the countries that are cybercrime havens. Progress in winning international agreement on norms, collective defense, cybercrime prosecutions, and IP protection will let us gauge international efforts to enhance cybersecurity. Cybereconomics is one important lever that governments will need to incentivize good cyberbehavior as well as deter the bad.

Reference

1. F.T. Sheldon and C.A. Vishik, “Moving toward Trustworthy Systems: R&D Essentials,” *Computer*, vol. 43, no. 9, 2010, pp. 31–40.

simplest spoken sounds or phonemes into words, which then are combined into phrases that form sentences and concepts); and

- feed-forward links to make appropriate connections among phonemes, words, phrases, sentences, and concepts in the context of previous knowledge.

There’s also feedback from higher to lower levels in the hierarchy for self-consistent extraction of knowledge in terms of known words (rather than nonsense words), proper syntax, correct grammar, filtering out an accent, situational context, and so on. Likewise, image processing identifies things such as points, lines, polygons, objects, familiar scenes, and scene changes.

The same neocortical processing paradigm extracts a hierarchical sequence of patterns for all time-serial sensory data, such as auditory and somatosensory data. Blind people can learn to read in braille, sense crude images via discrete touch points on the tongue, or sense soundscape images via stereo headphones. Understanding is the essence of intelligence, as is the ability to predict a situation correctly on the basis of previous knowledge. This hierarchical, brain-based paradigm differs considerably from the present program-counter-based programming paradigm and might provide insight for the data-to-information-to-knowledge processing paradigm of Figure 4.

This brain-computer metaphor assumes that the cell—in this case, a neuron—is the basic unit for information processing. This assumption stems from research by R. Quiñ Quiroga and his colleagues.⁹ They recorded the response of 137 human-brain neurons, 44 of which responded only to a specific object (for example, a picture of Jennifer Aniston). This response occurred for different views of the same object (for example, the front versus the side). These observations are consistent with Hawkins’ irreducible (invariant) memory representation.

The Human Immune System

Healthy humans can live for 70 or more years, while thwarting continuous attacks from diverse microbes, toxins, and health-endangering conditions. You could view each cell as an information processor that receives input, processes it, and produces some output. More than 200 human cell types combine to form a complex architecture of tissues, organs, organ systems, and whole-body systems-of-systems. This hierarchical architecture is scalable to approximately 10^{14} cells in a healthy adult.

All body systems participate in immune functions (see Table 1). Complex, adaptive human behavior arises from interactions among the tightly integrated, hierarchical components, which consist of massively parallel, cellular processors. Knowledge-based decisions can’t process arbitrary instructions and therefore aren’t hackable.

These examples might provide insight for scalable trustworthy computing via an integrated, active, distributed, hierarchical hardware-software composition (as we discussed earlier) with proper design, implementation, and “hygiene.” Perhaps, inherently scalable trustworthy systems are those with an architecture for

Table 1. Example human immune functions.⁸

Body system	Immune functions
Circulatory	Blood-distributed immune cells throughout the body
	Recovery of immune cells through lymphatic flow
Digestive	Continuous salivary cleansing of the mouth through lysozymes
	Pathogen destruction by HCl in the stomach
Endocrine	T-lymphocyte programming messages through thymus hormones
	Depression of immune activity through stress
Immune	Capture and destruction of pathogens at surface membrane barriers by phagocytes
	Natural-killer-cell attack of virus or cancer
	Inflammation to isolate site, attract phagocytes, dispose of dead cells, and promote repair
	Fever response by pyrogens to enhance repair and inhibit pathogens
	Apoptosis
	The major histocompatibility complex
Muscular	Movement to avoid or protect from pain, heat, or danger
Nervous	Fight-or-flight response
	Avoidance of unhealthy actions (for example, smoking) or promotion of healthy habits (for example, exercise)
	Enhancement or inhibition of immune functions through serotonin, norepinephrine, or epinephrine
	The blood-brain barrier
Reproductive	Inhibition of bacterial and fungal growth by the vagina's acidic mantle
Respiratory	A physical barrier for and entrapment of microorganisms by mucous (larynx, pharynx, and nasal cavity)
	Removal of debris-laden mucous from lower tract by cilia
	Filtering and entrapment of microorganisms by nasal hairs
Sensory	Cerumen and hairs as external barriers in the ear
	Foul tastes to prevent eating unhealthy food
	Continuous eye cleansing by tears with lysozyme
Skeletal	Production of blood (immune) cells in bone marrow
Skin	A mechanical barrier against entry of pathogens and toxins
	Perspiration as a bacterial growth inhibitor
Urinary	The acidic pH of urine as a bacterial inhibitor
	Cleansing of lower urinary tract with each voiding
	A bactericidal chemical in sebum
	Resistance against acids, alkalis, and bacterial enzymes in keratin

