PNNL-SA-107795

Pacific Northwest
NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

AiR: Asymmetry in Resilience

Report on the First Meeting on Asymmetry in Resilience for Complex Cyber Systems

December 2014

Authors: Christopher Oehmen, PhD
Nicholas Multari, PhD

U.S. DEPARTMENT OF ENERGY

Prepared for the U.S. Department of Energy
under Contract DE-AC05-76RL01830

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062;
ph: (865) 576-8401
fax: (865) 576-5728
email: reports@adonis.osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312
ph: (800) 553-NTIS (6847)
email: orders@ntis.gov <http://www.ntis.gov/about/form.aspx>
Online ordering: http://www.ntis.gov


This document was printed on recycled paper.
(8/2010)

PNNL-SA-107795

AiR: Asymmetry in Resilience

Hosts: Nicholas Multari, PNNL
Christopher Oehmen, PNNL

Co-Host: Marco Carvahlo, FIT

Participants/Contributors:

Janine Anderson, PNNL
Scott Borg, USCCU
Tom Clark, AFRL
Rich Colbaugh, Periander
Mike Corsello, PNNL
Anurag Dwivedi, JHU/APL
Barbara Endicott-Popovski, UW
Scott Godwin, PNNL
Will Hutton, PNNL
Craig Jackson Jr, Indiana U
Mary Lancaster, PNNL
Patrick Mackey, PNNL
David Manz, PNNL
Miles McQueen, INL
Rick Metzger, AFRL
Tyler Moore, SMU
Ann Nagel, UW
Charles Nelson, OSTP
Jeff Picciotto, MITRE
Pradeep Ramuhalli, PNNL
Craig Rieger, INL
Greg Shannon, CMU/CERT
Seth Shapiro, Kibble and Prentice
Von Welch, Indiana U

December 2014

Prepared for the U.S. Department of Energy under Contract
DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99352

Table of Contents

Acronyms and Abbreviations

AiR	Asymmetry in Resilience
APL	Applied Physics Laboratory
CMU/CERT	Carnegie Mellon University, CERT® Division
DARPA	Defense Advanced Research Projects Agency
FFRDC	Federally Funded Research and Development Center
IP	Intellectual Property
MIT	Massachusetts Institute of Technology
NIST	National Institute of Standards and Technology
OSTP	Office of Science and Technology Policy
PNNL	Pacific Northwest National Laboratory
SMU	Southern Methodist University
SP	Special Publication
U.S.	United States

Asymmetry in Resilience (AiR)

Hosted by Nick Multari and Chris Oehmen (Pacific Northwest National Laboratory)
Co-organized by Marco Carvalho (Florida Institute of Technology)

Crystal City, VA (September 17-18, 2014)

1. Purpose and Major Outcomes

Currently a worldwide collection of malicious actors has free reign to probe, attack, monitor, and manipulate networks including those crucial for supporting national critical infrastructure. Many research efforts are already underway to develop *resilient* cyber systems. However, a crucial component of making these a reality is the ability to measure and shift the cost balance that currently favors the attacker, in part using these emerging resilience techniques. Realizing that candidate resilience technologies have varying degrees of effectiveness and varying costs for using them (economic and resource), it is essential to understand the cost and benefit of defensive cyber technologies and techniques in light of their ability to provide an asymmetric advantage to defenders in order to maximize the effectiveness of cost applied toward winning cyber conflict.

AiR was a two-day meeting for luminaries from universities, government institutions, FFRDCs, and industry to identify the key challenges in shifting the advantage in cyber conflict to favor the defender in the context of resilient cyber systems. The primary outcome of this meeting was a prioritized consensus vision for realizing cyber asymmetry and a research roadmap. A secondary outcome of the meeting was the formation of provisional collaborative (multi-disciplinary, multi-institutional) teams to begin refining the concept or cyber resiliency and creating advocacy in the wider cyber resilience community.

The AiR meeting participants developed a consensus working definition of asymmetry as "*a disproportionate, exploitable imbalance between actors related to, but not limited to, resources, level of effort, risk, or consequences in an attack*". While this definition provides a framework to begin empirical study of asymmetry for cyber conflict, the AiR group acknowledges that a more formal definition may be required as a foundation for rigorous scientific exploration of the concept of asymmetry.

We also explored many facets of asymmetry to develop a conceptual basis for the concept as it applies to cyber conflicts. We examined the relational and adversarial dimensions of asymmetry, in fields ranging from economics, psychology/human factors, warfare, and immunology. As a source of inspiration, we explored state-of-the-art technologies that could be components in achieving asymmetry in cyber resilience and drew from instances of asymmetry in other domains such as biology.

