Maturing Cyber Security Using BioThreat Experiences and Resources

Norman Lee Johnson
Referentia Systems Inc
njohnson@referentia.com

Tim Williams
Referentia Systems Inc.
twilliams@referentia.com

Abstract

How does the current planning and response to cyber threats compare to biological threats planning and response? How do the resources of each compare? Biothreats have been a concern for millennia, and humans systems have had significant time and funding to develop a mature response. In this paper we observe that by comparison, cyber response is still in a relatively immature stage, possibly comparable to the state of public health protection prior to the implementation of safe water, sanitary conditions and vaccinations. Furthermore, we argue that because of the similarity between bio- and cyber systems, there are significant opportunities to advance the maturity of cyber research and response, either by using bio analogies for inspiration or by the direct transfer of resources. An analysis of existing cyber resources and gaps are compared to available bio resources. Specific examples are provided for the application of bio-resources to cyber systems.

1. Introduction

Cyber attacks are the most asymmetric of threats facing our nation today. A few individuals acting remotely can damage or destroy the operational capabilities of an entire government, military, and/or commercial sector – with minimal resources and preparation, with almost no risk during the attack, and with low likelihood of attribution. Our cyber vulnerability is partially persistent because of our limited success in managing our growing infrastructure complexity, in addition to the challenge of addressing known, resolvable cyber-security issues. Daily cyber attacks against commercial and government infrastructures are on the rise, and a report from 160 CEOs [1] suggest we are at risk of a “Cyber Katrina” unless action is taken. There are no shortages of studies identifying the problem and recommendations to solve it.[2] Policy statements, national position papers, and strategic federal agency plans have repeatedly identified strategic and operational cyber vulnerabilities, provided recommendations, and defined courses of actions over the last 5 years. The strongest recommendations are that:

- The greatest current challenge is our inability to address known vulnerabilities.
- Our information infrastructures, originally developed as security-neutral, must be transitioned to secure technologies, for example, making information assurance and identity management part of the infrastructure.
- The long-term management of the cyber challenge requires a system-wide engagement and commitment of all stakeholders, likely with a greater role for federal agencies.

The first two recommendations above are being addressed at some level by the nation and the armed services in the development of new cyber-security resources including detection, monitoring, analysis tools, training programs, and testbeds. But the final recommendation appears difficult to motivate and is illustrated by the observation that there is currently no capability to rank consequences against mitigation costs, particularly for high-impact but rare events such as a "Cyber Katrina." Another indicator of the lack of addressing the last recommendation is for preparedness planning: only one of the above-cited, high-level planning reports[3] call for predictive analysis technologies with risk assessment and consequence management to address the need for planning and response. Yet, predictive analysis technologies are central tools to other threat areas (chemical, biological, nuclear, radiological, etc.). This suggests that a major difference in maturity of planning and response systems exist between cyber and other threat areas.

The remainder of this paper examines the similarities between public and cyber health systems, how relatively mature the two domains are, and finally how activities in the bio-threat domain may help mature the cyber domain. For completeness we note that there are two application areas in the cyber domain which were inspired by the bio domain: computer security based on the adaptive immune systems [4] and simulations of the spread of computer viruses (or other replicating threats) based on epidemiology.[5] As will become obvious, these two areas of study, while important contributions to the cyber domain, represent a small part of the full opportunity.

2. The Difference in Maturation of Public and Cyber Health

A review of how public health resources has matured over time for biological threats is a helpful perspective for cyber preparedness. Figure 1 shows how until fairly recent times (150 years ago), public health experienced unstoppable and unexpected waves of epidemics, not too unlike our current experience with cyber threats. Removing these frequent epidemics from our society required that we develop healthy practices and infrastructures (safe water/food, sanitation) and specifically address certain known and reoccurring threats (smallpox, dysentery, bubonic plague, etc.) with vaccination or therapeutics. Once these preventative measures were operational, the public-health systems could focus on the relatively infrequent outbreaks of more difficult or unknown threats.