only “healthy” functions, rather than the patches-on-patches (PoP) approach to preventing future attacks. It’s questionable whether PoP can result in a smaller attack surface over the long term.

Analogies of Immune Function

The similarities between cybersecurity systems and the manner in which biological systems ward off threats have sparked research into specific applications. Examples include the immunocomputing and artificial immune systems that were investigated during the 1990s. From an information-processing perspective, several immunological principles make the analogy appealing, including distributed processing, pathogenic pattern

recognition, multilayered protection, decentralized control, and diversity and signaling.

We now consider relevant analogies for potential scalable trustworthy solutions, as an extension of present research.¹⁰ Each example addresses one or more of the problems we mentioned previously.

The Blood-Brain Barrier

The blood-brain barrier (BBB) is a three-layer membrane that controls the passage of substances between the central nervous system (CNS) and local blood vessels. A cyber analogy is physical isolation of the CPU from the rest of the cyberworld via a fast, in-line encryptor/decryptor chip (EDC). The BBB ef-

fectively protects the brain from infections by using carrier-mediated transporters (for example, glucose) to ferry low-atomic-weight substances (≤ 500 daltons) to and from the CNS. A cybersecurity analogy is short, encrypted packets sent via single-use keys. Strict physical isolation of the CPU could include a processor-resident OS on encrypted read-only memory that's distinct from applications. Tamper resistance in the CPU/EDC (not unlike the brain inside the skull, although distinct from the BBB) could shut down an always-on processor upon tamper detection, thus erasing the OS and any sensitive data. Answers to other questions, such as how the brain self-heals and restores lost memory, will certainly enable deeper understanding of intelligence and its cybersecurity analogies.

The Major Histocompatibility Complex

The major histocompatibility complex (MHC) distinguishes self from nonself. For instance, a blood-born immune cell, such as a leukocyte, encounters a foreign invader and engulfs and destroys it. It then displays random fragments or antigens on MHC molecules attached to its outer cell wall, so that other immune cells can learn the invaders' signature. Another instance involves an infected or cancerous internal cell that displays unusual, nonself antigens on its outer surface via the MHC. Such nonself-antigens stimulate an immune response against the cancerous cell, whereas the display of self-antigens elicits no such response.

Nonself is key to detecting and responding to malicious computational events. A cybersecurity analogy is the use of an encrypted certificate or security label for all approved hardware, software, data, and users. Indeed, global-scale identity management is needed to deny access by anonymous outsiders to sensitive data and to hold malicious insiders accountable for their actions.

Another approach, *dynamic program analysis*, reverse-engineers suspected malware into functional code fragments and searches for patterns identifying typical malware behavior. This approach uses behavior patterns to identify and thwart obfuscation techniques such as polymorphism or virtualization. (The Concordia architecture uses this approach.¹¹) MHC-like signatures of new attacks can then be quickly cataloged and distributed, providing a new architecture for automating the generalization of program structures and recognition of common patterns for malware analysis. Such a "Google for malware,"¹¹ combined with data provenance, would also provide benefits for attribution and situation awareness.

Conscious Decisions

Humans make conscious decisions that let them avoid dangerous situations and identify people by

intrinsic features such as the face, body features, mannerisms, the voice, body language, and specific knowledge, as well as extrinsic identifiers such as a badge or smart card.

Similarly, authentication mechanisms enable decisions by using something that the user

- knows, such as a user ID or static password;
- inherently has, such as one or more biometrics; or
- possesses, such as a token, smart card, or time-based password.

