A Vision for Scalable Trustworthy Computing

L.M. Hively[†], F.T. Sheldon[†] and A. Squicciarini[*]
[†]Oak Ridge National Laboratory, Oak Ridge, TN 37831-6418
[*]Pennsylvania State University, University Park, PA

Keywords: Cybersecurity, Trustworthy Computing, Scalability

Abstract
The cybersecurity landscape consists of an *ad hoc* patchwork of solutions [1]. Optimal cybersecurity is considered "hard," for various reasons: complexity, immense data and processing requirements, resource-agnostic cloud computing, practical time-space-energy constraints, inherent flaws in "Maginot Line" defenses as well as the growing number and sophistication of attacks. We begin by defining and abstracting the high priority problems including a crosswalk of the potential and co-opted solution space. Within that space, we claim that achieving scalable trustworthy computing and communications is possible via real-time knowledge-based decisions about cyber trust. Our vision is based on the human-physiology-immunity (HPI) metaphor and the human brain's ability to extract knowledge from data and information. We outline some future steps toward scalable trustworthy systems requiring a long-term commitment to solve the well-known "hard problems."

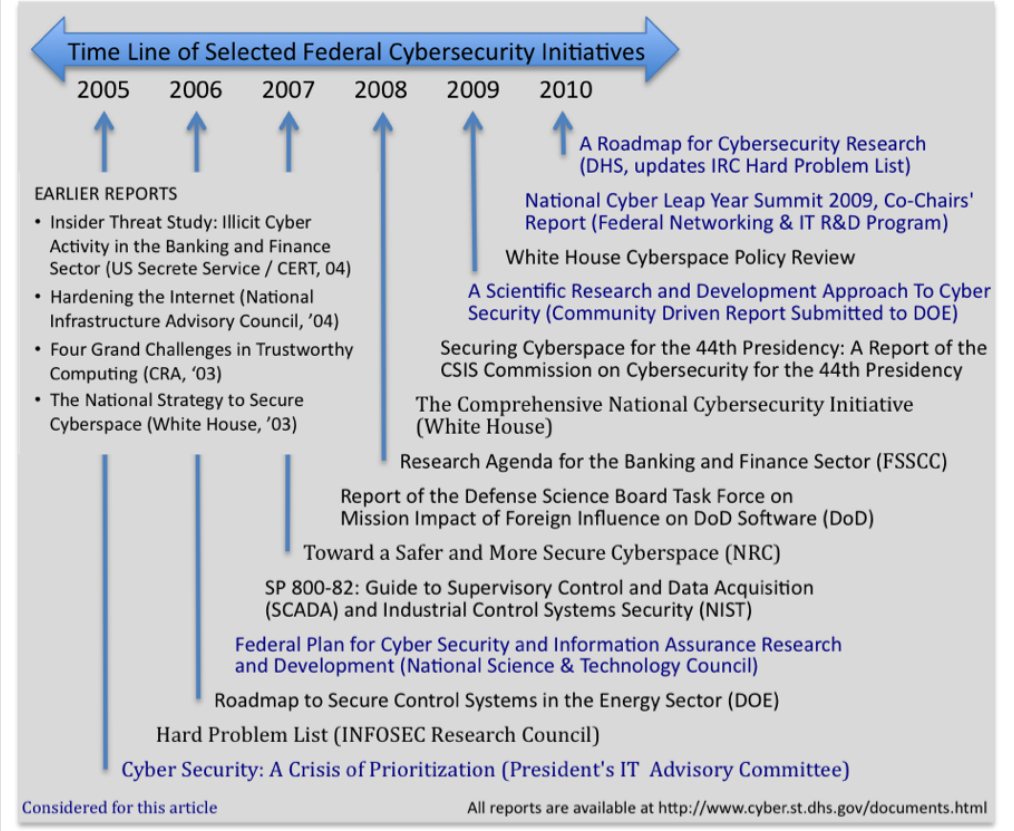

1. Introduction

Recent Federal policy documents have emphasized the importance of cybersecurity to society's welfare (Figure 1). One example is the President's *National Strategy to Secure Cyber Space* (2003), which describes *national response priorities*. *Cyber Security: A Crisis of Prioritization*[4] describes ten technologies needed for cybersecurity. *Federal Plan for Cyber Security and Information Assurance Research and Development*[3] developed 49 cyber security technical topics in 8 major R&D areas with corresponding funding priorities. DHS Roadmap for Cybersecurity Research[2] developed eleven "hard problems" (eight from the 2005 IRC Hard Problem List). National Cyber-Leap-Year (NCLY) Summit [5] discussed 5 cross-cutting solution themes. Table 1 maps the problem space to the solution space via analysis of the R&D priority documents. Indeed, a "leap forward" in cybersecurity is needed through "game changing" technologies [Aneesh Chopra, U.S. CTO]. Unfortunately, the realities are founded in a *Maginot Line* mindset, which have failed to prevent cyber crime/fraud losses estimated to have exceeded $1000B in 2008.