As we shift our cyber-health system by the implementation known countermeasures for common cyber threats, we will enter a similar phase of reduced "cyber epidemics."

In the above comparison of the development of biological and cyber health systems, the broad similarities are apparent. But, some might counter that there is a fundamental difference: biological systems have had the same host "technology" for millennia (our bodies), where technologies in cyber systems (host, networks, etc.) are constantly changing. This suggests that we could forever live in an epidemic-ridden cyber world, and never achieve the stable, mostly disease-free second stage found in public health.

A counterargument is that the bio-world is

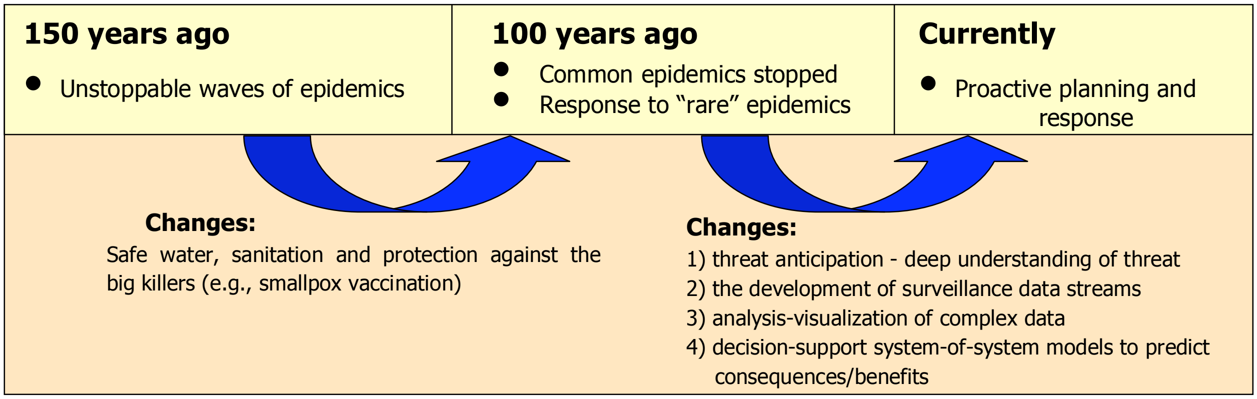

150 years ago
- Unstoppable waves of epidemics

100 years ago
- Common epidemics stopped
- Response to "rare" epidemics

Currently
- Proactive planning and response

Changes:
Safe water, sanitation and protection against the big killers (e.g., smallpox vaccination)

Changes:
1) threat anticipation - deep understanding of threat
2) the development of surveillance data streams
3) analysis-visualization of complex data
4) decision-support system-of-system models to predict consequences/benefits

Figure 1. How public health has changed over the last 150 years in the Western World.

equally adept at developing new "technologies" which exploit vulnerabilities, and our bodies have developed sophisticated multi-layered immune systems that have sustained the balance towards health. Furthermore, while the body "technologies" are unchanged, the interface between our bodies and our public health systems is complex and constantly changing as new health technologies are developed. In this argument, we are optimistic that comparable cyber-immune systems will be developed, and that a similar relatively "disease-free" cyber-health stage can occur.

Another lens on the relatively maturity of bio- and cyber-response programs is to examine defender activities that occur before and after an attack. One extreme is a purely responsive posture where you are primarily focused on containment of the attack and consequence management. The other extreme is where threat planning and response programs become more mature, as in the bio-threat space. Here, program activities address issues and opportunities well after an event (because they don't have to hunker down for the next attack) and well before (because of better preparation and understanding of the nature of the attacks and perpetrators). Post-event bio activities and programs include – listed from the event to much after – situational awareness, containment, consequence management, mitigation, forensics, remediation, and recovery. Pre-event bio programs include – listed from the event to much before – interdiction (stopping an attack closer to the source), anticipation, monitoring and detection, intelligence gathering on groups and possible resources, custom activities to limit entry of threats, export controls to limit technology leaks, and treaties and safeguards for nations to collaborate on reducing threats. All of these are on top of a public-health infrastructure which are coupled to these activities and minimize vulnerabilities.