Authentication should also include hardware, software, and data. Another analogy is the tracking or profiling of users' behavior for the purposes of deterrence, access, and forensic accountability of insiders.

Apoptosis

Apoptosis is programmed cell death to halt the spread of virus-infected cells and to halt a nonfunctional cell's use of resources. It removes cells that are damaged beyond repair, implying that cancer arises at least partly as a result of immune dysfunction. Apoptosis can be initiated by the cell itself, the tissues surrounding the cell, or the immune system. Typically, 50×10^9 to 70×10^9 cells (out of approximately 10^{14} total) die daily (approximately 0.06 percent per day) in a human adult.

So, apoptosis handles all combinations of good and trustworthy versus corrupted or malicious cells, which are analogous to cybernodes, scaling from user to computer to network. A more specific cyber analogy is the termination of network access for any node that displays unauthorized activity or violates the security policy.

Beyond the Human Analogy

Scalable trustworthy systems needn't rely only on an understanding of human physiology. A natural example involves the Komodo dragon's saliva, which contains *Pasteurella multocida*, a virulent strain of bacteria that quickly causes sepsis and death from a single bite. A component of the Komodo dragon's blood neutralizes these bacteria.¹² Other recent research shows that proteins in alligators' white blood cells have antibiotic properties, which protect the animal from fungi, yeast, and bacteria, even if the animal has had no previous exposure to these organisms. An understanding of such immune responses will likely be useful for inspiring cybersecurity research for years to come.

Next Steps

The vision we've discussed here is inherently long term, multidisciplinary, and certainly on the order of a grand challenge. Scalable trustworthy systems involve

needs beyond computer science and high-performance computing, including management of complexity at all scales, analysis of exabytes of data in near real time, and protection of existing infrastructure that's undergoing increasingly sophisticated attacks.

Requirements

These needs entail both functional and nonfunctional requirements. For example, the Furps+ approach analyzes software functionality, usability, reliability, performance, and supportability, plus design, implementation, interfaces, and physical constraints. It then captures requirements using specific, quantifiable metrics for testing, inspection, or analysis, to understand how well we're doing to enable continued, more effective improvements for both functional and nonfunctional aspects. (Examples of nonfunctional aspects include security properties that might be inspired by nature—for instance, the BBB, MHC, or CNS.)

Meeting Specific Needs

Use of the HPI metaphor suggests solutions for specific needs. For example, solutions could be based on a systematic understanding of the immune function for each human cell type as a basic component of bodily functions, and the body's immune defense systems in particular. Such an understanding involves

- characterization of each cell type's specific, quantifiable functions;
- hierarchical organization of cells into tissues, organs, organ systems, and the whole body; and
- identification of the diverse, distributed functions underlying this hierarchical organization that collectively create robust immunity through real-time, knowledge-based decisions.

The human body manages complexity by a rich synergy among hardware and software, specific functions for each cell type, hierarchical architecture, massive redundancy, and multiple feed-forward and feedback loops for signaling and control. We need to employ this strategy to create breakthrough cybersecurity approaches that ensure a chain of trust for only healthy functions and signals to eliminate whole classes of vulnerabilities.

Other solutions could be based on the abstraction of physiological functions as predictably composable components (for example, interoperable, provably secure, reduced-instruction-set code primitives). Such solutions would use cyber analogies to cell-based functions that

- avoid, detect, and eradicate attackers;
- recognize and thwart malicious users (for example,

- analogous to “spontaneous” remission of cancer);
- detect and heal underlying damage;
- restore normal functions; and
- prepare for efficient resolution of future attacks of a given type.

Implementing predictably composable components in the underlying hardware is another challenge to ensuring healthy functions, including

- platform independence;
- the ability to thwart all known and 0-day attacks, while avoiding PoP; and
- scalability across the infrastructure, in areas such as computers, sensors, embedded processors, routers, repeaters, firewalls, hubs, and instruments.

Modern software engineering has made substantial progress in writing secure code via structured or formal planning and methods, implementation, and testing. This long-term view would naturally be much more cost effective than PoP.