Table 1. Crosswalk of recent federal cybersecurity priorities.

Federal Problems Characterization Efforts (selected)			Solution Themes (†)
PITAC 2005 Cyber Security Priorities [4]	NSTC 2006 *Some* of the Top Cybersecurity / IA R&D Priorities [3]	DHS 2009 Roadmap for Cybersecurity Research (Hard Problem List V2) [2]	NITRD 2009 National Cyber Leap Year Summit [5]
P1 –Authentication (3)	N1 –Authentication, authorization, trust management, access control, privilege management (4)	D1 –Scalable trustworthy systems (including sys. arch. & requisite dev. methodology) (4)	(1) Hardware-enabled trust [knowing when you've been had]
P2 –Secure software engineering (2)	N2 –Large-Scale cyber situational awareness, automated attack detection, warning, response (3)	D2 –Enterprise level security metrics (measures of overall system trustworthiness) (3)	
P3 –Holistic system security (2)	N3 –Insider threat detection, mitigation, forensics, traceback, attribution (4)	D3 –System evaluation life cycle (including approaches for sufficient assurance) (2)	(2) Cyber economics [crime/ fraud does not pay]
P4 –Monitoring/detection (3)	N4 –Secure DNS/routing, protocols/process control systems (3)	D4 –Combating insider threat (3)	
P5 –Secure fundamental protocols (2)	N5 –Domain-Specific Security (e.g., wireless, RFID) (2)	D5 –Combating malware, botnets (3)	(3) Moving-target defense [attacks work once if at all]
P6 –Mitigation and recovery (1)	N6 –Detection of vulnerabilities, malicious code; metrics/software testing/ assessment (3)	D6 –Global-scale identity management (3)	
P7 –Cyber forensics (3)	N7 –Secure OS, software engineering, information provenance (3)	D7 –Survivability of time-critical systems (4)	(4) Digital provenance [basing trust decisions on verified assertions]
P8 –Modeling/testbeds (3)	N8 –Cybersecurity, IA R&D testbeds, IT systems, Internet modeling, simulation, visualization (3)	D8 –Situational understanding, attack attribution (2)	
P9 –Metrics, benchmarks, best practices (3)	N9 –Trusted computing Base architectures, composable, scalable, secure systems (3)	D9 –Provenance (relating to information, systems & h/w) (4)	(5) Nature-inspired cyber health [move from forensics to real-time diagnosis]
P10 –Non-technology issues (2)	N10 –Inherently secure, high-assurance, provably secure systems, architectures (3)	D10 –Privacy-aware security (3)	
	N11 –Trust in the Internet, privacy (3)	D11 –Usable security (3)	

† Progress in this solution theme area will support advances in the other problem areas listed; (#) frequency a priority appears in column four.

Innovative solutions are clearly needed. A long-term vision for *scalable trustworthy* systems requires solutions for all the problems listed in Table 1. *Trustworthiness* is a multidimensional measure of system requirements for usability, integrity, availability, survivability, confidentiality, performance, accountability, attribution, and other critical needs. Precise requirements for trustworthiness and corresponding measures are fundamental precursors to developing and operating trustworthy systems [2]. *Scalability* is the ability to satisfy requirements as systems, and systems of systems expand in functionality, capacity, complexity, and scope [2]. *Composability* is the ability to create systems and applications with predictably satisfactory behavior. To enhance scalability, high assurance systems should be developed from a set of composable components and subsystems, each of which is itself suitably trustworthy, within a system architecture that inherently supports facile composability [2].

Present architectures (hardware, operating systems, networks, applications) do not satisfy these combined requirements adequately. Scalable trustworthy systems should be composed of "suitably" trustworthy components, down to the most basic level, thus avoiding development of new methodologies at each successively larger scale. Moreover, scalability should enhance trustworthiness (e.g., constructive system design, meticulous use of best practices, error-correcting code to overcome unreliable communications and storage, encryption to protect integrity and confidentiality of insecure communications). Such techniques are incomplete, if they rely on the trustworthiness of developers, users, and administrators. The challenges are then: a) a sound basis for composability that scales to large, complex, trustworthy systems; b) trustworthiness evaluations of composite systems that are themselves composable and scalable; and c) development of components, analysis tools, metrics and testbeds for a+b.

Next, we consider the features of modern technology and their inherent threats. We discuss why cybersecurity is so hard. We propose a compelling basis for the human body's management of complexity, nonlinearity, and the immune response as a metaphor for scalable trustworthy systems. We give examples of human-body functions that hint at achieving scalable, trustworthy solutions. We outline a path forward to enable this new paradigm. Finally, we summarize our ideas and present our conclusions: *there is growing interest in human immune-inspired functions for construction of information systems*.