Interestingly, cyber programs do have a few examples of these pre- and post-event activities (e.g., export controls on encryption, surveillance/monitoring resources, etc.), but generally resources are deployed by companies rather than federal or international agencies, unlike for biothreats where federal and international programs are the main source of funding and regulation.

From the broad perspective above, cyber programs are far less mature than the bio programs. It is therefore no surprise that recent policy positions of greater federal involvement and international cooperation on cyber threats are important to maturing our cyber defense. But how was this final maturation of bio-threat response accomplished at a more technical level? The relevance of the final transition of public health in Figure 1, occurring in the last 10 years or so, to cyber health is the focus of this paper: public health is undergoing a transition from a responsive posture (create a system to deal with the unknown threats as possible) to a proactive, defensive planning and response posture. The viewpoint of this paper is that the research and development activities for maturing cyber systems can be greatly advanced by a comparison to bio-planning and response programs. Instead of reinventing the wheel (and the car and the supporting infrastructure) for cyber security, can we leverage the knowledge and resources from existing, effective bio-threat programs?

3. Process/Functional Similarities

There are many levels of similarity between the processes of cyber threats and bio-threats which enable the use of bio-threat solutions and tools as templates – if not actual resources – for cyber research and tools. The greatest similarity, and the one that drives our choice of vocabulary for cyber threats, is the infectious processes. This process can be made more general, again using the bio-threat understanding, by identifying a threat-host process (of which the infection process is a subset), because some threats do not involve infection, such as allergies or denial-of-service attacks. Familiar bio-cyber examples for the threat-host process are:

- The viral spread by a compressed code that highjacks host processes,
- The signatures in the "genetic" code that can be used for identification,
- Signatures of the threat from its activity within the host or between hosts (in cyber systems these are, for example, access logs or non-essential files; in bio-systems these are non-essential bio-compounds), and
- The self-destructive immune response in the host to the presence of a threat.

On the host-response side, there are strong similarities at the function and process levels:

- The host immune state – as determined by immunization or prior or current infections – determines susceptibility,
- The host defensive options are similar in form, function and process – firewall-cell wall with preferential transport, layered defense systems, innate (always active) and adaptive (takes time to be active) immune response, system isolation and, if necessary, death of the host.

4. System-Wide Consequence Similarities

There are also similarities of the consequences due to changes in host activity on system-wide functions from a system-of-systems viewpoint. These can be broken down into direct and indirect or secondary consequences.

Direct system-wide consequences reflect the impact of the reduced activity or removal of the host on the system: the host both performs activities useful to the greater system (as a DNS server or a soldier), as well as being a repository of information for the rest of the system. Direct consequences can have short, medium and long-term impacts on the greater system depending on their function and how long they are degraded or removed from service. This bio-cyber similarity may enable some cost-benefit analysis resources that are used in bio-systems to be applicable to cyber systems. This statement needs to be qualified somewhat because of the observation that human and cyber coupling has quantitative differences in coupling with the greater system: humans require extensive coupling with other systems (transportation, different places to work and live, etc.) in comparison to cyber systems (e.g., cyber hosts don't work and live in different environments). But even this observation is rapidly changing as greater interdependence is becoming core to host functions in cyber hosts, such as the trend toward cloud computing/storage.