Some Goals

Computer attacks against the Pentagon currently average 5,000 each day. As the National Cyber Leap Year Summit cochairs stated, we need R&D that translates biological immunity to digital immunity “to automatically detect situational changes, determine imminent danger and mitigate cyberattacks.”⁴ Here are some examples of their suggested research goals:⁴

- “Thwart malicious attacks through signaling, implementation of diversity and immunogenic detection as hardware-software solutions. Rapidly regenerate (self-healing) survivable capabilities in mission critical systems after a sophisticated attack.”
- “Evolve immunity to attacks through evolutionary computing to create new deceptions (gaming strategies) as new threats emerge. [Implement] self-learning while monitoring insider activity and develop profiles for appropriate and legitimate behavior (modeling).”
- “[Integrate] the many disparate security tools using both feed forward and feedback signaling mechanisms in a cyber defense system ... to ensure tolerance and identify attacks while minimizing false alarms.”
- “[Amalgamate immunologically inspired] distributed control mechanisms for learning, memory and associative retrieval to solve recognition and classification tasks. ... [The body handles] antigenic challenges through collaborative interaction. ... [Pursue a] similar strategy (distributed control mechanisms for monitor and response) ... to avoid a single point of failure and to enable robust decision making.”

Grand-challenge-class R&D is needed to address these long-standing and increasingly severe issues.

Using the HPI metaphor, we can address the hard problems discussed in the section “Cybersecurity—a Very Hard Problem.” How does the brain make knowledge-based decisions about trust? How does the brain do real-time processing of data into information and knowledge for these decisions? How does the brain manage the inherent complexity of this data-into-knowledge transformation across 10^{10} nodes (neurons)? How do the brain and immune system avoid cascading failures in the midst of ongoing attacks? Insights from these questions will undoubtedly be useful in developing far-reaching strategies to secure cyberspace and better deal with the hard problems. These insights will enable society to reduce the risk to highly critical systems and infrastructure, thwart the sophisticated, rapidly growing threat, and address other priorities such as the untold \$100+ billions of losses to cybercrime.

However, the HPI metaphor might not always scale to the fast-changing, ever-more-sophisticated arms race. Indeed, our vision is both necessarily and purposefully general and high level. This is because the challenge is immense, in terms of both reverse-engineering the brain (to include the various naturally evolved human defenses) and adapting that knowledge to achieving scalable trustworthy computing. Some thought-provoking questions remain about our metaphor’s suitability. In particular, the cybersystems’ increasing complexity will eventually surpass the scale of the human system. So, at some point, HPI systems might simply fail to scale to cybersecurity problems. However, it’s unclear when—if ever—we’ll reach this point. If we do, by then we’ll certainly have established a “Cyber Center for Disease Control.” □

Acknowledgments

UT-Battelle LLC manages Oak Ridge National Laboratory for the US Department of Energy, under Contract DE-AC05-00OR22275.