2. The current scenario: computational performance vs. ubiquitous insecurity

An important challenge for cybersecurity is keeping pace with the evolution of modern systems. Figure 2 depicts this evolution, as the speed of modern computers has increased by more than 10^4 over the last fifteen years. Performance is now limited by parallelization and energy consumption, rather than individual processor speed.

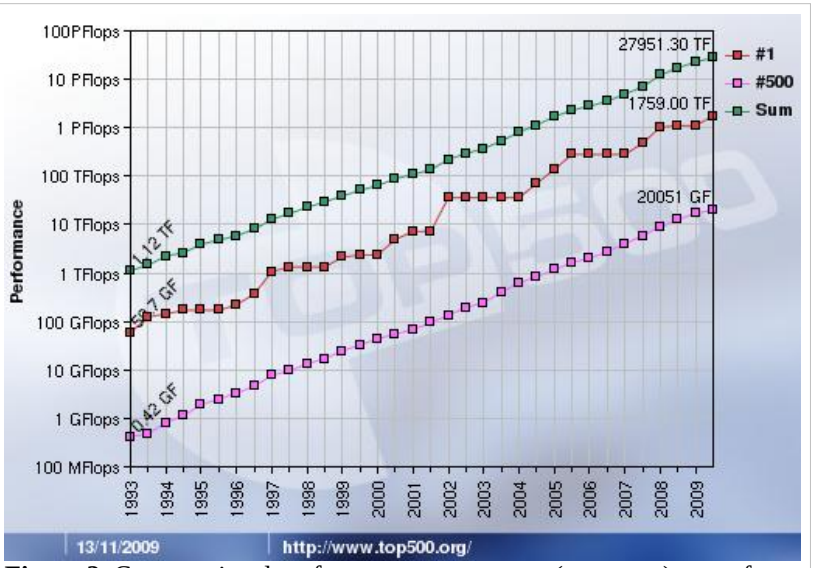

Figure 2. Computational performance versus year: (top green) sum of world's top 500 computers; (middle red) fastest computer; (bottom purple) slowest of the top 500 computers.

Computational improvements have been accompanied by ubiquitous insecurity in the cyber realm[6]. Malicious software (malware) is frequently used in attacks via the Internet (Figure 3), involving deliberate infiltration or damage to a computer system without the owner's informed consent. Attacks range from "low-and-slow" over a day (or more) to "fast-and-focused" at the millisecond level or faster. Such attacks are hidden in a sea of normal cyber activity. The threat posed by internal attacks can result in devastating consequences[7], including elevated privileges for malware that is directed by external agents.. Insider threat has many forms, including information extrusion, neglect (failure to follow policy or best practices), indifference or ignorance, and maliciousness (e.g., disgruntled employees and spies [exfiltration]). The greatest challenge is the continuous attack evolution. Previous solutions for known threats may not address the new attacks, which are hard to predict in terms of effectiveness and disruption. Traditional risk methodologies provide common-sense advice, but usually lack specific guidelines for the evaluation of emerging threats. Hence, better protection from future threats is needed at all sensitivity levels. Cybersecurity is a multi-faceted, hard problem for several reasons, as discussed next.