The indirect or secondary consequences – those system consequences that result indirectly from changes in the host activity or function, often due to interdependence of infrastructures – have greater similarity and, consequently, greater opportunities. A simple example is our power-generation and distribution systems: both rely on human and cyber support for continued operation. As human and cyber systems are compromised, the power grid becomes at greater risk of intermittency and possible collapse. Similar statements can be made for other infrastructures: banking, finance, water, food, transportation, etc. and many studies are being developed about the interdependencies of different infrastructures.[6] It is telling to note that critical infrastructure studies are only recently including cyber systems.[7]

5. Maturing the Domain: General Considerations

A detailed review of the mature programs and resources in responding to bio-threats (both those that naturally emerge, as in pandemic influenza, and those that are intentionally created, as in weaponized anthrax spores) is beyond the scope of this paper and is available elsewhere.[8]

In the previous discussion of the relative maturity of bio- and cyber-security programs, we observed that mature programs address the threat from end-to-end: from the control of technologies that can be used to develop threats, to the discovery and monitoring of potential attacking groups to addressing the long-term consequences of an event. Here we consider the relative maturity in more detail.

A specific example of mature bio-programs is the current planning, preparation and surveillance for pandemic influenza. The world has developed an extensive and cooperative sampling and surveillance system to monitor and warn of the expansion of "bird flu" and the occurrence of a contagious human form. Additionally, predictive planning and response tools were developed and used to assess different mitigation options and to deploy systems for the response to the pandemic. One tool[9] was developed and applied which simulated an epidemic in the entire U.S. population – 300 million people, the largest agent-based simulations used in production at the time – driven by census, workflow data and transportation data. Major changes in the mitigation strategies resulted from the application of this simulation tool in support of the White House's "National Strategy for Pandemic Influenza: Implementation Plan", May 2006.[10] No equivalent predictive planning resources or response plans exist or are being developed in the cyber realm.

Figure 1 identifies the resource components that brought about the most recent maturity of the bio-threat programs in order to transition from a responsive to proactive posture:

1) *Threat anticipation* – a deep understanding of the threat – its origins, forms, signatures, and, most importantly, potential variations;
2) *Surveillance of data streams* – providing indicators of the early stages of a possible attack and situational awareness of an ongoing attack;
3) *Analysis-visualization resources* of complex time-varying, heterogeneous data that result from 1 and 2 above; and
4) *Decision-support system-of-system models* to predict consequences/benefits/costs for planning and for forecasting the evolution of the current attack and assessing different mitigation options.

An analogy to a more simple threat system clarifies these resource components. Severe weather prediction, preparation and response have also undergone major advancements due to the development of the four components above, in particular, the development of data streams worldwide, simulation and analysis tools that are driven by these data streams, and decision-support tools.

An important observation is that the inherent, chaotic nature of weather systems requires a heavily data-driven approach – theory plus limited data is not sufficient. The same data-driven requirement is also true for bio-threats, both because of the inherent randomness of the system (such as the influence of random human-human contacts in the early stages on an epidemic), and because the attacker-protector dynamics (such as rapid change of virus from immune system pressure). Both of these sources of chaotic change can be observed in surveillance data, but are difficult to predict from theory (at best we can bound the degree of change – useful for planning but of limited utility in responding to a threat). In the absence of "theory" or detailed knowledge of the threat, then the data-driven approach becomes even more important.

Because of the similarities of weather-bio-cyber systems, we also expect cyber-security planning and response systems to equally require a data-driven approach. This approach includes using data streams for characterizing the range of threats and responses for planning, for surveillance of new threats, and for tracking the real-time system response to an evolving threat and attempted mitigations.

6. Maturing the Domain: Specific Guidelines from the Bio-Experience

Table 1 summarizes the resource components listed in the last section for the identification of resources and gaps to mature the cyber domain and then identifies the potential enabling bio-resources.

Because of the important role and opportunities of the different aspects of decision-support tools, three essential steps are identified in the maturation of decision-support tools for cyber programs:

1. The development of forecasting resources (typically in the form of simulations) – where we use the word forecast over prediction to indicate the chaotic nature of the systems and the need for a stochastic treatment,
2. The development of cost-benefit analysis resources (typically risk assessment and management tools) and
3. The development of integrated decision-support tools that combine all of the previous developments (data generation to analysis to prediction to cost-benefit).