Using this exploration as a backdrop, we defined technology gaps and priorities for research and development. The AiR group also developed a research roadmap designed to provide a foundation for rigorous research around asymmetry in conflicts for resilient cyber systems. This research plan consists of the following steps:

1. Rigorously Define Asymmetry for Cyber Resilience

2. Examine Asymmetry in Other Domains

3. Develop Asymmetry-Specific Case Studies for Cyber Resilience

4. Develop Cyber Benchmarks, Measure, and Evaluate

The remainder of this document contains: 1) a list of the represented institutions; (2) a description of the breakout groups; 3) a summary of findings; and 4) closing thoughts on next steps for research in cyber asymmetry. The appendix contains the agenda for the two-day meeting.

2. Meeting Participants

AiR meeting participants were chosen to represent a broad set of institution types and areas of expertise. The group spanned Department of Energy Laboratories and other FFRDCs, industry, government, and academia. Participants were also chosen to represent measures/metrics, risk management and models, economics of cybersecurity, high availability systems, and behavioral models as related to cyber resiliency.

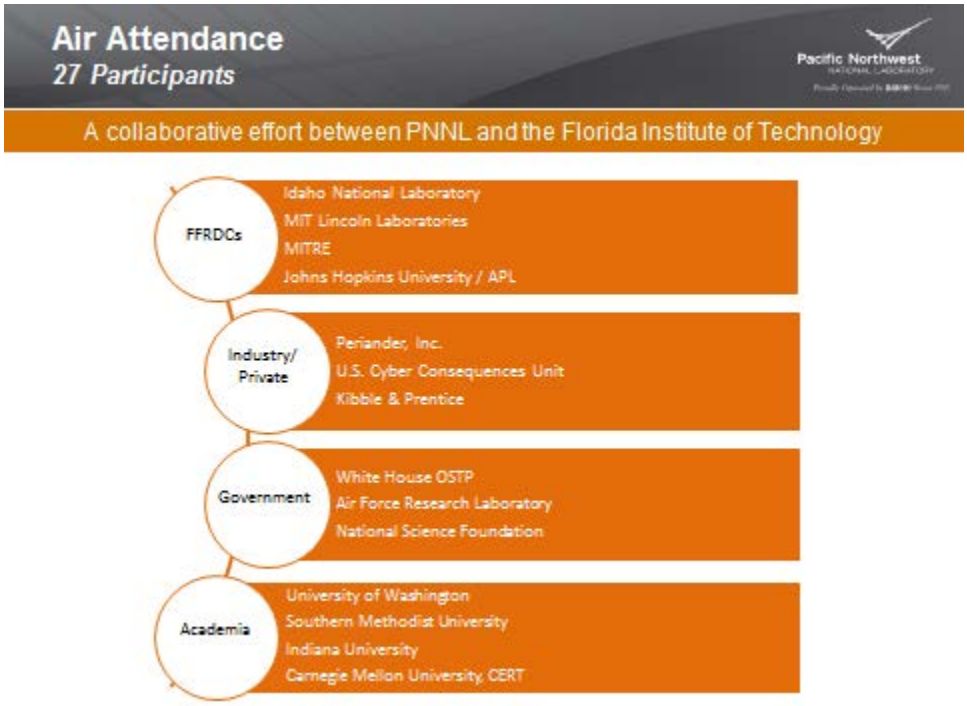

The 27 meeting participants divided into three working groups of 9 participants each. The goal of the working groups was to provide an environment conducive for discussion of each breakout topic. Working groups were constructed to have balanced representation from across the five areas of host institutions, FFRDCs, industry or private institutions, government agencies, and academic institutions. Each participant was assigned to a specific working group to ensure an even distribution of technical backgrounds across the groups.

Breakout topics were planned to ensure a logical progression from defining asymmetry in cyber resilience to establishing a research roadmap that would enable the evaluation of defensive cyber technologies. In each breakout session, participants were encouraged to share their point of view and collectively work toward assimilating these various perspectives into cohesive summaries for presentation in the plenary sessions.

3. Findings

3.1. Definition of Asymmetry

Finding 1: After examining a number of potential definitions of asymmetry in a resilient cyber environment, the group arrived on a consensus of the following definition:

- Disproportionate, exploitable imbalance between actors related to, but not limited to, resources, level of effort, risk, or consequences in an attack

Finding 2: There was no consensus on whether asymmetric is its own discipline or must be examined in context to resiliency and cybersecurity.

Each group developed a working definition for asymmetry as it pertains to resilience of cyber systems. We note that asymmetry is a well-understood concept in physical warfare (military history), but is less well understood in the cyber realm. The group generated several working definitions that are included here for reference. The last definition was the consensus definition agreed to by the group.