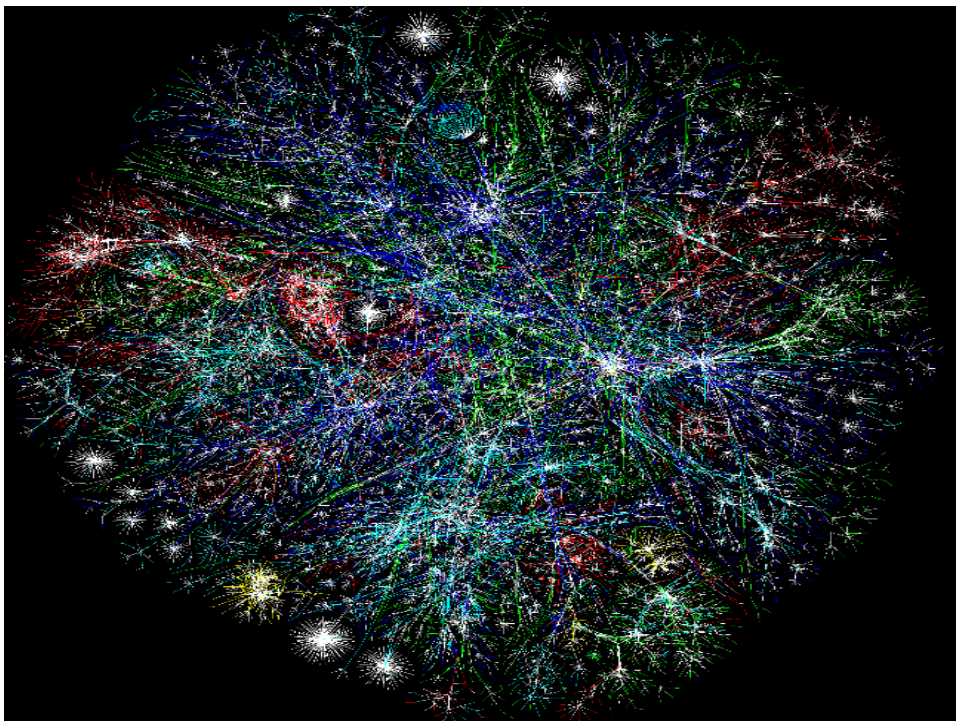

Figure 3: 1-day Internet map (23 Nov. 2008): Red: Asia/Pacifica; Green, Europe/Middle East/Central Asia/Africa, Blue: N. America; Yellow: Latin American and Caribbean, Cyan: Private Networks; White: unknown (www.opte.org/maps/tests/).

3. Cybersecurity is a Very Hard Problem

Complexity at all levels is one feature that makes cybersecurity hard. Figure 3 illustrates that the Internet is a very complex and seemingly scale-free[14]. All modern computers are themselves networks of systems (e.g., CPUs, memory, GPUs, storage, data busses, I/O devices, etc.). All modern software is a complex network of processing functions. The information infrastructure is a complex system of systems of hardware, software, operating systems, data, networks, and people. Complex interactions frequently produce

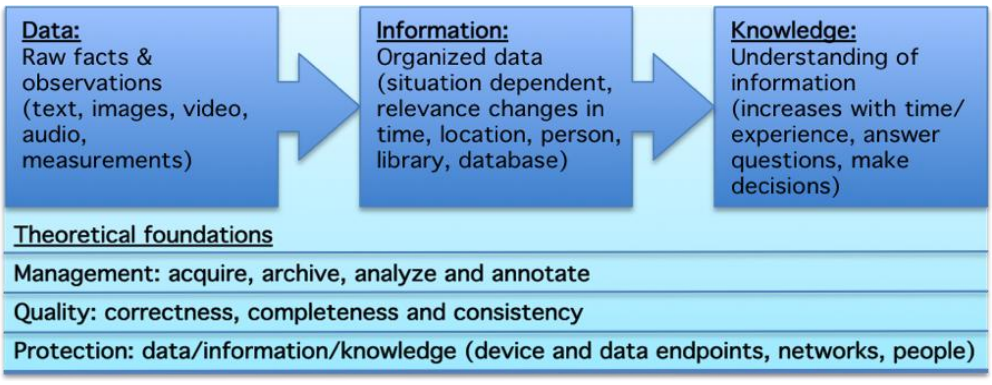

Figure 4: Conversion of raw data into information (data in the context of other data), hence into knowledge (information in the context of other information) for understanding and prediction.

emergent, unexpected, and potentially adverse behavior. Indeed, failure in such an infrastructure can be so complex that no one can presently determine the cause, let alone a cure. Scalable trustworthy systems must cope with complexity.

The amount of global data is immense, involving 451 Exabyte's (4.51×10^{21} bytes) or 72 GB for each person on Earth[8]. "Data" are all electronic forms of data, information, and knowledge. Scalable trustworthy systems must be able to process more of this tsunami of data in (near) real-time for attack characterization, situational understanding, attribution, and appropriate response.

Cybersecurity decisions require the conversion of data into information and hence into knowledge. Data analysis in the context of other data generates information, processed in the context of other information to create knowledge (Figure 4). Current systems cannot create knowledge, but rather rely on decisions by humans who cannot respond at computer speed (milliseconds or less). Moreover, a human cannot detect sparse anomalies in the knowledge discovery process. Accordingly, we view the "best practice" of the software-patch cycle (e.g., "Patch Tuesday"), and subsequently reverse-engineered exploits (e.g., "exploit Wednesday") as an ongoing admission of failure. Robust cybersecurity requires a new paradigm. Scalable trustworthy systems must process the tsunami of data in (near) real-time for knowledge-based decisions about cyber trust.

Cybersecurity has practical constraints, which include:

- Protection of private information (essential for public acceptance)
- Handling of imperfect data appropriately (errors, incompleteness, inconsistency, noise)
- 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 be inter-operable with legacy systems
  - Comply with mandates of law at all levels
  - Provide for graceful degradation of safe operation during failure
  - Minimally impact the users' ability to do real work
- Facilitation of open-source software use, parallelism, debugging, and software quality assurance
- Enabling multi-language development for multiple applications

Scalable trustworthy systems must interoperate with legacy systems within constraints that are reasonable and within the context outlined above.