Table 1 is far from being exhaustive and represents the authors' experiences (possibly myopic) into the bio- and cyber domains. Yet, even with this qualification, the gaps in a mature cyber-security programs are evident and intuitive. And, with some familiarity with the bio-threat resources, the possible opportunities for inspiration from the bio-domain, if not actual resources, are apparent. The next section provides definitions of bio-vocabulary or research areas that may be unfamiliar.

7. Useful Definitions in Bio-Threats

Threat Phylogeny: using the genetic code in the "genome" to determine the relationship between threats and their variations – often indicating their evolutionary linage and separation.

Virulence databases: a database (and understanding) of the genomic components that make a threat dangerous. For cyber it might be a "delete-all" call. Note that even though a genome or code may contain virulent factors, they may not be expressed.

Forensic tools: powerful analysis resources which connect the presence of a threat or characteristic to its source or history – perhaps the most developed application area for bio-threats outside of public health.

Syndromic surveillance: examines the statistics of symptoms appearing over time and location to identify health problems before physicians can diagnose them – these are becoming common in local public health departments and the military. Some bio-attacks can only be identified by this method.

Virulence change identification (ID): Identification of how a threat changes over time. We currently are tracking this for bird flu, to identify the remaining changes needed to observe a human epidemic.

Health metrics: measures of health of the public, etc.

Standardized threat scenarios: a set of scenarios (threats and deployments) that are broadly accepted by the community.

Threat anticipation: This is a very complex area. It can range from intelligence that tells you what your enemy is planning, to an analysis of your vulnerabilities and the resources available to identify where likely attacks could take place. The science-based form is essential for predicting the unexpected or unknown.

Table 1 - Illustration of how mature bio-threat resources can or may help fill gaps in cybersecurity

Cyber Resources Required for Mature Planning & Response	Existing Cyber Resources	Cyber Gaps: Needed Resources	Enabling Bio-Resources
Diverse cyber data: providing historical and real-time data of current network topology and traffic; enclave, component and user activity, access, status	**Rich and more in development -** Network flow traffic types/volume; component types & programs used	Status of components: susceptibility, symptoms of attack, readiness, activity, threat level	Genome” threat data bases, “virulence” databases, current threats, current news
Analysis and visualization of complex data streams: past and situational health, attacks, losses; global-to-local drill down, weak-signal precursors, threat ID and attribution, intuitive analysis of large data sets	**In development -** Large data set analysis identifying trends and precursors, anomalous behavior, ideally automated	Health of network and components, direct and inferred attack status, syndromic precursors to attack ID, forensics, threat attribution, …	Threat phylogeny, syndromic surveillance, health metrics, virulence change ID, forensic tools, responsiveness status, visualization resources
Predictive models of future state/losses from an attack given historical and current state, with transparency of outcome-to-cause and uncertainty quantification	**Scarce** - mostly academic simulations of network activity for limited threats; no exhaustive studies of tipping points	Databases of threats, standard threat models, emerging threat theory, effectiveness of response options	Epidemiological simulation resources, studies of mitigation options, coupled infrastructure sims, cost estimates,
Consequence - benefit resources including risk assessment, management and communication, expert-stakeholder conflict resolution, mission continuity	**Very limited for real-time response**; **limited for planning;** fundamental understanding limited	Metrics for mission readiness, threat-vulnerability mapping, integration of simulations	Standard threat scenarios for uniform preparedness, advanced risk assessment, adversary models,
Decision-support integration of above for planning and response: quantitative and transparent assessment of options, local-to-global cost-readiness tradeoffs, acquisition guidance, etc.	**Very limited** - currently wet-ware (human) based, no policy-level guidance on infrastructure acquisition, no operations support tools	Cost-benefit analysis of “what if” scenarios and response options; Risk management and communication	Threat anticipation-prediction, risk-based training, multi-stakeholder net-assessment studies, acquisition tools

8. Examples of Mapping the Bio to Cyber

Many of the “enabling bio-resources” in the right column of Table 1 require a lengthy discussion to explain and to exploit the perceived opportunities. The following highlights a few of the more easily communicated opportunities. The next section provides one detailed example of mapping a specific resource. A web search on key phases will lead the reader to more information. Please contact the authors for questions or assistance.