- Disproportionate, exploitable imbalance of resources, cost, speed, information awareness, and impact (including social and legal) among competing parties
- Given a conflict between two or more adversaries where each of the adversaries have well defined rules and goals, one organization is threatened with serious damage by an adversary with very different size and capability
- Quality that creates an imbalance in the level of effort in terms of cost, time, or consequences
- Inequality of resources, capabilities, knowledge, tools, and/or motivation between attacker and defender
- Non-reflective image of opposing sides
- A quality that creates imbalance between actors in the resources, level of effort, risk, or consequences in an attack
- Enormous value creation with limited means
- Extreme efficiency that allows the asymmetry of capabilities
- Overwhelming inequality between cost/value ratios relative to attack/defense mechanisms
- Crafted overwhelming inequality through terrain manipulation and innovation
- Disproportionate, exploitable imbalance among competing parties
- Disproportionate, exploitable imbalance between actors related to, but not limited to, resources, level of effort, risk, or consequences in an attack

One question that arose during the meeting regarding asymmetry as it applies to resilient cyber systems focused on whether asymmetry is its own discipline, or whether it is a modifier of existing disciplines. There was no immediate agreement on whether the asymmetry concept is sufficiently well developed to apply ‘as-is’ to resilience-focused work versus in need of or worthy of its own research agenda. If the former, then we must begin with the context of resilient cyber systems and derive foundations for asymmetry in this context. On the other hand, if the latter, our goal would be to develop a foundation for asymmetry in general with an emphasis on determining its abstract properties and invariants and applying the result to resilient cyber systems.

3.2. Dimensions of Asymmetry

Finding 3: Asymmetry is only meaningful in the context of a *relationship* between cyber adversaries and defenders

Finding 4: Capability asymmetry is a fundamental game-changing imbalance asymmetry that we wish to realize for cyber defenders.

Finding 5: A less-resourced defender can enjoy an asymmetric advantage over a well-provisioned attacker.

Finding 6: If we can understand the nature of the conflict environment (terrain), we may be able to find simple methods to influence it to disproportionally favor defenders

Asymmetry is only meaningful in the context of a *relationship* between cyber adversaries and defenders. Asymmetry is not an inherent attribute of a resilient system. The dynamic nature of actors' interaction means that asymmetry drives a *coevolution* that occurs during a cyber conflict. We can glean two main forms of asymmetry from physical warfare tactics: numerical asymmetry and capability asymmetry. *Numerical asymmetry* is simply numerical superiority in the sense that having ten weapons would be better than having only one weapon. Capability asymmetry is a fundamental game-changing imbalance more akin to having a gun vs. a knife. It is capability asymmetry that we wish to realize for cyber defenders. This *material* advantage applied to the cyber world would be the presence of technologies or tactics that give a fundamental, disproportionate advantage, in our case, to the defender. This advantage would come in the form of a *force multiplier* as opposed to just a numerical advantage.

The group also explored asymmetry as an imbalance between adversaries and defenders in terms of organization size, resources, and objectives. Discussions concerning organizational size focused on cases where a nation state or criminal organization sponsored the adversary. In these cases, it is typical that the adversary has access to various experts and the resources to fund them. This is in contrast to the defender's organization that may consist of one to a few experts who are resource constrained. However, research in this type of asymmetry in other fields (Arreguin-Toft 2005) can provide insights into how a less resourced defender can still enjoy asymmetric advantages.

An organization's objectives are a commonly raised issue as they relate to cyber resiliency and security. One aspect of asymmetry currently working against defenders is the limited and pinpoint objectives of the adversary versus the defender's requirement to continuously protect all critical resources.

An additional topic introduced at the workshop that merits further exploration is that of "terrain." While it can be extremely difficult to define the terrain of the cyber battlefield (the concept of distance has little meaning, so what do we really mean by "terrain?"), the group believes that if we can understand the nature of the conflict environment, we may be able to find simple methods to influence it to disproportionally favor defenders. This could take the form of defenses that impose on the adversary an extremely high cost or number of resources to overcome or take the form of low-cost defensive measures (such as cryptography or air-gapped systems) defeating high-cost attacks.

3.2.1. Economic-Related Attributes

Finding 7: The attacker and defender each assign their own value to the contested system. The fact that these values are unlikely to be the same is significant.

Finding 8: While the relative cost and benefit can serve as a disproportionate incentive, the ability of parties involved in cyber conflict to inflict consequences on one another could serve as a deterrent.

Economic attributes of asymmetry relate to the notion of *cost* paid or *resources* consumed by attacker and defender, the *value* of the contested systems and data, and the *consequences* associated with attack and defense. Cost includes the monetary cost of systems, time, exposure (a cost that the adversary may perhaps want to minimize even more than monetary expense), and the indirect cost of other priorities that are not met when focusing on a particular defense. The notion of cost therefore applies to any valuable *resource* that is brought to bear for defender or attacker.