A further challenge is the inadequacy of perimeter defenses in our networked world. Traditional approaches focus on a "layered defense," or "defense in depth" to protect the "crown jewels" by physical or cyber walls and fortifications that form "air gaps" between the layers.[1]. This "Maginot Line" approach [9] cannot protect the "inside" from the "outsiders," which are inherently connected in a networked world. 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 does not fortify the network, just as the fortification along the Maginot Line was insufficient during the World War II blitzkrieg. 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 are very difficult to hack;
- Systems that determine trustworthiness of hardware, software, network, and users (e.g., cataloging).

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

Cyber attacks are growing in number and sophistication. Recent examples include:

- Organized nation-state attacks against the Pentagon and other facilities in the US;
- Organized nation-state attacks on the countries of Estonia and Georgia;

- 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);
- Use of botnets (and other organized Internet exploits);
- Website and web application exploits;
- Compromise of unsecured data.

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 (malicious or honest mistakes) in complex systems are very difficult to detect, analyze, and correct. Thus, all modern complex systems have vulnerabilities. Updates compound this complexity. Ubiquitous networking opens a vulnerable 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. This logic concludes that the root cause of vulnerabilities is the always-imperfect software (and hardware and networks) that can never be totally secure. *We refute this argument via examples of complex living organisms that manage complexity and provide secure, real-time responses.*

4. Compelling Reasons that Scalable Trustworthy Systems Are Possible

One compelling motivation is the human brain's 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 (≤1 second) via neuron-based processing with a single-neuron cycle time of ~10 milliseconds, corresponding ~100 neural hops (10^2 processing cycles) per second over ~10^{11} brain neurons, for a net processing power of ~10^{13} cycles/second. Modern high-performance computers run at >10^{15} operations/second, or >10^2 more processing power than the brain, yet cannot perform "intelligent" real-time processing of the same data. Accordingly, we view the human brain's capacity for intelligent, real-time, knowledge-based decisions as a basis for envisioning scalable/secure situation awareness.

Jeffrey Hawkins' 2004 book, On Intelligence, focuses on the brain's neocortex, which has ~10^{10} neurons, and ~10^{14} connections. The key features are (1) an *irreducible representation* for each item in memory; (2) *auto-associativity* among items (e.g., recalling one line of a song leads to the remainder), because "a memory" is recall of a time-serial sequence of stored items; (3) *hierarchical* processing, (for example) combining the simplest spoken sounds [phonemes] into words, which then are combined into phrases that form sentences, and concepts; and (4) *feed-forward* links to make appropriate connections among phonemes, words, phrases, sentences, and concepts in the context of previous knowledge. There is also *feed-back* 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 any accent, situational context, etc. Likewise, image processing extracts (for example) points, lines, polygons, object identification, scene familiarity, and scene changes. Indeed, the same neocortical processing paradigm extracts a hierarchical sequence of patterns for all time-serial sensory data (auditory, somatosensory, etc.). Blind people can be trained to read brail by their fingers, to "see" crude images via discrete touch points on the tongue, or to "see" sound-scape images via stereo headphones. "Understanding" is the essence of intelligence, as is the ability to predict a new situation correctly on the basis of previous knowledge. This hierarchical, brain-based paradigm is very different from the present program counter (PC) based programming paradigm, and may provide insight for the data-to-information-to-knowledge processing paradigm of Figure 4.

This metaphor assumes that the cell (e.g., neuron) is the basic unit for information processing. This assumption stems from work [15] by Quiroga *et al*. These experiments recorded the response of single human-brain neurons, showing 44 (out of 137) that were selective to a unique object (e.g., picture of

Jennifer Aniston). This response occurred for different views of the same object (e.g., front versus side view). These observations are consistent with Hawkins' irreducible (invariant) memory representation.

Healthy humans can live for 70+ years, while thwarting continuous attacks from diverse microbes, toxins, and health-endangering conditions. Each cell can be viewed as an information processor that receives input, processes 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 ~10^{14} cells in a healthy adult. Indeed, all body systems participate in immune function (Table 2). Complex, adaptive human behavior arises from interactions among the tightly integrated, hierarchical components, which are composed of massively parallel, cellular processors. Knowledge-based decisions cannot process arbitrary instructions, and are therefore not hackable. These examples may provide insight for scalable trustworthy computing via an integrated, active, distributed, hierarchical hardware-software composition (as discussed above) with proper design, implementation, and "hygiene." Perhaps, inherently scalable trustworthy systems are those with an architecture for only "healthy" functions, rather than the patches-on-patches (PoP) approach to preventing (further) attacks (e.g., do we even know if, in the long run, if PoP can actually result in a smaller attack surface?).