The opportunities that are most apparent for the cyber domain from the authors’ perspective are (in the order of top-to-bottom in Table 2):

- Development of cyber-threat databases that are based on the code content, independent of the expression/use of the code, to allow quick assessment of the threat potential. A subset of this database is a virulence database that contains coding that makes a threat "virulent" or destructive to the host or system.

- Threat phylogeny examines the evolution of threats based on their coding, to understand their origin and possible activity in certain hosts. Engineered threats make these evolutionary studies less useful, because large changes typically occur in engineered threats (even in bio-systems) in comparison to evolutionary changes. Syndromic surveillance examines the symptoms of host systems to detect a threat based on its effects rather than its direct presence. This type of surveillance is particularly useful in detecting unknown threats where the "genetic" coding is not known. Cyber surveillance has crude forms of this approach, such as observing unexplained increases in computation burden or number of files.
- Significant resources for epidemiological simulations over many scales (spatial and functional) are available in bio systems.[11] Some of these resources may be useful for cyber-system modeling.
- Standardized threat scenarios are useful in a maturing program to focus researchers, government and industry in developing countermeasures. For bio-systems collections of scenarios that spanned the range of threat types and consequences were particularly useful maturing awareness and focusing the discussion in a complex environment. A caution is necessary from the bio experience: standardized scenarios are good in early planning but their extended use can cause inability to adapt to new threats or developing a broader threat scope.

Once the threats are well characterized and their activity in the host is well understood, programs of threat anticipation can be developed that match the threat space to current vulnerabilities to anticipate where the next threat may occur and of what type. This level of understanding can be statistical if rich data is known on threat occurrences, can be based on intelligence information of groups in the process of developing threats, and/or can be based on a deep understanding of what threats are possible for given host systems. Threat anticipation represents a current research area for bio-threats and is quickly maturing.

9. A Specific Example of Bio-Cyber Mapping: Categorizing Threats

One of the core challenges in responding to a complex threat space (true for bio- and cyber domains) is to find some categorization of the threat space that helps in the planning of response options. We know that not all threats are equal in severity, sub-systems attacked, systems affected, etc., yet the complexity of the threat-host response can prevent "getting out of the weeds" and results in treating them all equally, at worse, or crude categorization into severe threats that must be addressed and others which can be deferred, at best. For bio-systems, the threat space is very complex and for a long time the complexity limited the planning possible. As suggested in Figure 1, experiences in threat and public health did finally lead to developing healthy living conditions and addressing the severe, reoccurring threats as possible (some threats, such as influenza, defy a general solution even though each year it kills many 10s of thousands of people in the U.S. and is costly from its impact on the workforce). A common view within the bio world is that public health programs removed the most dangerous threats as was possible for reoccurring and emerging (and possibly engineered) threats by the mechanisms listed in Figure 1. Even though this first revolution in public health reduced the expected epidemics, there remains great complexity in the threat space, and this limits the ability to develop additionally required operational responses.

One approach to simplify the threat space was proposed in a recent National Academy of Sciences report on chemical and biological threats: to divide the threat space by the responder's ability to quickly detect the threat and the ability to quickly treat the threat, as illustrated in a cyber version in Figure 2.

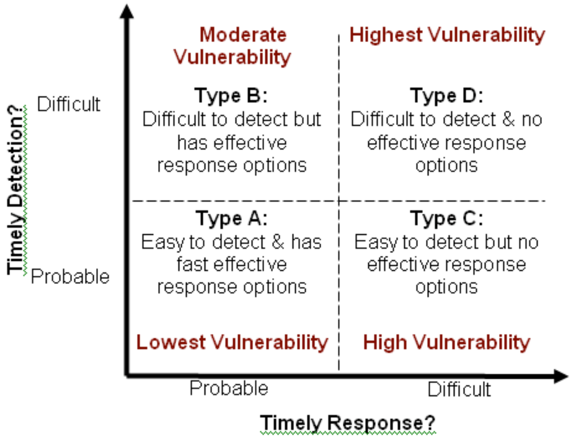
Moderate Vulnerability
Highest Vulnerability
Difficult
Type B:
Difficult to detect but has effective response options
Type D:
Difficult to detect & no effective response options
Timely Detection?
Type A:
Easy to detect & has fast effective response options
Type C:
Easy to detect but no effective response options
Probable
Lowest Vulnerability
High Vulnerability
Probable
Difficult
Timely Response?