Balancing the notion of cost is the notion of *value* or importance of the contested components to the overall success of the organizational mission. For example, two servers may be composed of the same hardware and operating system and therefore have very similar cost. However, if one is hosting sensitive intellectual property and the other stores payroll information, then the relative value they have in supporting an organization's mission of generating new intellectual property is very different. The attacker and defender each assign their own value to the contested system. The fact that these values are unlikely to be the same is significant. For example, the defender might dilute the value of stolen data for the attacker by deliberately introducing false or partial data that could be resolved by the defender given their more complete context.

Supplementing the concepts of value and cost is the concept of *consequences* or *impact.* In some cases, consequence is just another form of cost. For instance, losing an email server has an associated cost of lost productivity. In this case, the cyber system is itself the supported mission. However, in cases where systems are co-dependent (such as electric power and communications critical infrastructures), loss of one has consequences beyond cost incurred by the victim organization. In this case, consequence is more than just a cost function. While the relative cost and benefit can serve as a disproportionate incentive, the ability of parties involved in cyber conflict to inflict consequences on one another could serve as a deterrent. However, this relies on the ability to reliably attribute malicious behavior to individuals or their stakeholders, and, for law-abiding defenders, clear legal authority to act.

3.2.2. Additional Behavioral Factors

Finding 9: There may be an imbalance in an adversary's and defender's *willingness* to engage in certain activities

Meeting participants also identified attributes of asymmetry that relate specifically to the human aspect of cyber conflict. While there is a wide range of socio-cultural factors that bear on asymmetry, we focused on the fact that there may be an imbalance in an adversary's and defender's *willingness* to engage in certain activities. In some cases, this imbalance comes from regulations or legislation that governs defender behavior. We assume, in general, the defender is attempting to be lawful, meaning that they will have a low willingness to break rules in pursuit of a strong defense. By contrast, attackers may not have these same constraints allowing them a disproportionate range of possible actions.

3.2.3. Possible Means to Achieve Defensive Asymmetry

Finding 10: Several modes of developing and exploiting a disproportionate advantage are available to defenders. Such modes include resource, capability, and speed or agility advantages.

Several modes of developing and exploiting a disproportionate advantage are available to defenders. In some cases, defenders may enjoy a resource advantage making it possible to overwhelm attackers with a hostile environment. Alternatively, the defender may be able to exploit an imbalance in *capability*, for instance, by introducing hard problems to solve or by using alternative technologies that defeat common computational capabilities (like quantum computing). While resource and capability advantages often exist for very large cyber institutions, smaller ones might exploit imbalances in *speed* or *agility*. Moving (or giving the illusion of motion) or dynamic reconfiguring may not be tenable at a global enterprise scale. However, since adversaries rely upon a static environment in order to ascertain the targets and exploit their vulnerabilities, these may be cost-effective techniques for small installations. Larger enterprises may also employ these techniques for high-value subsystems. In many cases, one goal is to remove the time advantage that the adversary currently enjoys. By shifting the *time scale* in which dynamic resilient actions occur, a defender may be able to prevent adversaries from making reliable assumptions about the target environment.

3.3. State of the Art for Asymmetry

Finding 11: There is no significant existing body of research or technologies specifically focused on formalizing an approach to using what is known about asymmetry to achieve resilience.

Finding 12: The current preferred tools that support defensive activities toward an asymmetric advantage span several categories, but are, in general, point solutions.

The AiR group noted that there is no significant existing body of research or technologies specifically focused on formalizing an approach to using what is known about asymmetry to achieve resilience. One important aspect of achieving asymmetric advantage for the defender in the cyber conflict is the ability to assess the space of adversarial actions. In some cases, this is accomplished using standard methods for assessing *risk*, such as NIST SP 800-30 rev 1, “Guide for Conducting Risk Assessments” (National Institute of Standards and Technology September 2012). There are also future-looking techniques to “play through” ensembles of scenarios to assess the relationships between adversary and defender using *game theory*.

The current preferred tools that support defensive activities toward an asymmetric advantage span several categories, but are, in general, point solutions that confer advantage to specific defender subsystems for given attack scenarios. Protection of the data layer is the most mature, having several well-defined techniques. For instance, *cryptography* (Van Leeuwen 1990) is a well-established approach to protect the confidentiality of data by introducing a disproportionate problem for attackers to solve to make sense of the raw data. In fact, the fields of *information security* and *information assurance* (Bidgoli 2006) have yielded many tools and techniques that are applicable to defending data.

Similarly, *formal methods* (Rushby 1997) is a branch of mathematics that can be used to create and verify mathematically correct subsystems. Formal methods can be very costly but may only need to be applied

once, so their high cost is amortized over the lifetime of the proven system resulting in a cost-effective defense.

One interesting avenue for achieving an asymmetric advantage in cyber systems is the drive to find *novel signatures* or *behavioral indicators* (Lazarevic 2003). In some cases, successfully developing signatures using previously untapped data sources presents defenders with a method for achieving game-changing situational awareness. Often such advantages are short-lived because attackers who are thwarted by these methods will probe and test to find out which of their behaviors are being newly discovered and alter their behavior accordingly.