Table 2: Examples of Immune Functions by the Human Body (adapted: Table 22.2, Ref. 10)

Body system	Brief description of specific immune functions
1 Circulatory	Blood-distributed immune cells throughout the body; recovery of immune cells via lymphatic flow
2 Digestive	Continuous salival cleansing of mouth via lysozyme; pathogen destruction by HCl in stomach
3 Endocrine	T-lymphocyte *programming* messages via thymus hormones; depression of immune activity via stress
4 Immune	Capture/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, promote repair; fever response by pyrogens to enhance repair and inhibit pathogens; apoptosis; major histocompatability complex
5 Muscular	Movement to avoid or protect from pain, heat, danger
6 Nervous	Fight-or-flight response; avoid unhealthy actions (e.g., smoking), pursue healthy habits (e.g., exercise); enhancement/inhabitation of immune functions via serotonin, (nor)epinephrine; blood-brain barrier
7 Reproductive	Inhibition of bacterial/fungal growth by acidic mantle of vagina
8 Respiratory	Physical barrier/entrapment of microorganisms by mucous (larynx, pharynx, nasal cavity); removal of debris-laden mucous from lower tract by cilia; filtering/entrapment of microorganisms by nasal hairs
9 Senses	Cerumen and hairs as external barriers in the ear; foul tastes to prevent eating unhealthy food; continuous eye cleansing by tears with lysozyme
10 Skeletal	Production of blood (immune) cells in bone marrow
11 Skin	Mechanical barrier against entry of pathogens/toxins; perspiration as bacterial growth inhibitor
12 Urinary	Acidic pH of urine as bacterial inhibitor; cleansing of lower urinary tract with each voiding; bactericidal chemical in sebum; resistance against acids, alkalis, and bacterial enzymes in keratin

5. Analogies of Immune Function

The compelling similarities between cybersecurity and biological systems have sparked research into specific applications (e.g., immuno-computing and artificial immune systems in the 1990s). From an information-processing perspective, there are several immunological principles that make the analogy appealing, including distributed processing, pathogenic pattern recognition, multi-layered protection, decentralized control, as well as diversity and signaling. We next consider relevant analogies for potential scalable trustworthy solutions, as an extension of present research [12][13]. Each of the examples addresses one (or more) of the above-mentioned problems.

One immune example (Table 2: row 6: nervous system) is the blood-brain barrier (BBB), which 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 central processor unit (CPU) from the rest of the cyber world via a fast, in-line encryptor/decryptor chip (EDC). The BBB effectively protects the brain from infections by using carrier-mediated transporters (e.g., glucose) to ferry low-atomic-weight substances (≤500 daltons) into and out of the CNS. A cyber analog is short, encrypted packets via single-

use keys. Strict physical isolation of the CPU could include a processor-resident operating system on encrypted read-only-memory that is distinct from applications. Tamper resistance in the CPU/EDC (not unlike the brain inside the skull, though distinct from the BBB) could shut down an always-on processor on tamper detection, thus erasing the operating system and any sensitive data. Answers to other questions (e.g., how the brain self-heals and restores lost memory) will certainly enable deeper understanding of the intelligence and cybersecurity.

The major histocompatability complex [10] (MHC) distinguishes self from non-self (Table 2: row 4: immune system). For instance, a blood-born immune cell (e.g., leukocyte) encounters a foreign invader, engulfs and destroys it, and displays random fragments (antigens) on MHC molecules that are attached to its outer cell wall, so that other immune cells can learn the signature. Another case involves an internal cell that is infected or cancerous, and displays unusual, non-self antigens on its outer surface via MHC. Such non-self antigens stimulate an immune response against the cancerous attack, while the display of self-antigens elicits no such response. Non-self is key to detecting and responding to malicious computational events. A cyber analog 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, known as dynamic program analysis, reverse-engineers suspected malware into functional code fragments and searches pattern-identifiers of typical malware behavior (i.e., Concordia: Google for Malware), thus thwarting obfuscation techniques (polymorphism/virtualization). MHC-like signatures of new attacks can then be quickly catalogued and distributed providing a new architecture for automating generalization of program structures and recognition of common patterns for malware analysis. A Google for malware combined with data provenance would also provide benefits for attribution and situation awareness.

Conscious decisions (Table 2: row 6, nervous system) allow us to avoid dangerous situations and to identify people (e.g., intrinsic face and other body features, mannerisms, voice, body language, 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: (1) knows (e.g., user ID, static password), (2) is (e.g., one or more biometrics), and has (e.g., token, smartcard, time-based password). Authentication should also include hardware, software, and data. Another analogy is the way users' behavior is tracked (or profiled) for the purpose of deterrence, access and forensic accountability of insiders.