Figure 2 - An approach to the simplification of the cyber-threat space, as inspired by the approach for the bio-threat space in a National Academy study on building protection

Figure 2 is a powerful threat characterization because it:

- Puts the complex variety of threats in a comparable and understandable basis (for example it can apply to both chemical and biological threats),
- Links measurable attributes (timely detection and response) to outcome: vulnerability and consequences, and
- Points to where the biggest challenges occur: difficult detection and slow response.

While these conclusions may seem obvious, they can be difficult to communicate to the less knowledgeable. The categorization can be useful in justifying a course of actions when budgets are limited.

The next stage in the application of the threat characterization landscape is to propagate the figure with known threats and their variations. This would help in identifying an existing threat that could be modified either to be more difficult to detect or more difficult to respond to, thereby increasing its consequences.

Another application of the threat characterization landscape would be to extend its characterization by adding a third dimension to include consequences of response options (high/low), because threats that have similar timely detection and response options could differ greatly by the consequences of the mitigation, e.g., continued normal operations or suspend all operations. This axis could include "levels of regret" as used in the bio-domain, to describe unavoidable consequences from a mitigation action even in the absence of the threat, as for example, establishing a preventative quarantine or taking a host offline.

10. Conclusions

The main objective of this paper is to present the cyber security researcher a broader perspective of their activities, as seen from the lens of the complex, but more mature field, of biothreat research and programs. The full breadth of such an inquiry is not possible in this short paper, but even at a summary level many possible opportunities for new areas of research become apparent. And, just as importantly, the comparison of the two domains provides the beginnings of a roadmap for how to mature cyber security, both for research and policy. Within this context and in developing this paper, the authors were reminded how the cyber community as a whole may be excessively focused on short-term concerns and miss the opportunities at the horizon, which may lead to long-term resolutions of current challenges.

Likely each reader of this paper will see different opportunities from the bio-threat arena, but as a summary the following were significant to the authors:

- The current emphasis of policy is aligned with the immediate needs in cyber security, e.g., addressing known vulnerabilities – and rightly so, but there is a noticeable absence of planning what comes next, once a stable, lower incident environment is achieved. Now is the time to begin investing in the next stages of threat characterization, to discover what the bounds of threats are, threat anticipation, and to identify future threats and their mitigations.
- Similarly there is a push towards more federal and international engagement in cyber security, as occurred in the bio-threat domain development. Many aspects of the bio-threat programs and specific technologies can be borrowed from the bio-threat areas, as for example, global surveillance of the outbreaks of threats or the monitoring of the "syndromic" signatures that suggest the presence of a unidentified threat.
- Many technologies or approaches can be transferred directly to the cyber domain, for example, the development of threat virulence databases, simulations for planning and response, forensic resources, and particularly decision support tools for the evaluation and selection of different response and mitigation strategies.

As a final remark, there are research areas where progress can greatly benefit both cyber and public health. The prime example is the importance of human factors (cultural, social, behavioral) on the formation, spread and response of bio- and cyber threats. For example, in biothreats the greatest source of uncertainty during an outbreak is how individuals will respond. Will they panic, possibly making the problem worse or will they follow directives from authorities? Little progress has been made to reduce these uncertainties (as illustrated that behaviors in simulations are prescribed rather than adapted to the current situation[9]), making planning for outbreaks challenging. Similar arguments can be made for cyber systems. How do users respond to a real or threatened attack? Do they make the problem worse if they panic? How can they sustain their productivity in the presence of mitigation responses to an attack? At best, currently studies can be done to bound the effects of behavior, but true forecasting of cyber or bio events for either planning or response requires knowledge of how the attackers, defenders and users behave.