The early days of email filtering technologies provide an example of this back-and-forth. At first, unwanted advertising email arrived at target email addresses unchecked. At some point defenders learned to analyze the sender's attack attributes to discover novel signatures of unwanted email. When the success rate of attacker emails getting through dropped too low, attackers changed their tactics and spoofed the sender, successfully defeating the signature. Defenders then began to look through the email body for keywords that were frequently associated with unwanted email, again gaining the upper hand for a time. Attackers again changed their tactics by disguising keywords with misspellings and other techniques that human readers could still interpret, but which would not match terms in keyword lists. As email filtering technologies continue, the game-changing methods for finding new signatures often represent fundamental shifts in the terrain over which this conflict is conducted, but the co-evolutionary nature of the conflict can make these novel signatures short-lived.

There are some analogs to cyber defense for which concepts of asymmetry are established and could be used as a starting point to develop foundations for cyber systems. In *physical warfare*, the notion of asymmetric force is a well-established doctrine that describes the advantage enjoyed by game-changing technologies such as nuclear technologies that render competing technologies irrelevant for defined scenarios (Metz 2001; Hartman 2002). Similarly, naturally occurring biological systems exhibit a wide range of resilience responses that are meant to provide asymmetric advantages (Rahme 2009; Rewald 2009; Mingo 2009; Gordon 2009). For example, viruses affect cells by co-opting the cellular machinery normally used to sustain life (Black 2012).

4. Gaps in Capabilities and Strategies

The AiR group identified a wide collection of gaps in current technologies and practices that must be addressed to shift the asymmetric advantage in cyber conflict in favor of the defender. These technology areas are reported here in two categories—essential and desired. *Essential* technologies are those that we feel are required for developing sound approaches to realizing asymmetry. *Desired* technologies are those that would increase the effectiveness of asymmetry-related technologies, or are parallel activities that would be necessary for dissemination or effective use of these technologies.

In general, these technologies should help us move away from the vulnerability-centered approaches that currently dominate the landscape. In conjunction with that shift, user-focused design and implementation would promote practical implementations of these concepts for use by defenders. The ultimate goal is to engineer *security by design* into systems that is usable and informative to practitioners.

4.1. Essential Capabilities and Strategies

During the meeting, the participants identified a number of capabilities deemed essential to support moving the asymmetric advantage from the attacker to the defender. These include capabilities that provide an understanding of the events within and external to pre-planning strategies to increase the cost to the attacker at a cost acceptable to the defender.

Metrics. Measureable progress in leveraging the asymmetry concept requires a formal definition of asymmetry. In line with this definition is a clear delineation of the relationships between asymmetry, resilience, and cybersecurity. An additional requirement is the ability to measure attributes of cyber systems as they apply to the relationship between attacker and defender. This includes model abstractions spanning the space of what is knowable as well as a clear assessment of the *cost* and *benefit* of various actions. Metrics will also be required for quantitative self-awareness through, for example, *attack surface quantification*. In addition, metrics will be required for quantifying and comparing the degree and value of asymmetric advantage for given conflicts.

Situational awareness. Closely related to metrics is the need for real-time situational awareness. This would include intra- and inter-organizational information sharing and legally compelling attribution of all packets. Such awareness might come from a variety of sources including *streaming, distributed sensors* that are designed to collect data that is demonstrably relevant for assessing asymmetric advantage. In parallel with sensors on the defended system, models of the system would give rise to additional measures via methods like *distributed graph analytics* operating on models of complex cyber systems. Because they would model dynamic environments, these graphs would have to be linked to *automated discovery of assets and their relationships*. This includes resolution of the many *complex interdependencies* typically present in complex cyber environments as well as the relative *value* and *criticality* of the cyber assets.

Both sensor- and model-based measures would be potential sources of *novel signatures of behavior*. However, for this awareness to be useful, it will be necessary to augment existing methods to capture

mathematically the degree and impact of uncertainty that arises in measurements and processing. This process of *uncertainty quantification* would need to relate sources of uncertainty with their impact on downstream processing and decision-making. Understanding the provided measures and their degree of uncertainty are key components in enabling both automated and human-assisted decision processes. This information must be in a form that takes into account how the information is used by either the automated system or the human defender. Beyond understanding how humans act in the *role of a defender*, though, an understanding of how *humans act as attackers* and (sometimes passive) *facilitators of attacks* is also required to make optimal and informed decisions for resilient responses that lead to an asymmetric advantage.

Pre-attack detection and strategies. Strategies that promote activities far in advance of an attack are also of utmost importance. In this sense, the notion of asymmetry would be with respect to classes of threats that may be utilized by adversaries. One class of activities that supports this shift is information sharing, likely *threat intelligence sharing* at the classified level, as close to real-time as possible. This is a special category of previously mentioned information sharing that is geared toward discovering targeting and other upstream precursors of a cyber-attack. To make this possible for most organizations, the challenge would be to obfuscate data to make it suitable for sharing without removing the patterns of interest for discovering threats.