Apoptosis is another immune example (Table 2), which manifests as a form of programmed cell death to halt the spread of virus-infected cells and to halt the use of resources by a non-functional cell. Apoptosis removes cells that are damaged beyond repair, implying that cancer arises at least in part as a result of immune dysfunction. Apoptosis can be initiated by the cell itself, by surrounding tissue, or by the immune system. Typically, 50 to 70 x 10^9 cells (out of ~10^{14} total) die daily (~0.06%/day) in a human adult. This approach handles all combinations of good/trustworthy versus corrupted/malicious cells, which are analogous to cyber nodes (e.g., user, computer, network). A more specific cyber analogy is the termination of network access for any node that displays unauthorized activity or violates security policy.

Scalable trustworthy systems need not only rely on an understanding of human physiology. A natural example involves the Komodo dragon's saliva, which has a very virulent strain of bacteria (*Pasteurella multocida*) that quickly causes sepsis and death from a single bite. A component of the Komodo dragon's blood neutralizes these bacteria [11]. Other recent research shows that proteins in the white blood cells (leucocytes) of alligator blood have antibiotic properties against fungi, yeast, and bacteria (including antibiotic-resistant *Staphylococcus aureus*) without having previous exposure to them. An understanding of such immune responses will likely be useful for inspiring cybersecurity research for years to come.

6. Next steps

Solutions must address the growing list of "hard problems," some of which are discussed here. The vision discussed here is inherently long-term, multi-disciplinary and certainly grand-challenge class. 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 while under increasingly sophisticated attack. These needs involve both functional and non-function requirements. For example, the FURPS+ approach uses functionality, usability, reliability, performance, and supportability, plus design, implementation, interface(s), and physical constraints. Requirements are then captured by specific, quantifiable metrics for testing, inspection, or analysis to understand how well we are doing to enable continued and more effective improvements.

Use of the human-physiology-immunology (HPI) metaphor suggests solutions for specific needs. One novel feature is a *systematic understanding of the immune function for each human cell type* as basic components of bodily functions, and immunity in particular. Implementation of this approach involves: (1) characterization of the specific, quantifiable function(s) of each cell type; (2) hierarchical organization of cells into tissues, organs, organ systems, and the whole body; and (3) identification of the underlying, diverse, and distributed functions from (2) that collectively create robust immunity via 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 feed-back loops for signaling, and control. This same approach is needed for game-changing approaches that ensure a chain of trust for only "healthy" functions/signals to eliminate whole classes of vulnerabilities.

Another novel feature is abstraction of physiologic functions as *predictably composable components* (e.g., interoperable, provably secure, reduced-instruction-set code primitives). This feature uses cyber analogs to cell-based functions that: (1) avoid, detect, and eradicate attackers; (2) recognize and thwart malicious users (e.g., analogous to "spontaneous" remission of cancer); (3) detect and heal underlying damage; (4) restore normal functions; and (5) prepare for efficient resolution of future attacks of a given type.

Implementation of *predictably* composable components in the underlying hardware is another challenge to assure healthy functions, including: (1) platform-independence; (2) ability to thwart all known and zero-day attacks, while avoiding the present (failed) PoP approach; and (3) scalability across the infrastructure (e.g., computers, sensors, embedded processors, routers, repeaters, firewalls, hubs, instruments). Indeed, modern software engineering has made substantial progress in revealing the "secrets" of writing secure code via structured/formal planning/methods, implementation and testing. This long-view would naturally be much more cost effective than the present PoP approach.

Dennis Blair (Director of the Office of National Intelligence) advised Congress in February 2010 that "malicious cyber-activity is growing at an unprecedented rate" and that efforts to defend against cyber-attacks "are not strong enough." An "explosion" of computer attacks against the Pentagon, currently averages 5,000/day. R&D is needed to translate from biological to digital immunity for automatically detecting situational changes, determining imminent danger and mitigating cyber attacks, for example [5]:

- Thwart malicious cyber-activity through signaling, implementation of diversity and immunogenic detection as combined hardware-software solutions. Rapidly regenerate (self-heal) 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. Self-learning while monitoring insider activity and develop profiles for appropriate and legitimate behavior (modeling).
- Assimilate the many disparate security tools using both feed forward and feedback signaling mechanisms in a cyber defense system to help 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 (decentralized control). The body handles antigenic challenges through collaborative interaction. A similar distributed controls strategy (monitor and response) may be more resilient by avoiding single point failures enabling more robust decision-making.

7. Conclusions

We use the HPI metaphor as a vision for development of scalable trustworthy systems. Healthy humans live and flourish for 70+ years, while making real-time knowledge-based decisions to defeat continuous attacks by pathogens, toxins, and other environmental assaults. How does the body cope? Answers should provide insights with an emphasis on the ongoing exploration of this area, especially cross-disciplinary research bringing together computer scientists, biologists and immunologists.