References

1. President’s Information Technology Advisory Committee, *Cyber Security: A Crisis of Prioritization*, US Nat’l Coordination Office for Information Technology Research and Development, Feb. 2005; www.nitrd.gov/pitac/reports/20050301_cybersecurity/cybersecurity.pdf.
2. Inter-agency Working Group on Cyber Security and Information Assurance, *Federal Plan for Cyber Security and Information Assurance Research and Development*, US Nat’l Science and Technology Council, Apr. 2006; www.nitrd.gov/pubs/csia/csia_federal_plan.pdf.
3. *Roadmap for Cybersecurity Research*, US Dept. of Homeland Security, Jan. 2009; www.cyber.st.dhs.gov/docs/DHS-Cybersecurity-Roadmap.pdf.
4. *National Cyber Leap Year Summit 2009 Co-chairs’ Report*, Networking and Information Technology Research and Development (NITRD), 16 Sept. 2009; www.qinetiq-na.com/Collateral/Documents/English-US/InTheNews_docs/National_Cyber_Leap_Year_Summit_2009_Co-Chairs_Report.pdf.
5. Michael Näf, “Ubiquitous Insecurity? How to ‘Hack’ IT Systems,” *Information & Security*, vol. 7, 2001, pp. 104–118.
6. “The Digital Universe Is Still Growing,” EMC Corp., 2011; www.emc.com/leadership/digital-universe/expanding-digital-universe.htm.
7. G.V. Hulme, “If Stuxnet Was Act of Cyberwar, Is U.S. Ready for a Response?” *Computerworld*, 3 Mar. 2011.
8. J. Hawkins and S. Blakeslee, *On Intelligence*, Henry Holt, 2004.
9. R.Q. Quiroga et al., “Invariant Visual Representation by Single Neurons in the Human Brain,” *Nature*, 23 June 2005, pp. 1102–1107.
10. D. Dasgupta and F. Nino, *Immunological Computation: Theory and Application*, CRC Press, 2008.
11. T. Daly and L. Burns, “Concordia: Google for Malware,” *Proc. 6th Ann. Workshop Cyber Security and Information Intelligence Research*, ACM Press, 2011, article 30; http://daly.axiom-developer.org/TimothyDaly_files/publications/sei/CSIIRW10/concordiaSheldon.pdf.
12. “Gator Blood Destroys Deadly Superbugs,” *Scientific Computing*, 2011; www.scimag.com/Gator_Blood_Destroyes_Deadly_Superbugs.aspx?terms=alligator.

Lee Hively is a senior researcher at Oak Ridge National Laboratory. He has a PhD in nuclear engineering from the University of Illinois at Urbana-Champaign. Contact him at hivelylm@ornl.gov.

Frederick Sheldon is a senior research scientist at Oak Ridge National Laboratory. His research has dealt with developing and validating models, applications, methods, and tools for creating safe, secure, and dependable systems. Sheldon has a PhD in computer science from the University of Texas at Arlington. He’s a senior member of IEEE. He has received a Sigma Xi outstanding dissertation award and a key contributor and significant event award from UT-Battelle for excellence in technology transfer. Contact him at sheldon@ieee.org.

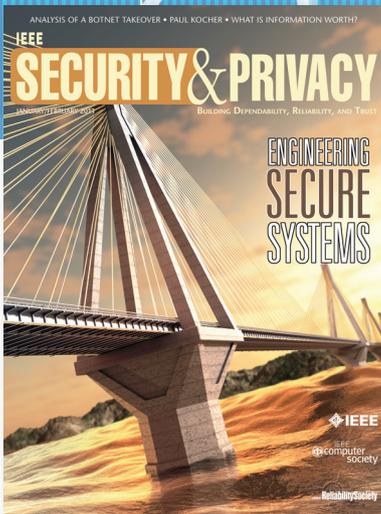
Anna Cinzia Squicciarini is an assistant professor at Pennsylvania State University’s College of Information Sciences and Technology. Her main research interests include trust negotiations, security in cloud computing, privacy, and access control for grid computing systems. Squicciarini has a PhD in computer science from the University of Milan. She’s a member of IEEE. Contact her at acs20@psu.edu.

IEEE

SECURITY & PRIVACY

\$19⁹⁵

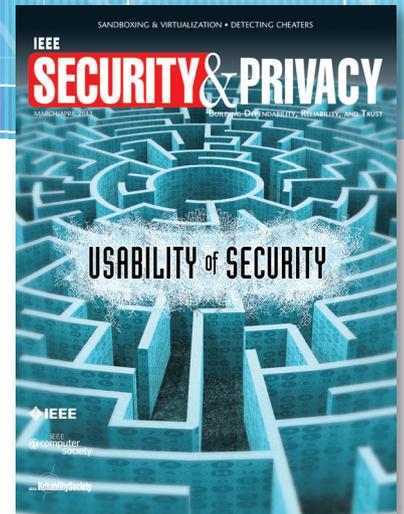
Subscribe to *IEEE Security & Privacy*!



Protect Your Network

Stay Ahead of the Competition

Prevent Attacks



IEEE Security & Privacy is the publication of choice for great security ideas that you can put into practice immediately. No vendor nonsense, just real science made practical.



—Gary McGraw,

CTO, Cigital, and author of

Software Security and Exploiting Software

www.qmags.com/SNP
for your digital subscription