A second class of activities employs *predictive analytics*. We posit that there are discoverable and reliable pre-indicators of imminent threat that can be used to inform the selection and implementation of pre-positioning to give the best asymmetric advantage – actions that change the landscape that adversaries would encounter, giving the defender an asymmetric advantage. Advanced methods for discovering these indicators would need to be developed.

Post- and during-attack strategies. Given that some attacks will be successful in penetrating installed defenses, several categories of desired capabilities exist that would support rapid awareness and regrowth of functionality. Systems with the ability to *self-heal* might limit the effectiveness of adversary actions and at the same time promote the functionality of a cyber system – in particular, the subsystems that allow the system to repel an attack. Self-healing includes the functionality of systems and applications, as well as the underlying data itself. The notion of *self-healing data* as a strategy could eliminate the effectiveness of data manipulation. Self-healing in general could be coupled with *dynamic reconfiguration* that results in constantly improving defensive posture, as opposed to healing that returns to a last known good state.

Understanding *attacker strategies* and methods as they unfold during a cyber conflict and reliable *attribution* can give the defender an asymmetric insight into the capabilities of the attacker (in much the same way that attackers now use knowledge of the defender's environment to their advantage). *Low-profile mitigation* is one approach to managing attacker strategies that manipulates the adversaries' tradecraft without revealing that their approach is not working. While target-of-opportunity adversaries are best defended by actively blocking their efforts, the goal behind low-profile mitigation is to have determined adversaries use ineffective strategies as long as possible. If our response to an attack reveals to our adversaries that they have been discovered or that we have an effective countermeasure, it would force them to evolve, potentially beyond our ability to detect or mitigate their methods. Ultimately, the goal is to shift the balance of the rate at which defenders and attackers evolve by speeding up evolution for the defender and slowing it down for the attacker.

4.2. Contributing Capabilities

Along with the technologies described above that we believe are essential to realizing asymmetric advantage for cyber defenders, there are several issues that, while not central to solving the asymmetry challenge, may contribute fundamentally to success or failure of developed approaches.

Transparency in complex interactions between cyber systems. Complex cyber systems do not exist in a vacuum and boundaries between institutions are not always well defined. While individual systems (e.g., hardware) can be associated with a discrete cyber system, these cyber systems depend on other systems, in particular critical infrastructure cyber systems. This means that there are cascading interdependencies that cross-domains and organizations, and should figure into resilience decisions. Adversaries might choose to attack indirectly if they perceive an advantage by attacking infrastructure on which an institution relies and for which it does not have ownership, insight into, or direct means to defend. The lack of legal constraints for some adversaries and their willingness to cross complex organizational or legal boundaries may introduce asymmetric advantage against defenders who must coordinate. This is made more complex by the possibility that an effective defense in one context may actually be detrimental to the larger system in which it exists (and on which others rely). The fact that transiently connected, personally owned devices now comprise a significant portion of many cyber infrastructures further blurs the line between systems. Taking this a step further, we must also account for the unintended negative side effects (cyber fratricide) of resilience actions that are meant to confer asymmetric advantage for a defender.

Awareness of Cyber Operators. This raises several questions that need to be addressed.

- Is the defender a cyber defense unit at a particular institution or part of a larger consortium? If the former, the advantage is that each defender is smaller than the collective is and therefore more nimble. The downside of this approach is the possibility of defenders creating negative consequences on each other.

- How do they communicate what they perceive to other defenders moving in the same terrain to prevent interference, and how do multiple cyber units coordinate with defenders of a common infrastructure? If instead the defender is defined as a consortium or “the community of cyber defenders at large,” then actions taken by individual organizations would consider their impact on the larger super-system, ideally resulting in resilience of the U.S. (or global) infrastructure as a whole, but potentially at the expense of individual institutions.

- How do we ensure that local perception and actions are communicated and coordinated across geographically and institutionally isolated subunits? Whether we take an institution-centric or communal view of the defender’s context, it will be essential to share data effectively without allowing this data, or the act of sharing it, to introduce new vulnerabilities. It will also be important that laws and incentives be in place to promote appropriate sharing and coordination.

Balance between automation and human-driven response. Some phases of cyber conflict can unfold with amazing speed as automated systems make sub-second scale decisions. Other phases of cyber conflict occur over months or years as adversaries observe, probe, and plan to attain their ends as needed. It is true that in order to keep up with adversaries engaged in the agile stages of conflict, we would need to shorten our detection and mitigation strategies timeframe, in many cases automating responses that

currently operate at human speeds. However, some AiR participants stated the notion of fully automating systems is probably not the right model either, in part because human insight is often a critical component of understanding what is happening during a conflict and how best to react. We propose an approach to cyber resilience in which 1) processes that can be reliably (and perhaps provably) automated are implemented to enable machine-speed sensing and reacting to some phases of cyber conflict; and 2) humans steer these systems at a high level, informed by streaming awareness of system state and automated decisions being made. It is an open research question whether such an approach can be constructed to increase the efficiency of human-steered machine responses and shift the asymmetric advantage in favor of the defender.