Furthermore, we discuss some hard problems, including complexity, data requirements, data processing into information and knowledge, practical constraints, all users as insiders in a networked world and the growing number and sophistication of attacks. Under the HPI metaphor, the strategies to address these hard problems can lead to specific solutions. "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 systems 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; that will enable society to: (1) reduce the risk to highly critical systems and infrastructure, (2) thwart the sophisticated rapidly growing threat and (3) address other sector priorities such as eCrime-fraud.

Notwithstanding, the HPI metaphor may not always scale to the fast changing ever more sophisticated arms race. Indeed, our vision is both necessarily and purposefully general and high-level, because the challenge is a grand, both in terms of reverse engineering the brain (to include the various naturally evolved human defenses). and to achieving our vision. Yes, there are a number of provoking open questions to about the suitability of our metaphor. In particular, the scale of the human system will eventually be surpassed by the increasing complexity of cyber systems. Consequently, there may be a crossover point where HPI systems will simply fail to scale to cybersecurity problems, and it is unclear when –if ever– this point will be reached[1]. If so, certainly by then we'll have established a Cyber Center for Disease Control (C^2DC).

Acknowledgement. Oak Ridge National laboratory is managed by UT-Battelle LLC for the U.S. Department of Energy, under Contract No. DE-AC05-00OR222725. This submission was written by the authors, acting in their own independent capacity and not on behalf of UT-Battelle LLC, or the U.S. Department of Energy.

8. References

[1] C. Catlett, et al, "A Scientific R&D Approach to Cyber Security," Community-driven report submitted to the DOE, Dec. 2008 http://www.er.doe.gov/ascr/ProgramDocuments/Docs/ CyberSecurityScienceDec2008.pdf

[2] DHS S&T "Roadmap for Cybersecurity Research," Jan. 2009, http://www.cyber.st.dhs.gov/docs/DHS-Cybersecurity-Roadmap.pdf.

[3] National Science and Technology Council, "Federal Plan for Cyber Security and Information Assurance Research and Development," *Inter-Agency Working Group on Cyber Security and Information Assurance*, Apr. 2006.

[4] President's Information Technology Advisory Committee. Cyber Security: A Crisis of Prioritization, Feb. 2005, http://www.nitrd.gov/pitac/reports/20050301_cybersecurity/cybersecurity.pdf

[5] National Cyber Leap Year Summit 2009 Co-Chairs Report, http://www.qinetiq-na.com/Collateral/Documents/English-US/InTheNews_docs/National_Cyber_Leap_Year_Summit_2009_Co-Chairs_Report.pdf

[6] Michael Näf, "Ubiquitous Insecurity? How to 'Hack' IT Systems," *Information & Security* 7 (2001) 104-118.

[1] There is a lot of research left to better understand the disconnects between the suitability of the HPI model and for example the crossover point where natural/human systems fail to scale to cyber problems.

[7] M. Keeney, *et al.*, "Insider Threat Study: Computer System Sabotage in Critical Infrastructure Sectors," U.S. Secret Service and CERT Coordination Center (Software Engineering Institute of Carnegie Mellon University) May 2005.

[8] EMC^2 "The Digital Universe is still growing ". http://www.emc.com/leadership/digital-universe/expanding-digital-universe.htm. Accessed May 31th 2009.

[9] FT. Sheldon, S.J. Prowell, A. Krings, and R. Abercrombie (Eds.). (2010) *Cyber Security and Information Intelligence Challenges and Strategies*, ACM Proc. Sixth Annual Cyber Security and Information Intelligence Research Workshop, Oak Ridge National Laboratory, Oak Ridge, 21-23 April 2010.

[10] E.N. Marieb, *Human Anatomy & Physiology*, 5th Ed., Addison Wesley Longman Inc. Publ. (2001)

[11] Scientific Computing. "Gator Blood Destroys Deadly Superbugs" http://www.scimag.com/Gator_Blood_Destroys_Deadly_Superbugs.aspx?terms=alligator. 2008. Accessed June 1, 2009.

[12] Dipankar Dasgupta, Fernando Nino, Immunological Computation: Theory and Application. CRC Press, 2008.

[13] S. Forrest and C. Beauchemin "Computer immunology". Immunological Reviews 216 (1), 176.197, 2007

[14] Stacy J. Prowell, Rob Kraus and Mike Borkin, Seven Deadliest Network Attacks. Syngress: Boston, 2010.

[15] R.Q. Quiroga, L. Reddy, G. Kreiman,, C. Koch, and I. Fried, "Invariant visual representation by single neurons in the human brain," *Nature* 435 (23 June 2005) 1102-1107: [C.E. Connor, "Friends and grandmothers," 1036-37].