[1] "Essential Steps to Strengthen America's Cyber Terrorism Preparedness: New Priorities and Commitments" Business Roundtable's Security Task Force, June 2006.

[2] ***Presidential directives*** (NSPD-38, NSPD-54, HSPD-12 and HSPD-23) http://www.fas.org/irp/offdocs/nspd/index.html, C***ongressional studies*** (e.g., CRS's Economic Impact of Cyber Attacks April 2004 (http://wikileaks.org/leak/crs/RL32331.pdf), IP3's National Cyber Security Research and Development Jan 2009 (http://www.thei3p.org/docs/publications/i3pnational cybersecurity.pdf), Information Security GAO Jun04, Cyber Analysis and Warning July 2008, (http://www.gao.gov/products/GAO-08-588) Infrastructure Protection GAO-08-825 Sept 2008, (http://www.gao.gov/products/GAO-08-825) Critical Infrastructure Protection GAO-08-1157T Sept 2008, (http://www.gao.gov/products/GAO-08-1157T) ***Interagency councils*** (e.g., National Cyber Study Group, Joint Interagency Cyber Joint Task Force), ***Intelligence studies*** (e.g., National Intelligence Council's Global Trends 2025 (http://www.dni.gov/nic/NIC_2025_project.html)), ***Public policy institutions*** (e.g., CSIS's Securing Cyberspace for the 44th Presidency Dec 2008, (http://www.csis.org/media/csis/pubs/081208_securin gcyberspace_44.pdf)), ***Federal agency studies*** (e.g., DHS's National Strategy to Secure Cyberspace (http://www.dhs.gov/xprevprot/programs/editorial_0 329.shtm), DOE's Scientific R&D Approach to Cybersecurity Dec 2008 (http://www.er.doe.gov/ascr/ProgramDocuments/Doc s/CyberSecurityScienceDec2008.pdf)), university centers (e.g., GTISC's Emerging Cyber Threats Report for 2009

(http://www.gtiscsecuritysummit.com/pdf/CyberThreatsReport2009.pdf)– Jan 2009)

[3] (GAO Jul 2008)

[4] S. Forrest and C. Beauchemin. Computer immunology. Immunol. Rev., 216, pp. 176-197, 2007.

[5] M. Newman, S. Forrest, J. Balthrop, Email networks and the spread of computer viruses. Physical Review E, 66, 035101:1-4, 2002.

[6] J. Grenier, "The Challenges of CIP Interdependencies", Conference on the Future of European Crisis Management (Uppsala, 19-21 March 2001)http://www.ntia.doc.gov/cip/workshop/cipft_files/frame.htm.

[7] M. Dunn and I. Wigert, International Critical Information Infrastructure Protection (CIIP) Handbook 2004. Swiss Federal Institute of Technology, Zurich, 2004

[8] L.E. Lindler, F.J. Lebeda, and G. Korch Biological Weapons Defense: Infectious Diseases and Counterbioterrorism, Humana Press, New York, 2005.

[9] T.C. Germann, K. Kadau, I.M. Longini, and C.A. Macken, "Mitigation Strategies for Pandemic Influenza in the United States," Proceedings of the National Academy of Sciences 103, 5935-40, 2006.

[10] http://www.pandemicflu.gov/plan/federal/pandemic-influenza-implementation.pdf

[11] L. Sattenspiel, A. Lloyd. "Modeling the Geographic Spread of Infectious Diseases: Report on the Critical Review of Geographic Epidemiology Modeling Study." Prepared for the Defense Threat Reduction Agency, DTRA01-02-C-0035. April 2003. http://www.dtra.mil/asco/ascoweb/CompletedStudies.htm