Policy and human resource issues. The AiR group identified several policy-related aspects to realizing asymmetric advantage for resilient cyber systems. First, international *policy* must be in place that clearly defines roles and responsibilities for coordination among cyber defenders. This would include *incentives* for implementing resilient cyber systems that provide asymmetric advantage to defenders and a framework for compliance with such an approach. Any compliance-based method would have to ensure that the incentivized behaviors lead to the intended outcome. Mechanisms for sharing data must ensure that competitive advantage is not lost in the act of sharing the data. Likewise, in many domains anti-trust or anti-collusion laws and regulatory oversight explicitly forbid sharing certain data. Effective coordination for cyber defense will only be cost effective (and hence contribute to asymmetric advantage for the defender) if a conducive environment can be established through a clear policy and incentive framework.

Second, in addition to protecting the confidentiality of shared intra- and inter-organizational information, the privacy of individual users must also be protected. In this context, privacy refers to an individual's right to be free from unwanted surveillance, and the ability to maintain and control the confidentiality of some personal information. To provide privacy while also providing for organizational confidentiality, we must find ways to ensure that sharing data for the purposes of coordinated defense does not lead to inappropriate intrusions into personal activities.

As many of these represent cultural shifts, we will also need processes in place to effectively *train a workforce* that understands and can practice cyber defense in this new defense perspective.

5. Proposed Roadmap and Research Priorities

In this section, we propose a research roadmap to address the gaps identified above. Taken as a whole, we feel that this proposed research roadmap would create a foundation for applying scientific rigor the properties of asymmetry as they relate to cyber conflict. Ultimately, the desire is that disciplined scientific practices applied to this line of research would result in useful approaches with well-characterized benefit, cost, and impact.

5.1. Step 1: Rigorously Define Asymmetry for Cyber Resilience

The AiR participants generally agreed that a formal definition of asymmetry is required for meaningful research to commence on the topic. While we developed a working definition, it lacks the mathematical basis needed to drive scientific research. To refine our working consensus definition into such a formal definition would require defining the scope of asymmetric attributes relevant in the context of a cyber conflict. It would also require a refinement of the dimensions of asymmetry described above along with understanding of which dimensions are related to predictive outcomes of a conflict. Once a rigorous definition is established, measures and metrics for system state, resilience, and defensive posture will need to be developed and validated.

A key activity in defining asymmetry is to draw linkages between the goals of asymmetry and how they meet the needs of potential stakeholders. Cyber defenders, system planners and architects, and government policy makers will all need to understand the impact that asymmetric concepts have in shifting the way cyber conflict is carried out, in the context of both planning activities and real-time response tactics. Regardless of the context, the goal will be to increase effectiveness of resilient defenses through prioritization of resources.

5.2. Step 2: Examine Asymmetry in Other Domains

Once a formal definition for asymmetry in cyber resilience is developed, we recommend analyzing examples of asymmetry in other established application domains, such as physical warfare or cryptography, and emerging domains, such as quantum computing. The goal would be to determine if implementations in those domains have an analog in cyber resilience. This would require studies of a broad selection of other disciplines for features of asymmetry. In each case, the goal would be to identify needs and requirements for realizing asymmetry. This analysis should be performed in light of needs of both private and public cyber environments.

The capacity to scale will need special attention. While some techniques may provide localized advantage, complex cyber systems can be globally distributed and massive in scale, necessitating defensive technologies that scale.

5.3. Step 3: Develop Asymmetry-Specific Case Studies for Cyber Resilience

Retrospective studies. We need to understand the cyber domain of today so that we can develop paradigm shifts in the way defenders protect the cyber domain. One way to do this is to engage in retrospective studies. Such studies would yield substantial insight into existing needs of cyber defenders. One challenge here is that gaining this insight will require the defenders to share potentially large amounts of information relating to specific cyber conflicts. Several different types of conflict should be studied to sample a wide range of parameters of scale and sophistication. These should include 1) traditional cyber criminals engaging in activities like theft of IP and resources against private citizens; 2) nation state actors engaged in attempts to destroy or disrupt infrastructure of other nation states; and 3) citizens using social media to bypass filtering against a nation state.

Prospective studies. As a complement to retrospective studies, we also recommend prospective studies that combine the analysis of live capture data containing instances of cyber conflict with experimentation in synthetic environments. This would undoubtedly have a modeling and simulation component for which threat and response scenarios would need to be developed. Red teaming, tabletop exercises, and game theory are examples of methods to assess and study features of asymmetry in cyber conflict; this process ideally would include cases where not all behavior is rational.

5.4. Step 4: Develop Cyber Benchmarks, Measure, and Evaluate

The final recommendation for creating a scientific foundation for studying asymmetry for conflict in resilient cyber systems is to develop sharable benchmarks for validation. The goal would be to make it possible for all asymmetry research to be compared against state of practice using a repeatable experimental protocol. Emphasis placed on reproducibility of results will incentivize data sharing and methods. Results must be repeatable even under the scenario that the adversary knows of the methods being tested; we cannot rely on (asymmetric) security by obscurity. Not all states of systems can be measured, so the work in measures and metrics will need to draw linkages between states of interest and phenomena that are observable (or model-able).

6. Concluding Thoughts

As complex interconnected cyber systems continue to be driven by increasingly resilient operational requirements, it is important to have a fundamental understanding of the underlying principles of cyber resilience. Candidate resilience technologies will have varying degrees of effectiveness and varying costs for using them (where cost includes economic cost as well as any other resources that would be required to implement and sustain a resilience approach). In particular, it will be essential to understand the cost and benefit of defensive cyber technologies and techniques in light of their ability to provide an asymmetric advantage to defenders in order to maximize the effectiveness of cost applied toward winning cyber conflict. Toward this end, the AiR participants put forward the notion of asymmetric cyber advantage as a guiding principle for evaluating and selecting resilience technologies.

Our working definition of asymmetry in resilience is "a disproportionate, exploitable imbalance between actors related to, but not limited to, resources, level of effort, risk, or consequences in an attack." While this definition provides a framework for exploring the principles of asymmetry, we recognize that a more mathematically formal definition is needed as the foundation for scientific research in this area.

The participant identified research areas critical for developing a rigorous approach to asymmetry in cyber conflict. Metrics are needed for quantifying degree of resilience, security posture, state information in complex cyber systems, and cost/benefit for defenders and adversaries. Metrics and mission operations workflow modeling would provide a critical link between dynamic responses in resilient systems and awareness of the state of systems and assets in the context of how they support the mission that is served by the cyber system. Comprehensive (real-time) situational awareness augments the perspective that would be gained from metrics and workflow modeling with sensing of the state of systems and their functionality with the presumption that at any given time adversaries would have a foothold in the defender's cyber system. We see adversary characterization as an important component of asymmetry because insight into their techniques and motivations may lead to strategic decisions that shift the advantage back in favor of the defender. Ultimately, humans are behind the motivations of cyber systems and the conflicts that play out on them. As a result, we see understanding the many roles of humans in cyber conflict to be an essential component in the study of asymmetry in general. This includes humans as attacker and defender, but also as the terrain over which the battle is being fought and hence as (sometimes unwitting) facilitators of cyber conflict.

We propose developing a foundation for scientific research in cyber asymmetry that has four major components: 1) a rigorous definition for asymmetry that includes a sufficient level of mathematical precision for use as a guiding principle; 2) exploration of asymmetry in other domains and analysis of the applicability of these asymmetry concepts to cyber resilience; 3) case studies on cyber conflicts to form an empirical foundation for analysis of cyber conflict; and 4) benchmarks and evaluation methods to ensure repeatability and a basis for comparison across methods.

This shift in cyber culture, the movement from a posture of defensive reaction to one of proactive and reactive resilience, is as much a technical challenge as it is a social one, so its success as a long-term strategy will rely on an effective partnership between the research community, technology transition partners, cyber practitioners, policy makers, end users, and those charged with training a new generation of workforce who is well versed in asymmetry principles for resilience. The ultimate goal is to realize resilience in a way that provides asymmetric advantage for the complex cyber systems on which the U.S.

relies for its most fundamental activities. This will require a clear policy and legal framework that incentivizes decisions that align with this asymmetric vision. Asymmetry is a challenging aspiration, but an essential component for ensuring that resilience is achieved in a way that truly changes the nature of cyber conflict in favor of defenders.

7. Works Cited

Arreguin-Toft, Ivan. *How the Weak Win Wars: A Theory of Asymmetric Conflict.* Cambridge University Press, 2005.

Bidgoli, H. ed. *Handbook of INformation Security, Information Warfare, Social, Legal, and INternational Issues and Security Foundations.* Hoboken: Wiley & Sons, Inc., 2006.

Black, Jacquelyn G. *Microbiology: Principles and Explorations, 8th Edition.* Hoboken: Wiley & Sons., 2012.

Gordon, D.M., Pilko, A., De Bortoli, N., Ingram, K. K. "Does an ecological advantage produce the assymetric lineage ratio in a harvester ant population?" *Oecologia*, 2013: 849-857.

Hartman, William J. *Globalization and Asymmetrical Warfare.* Maxwell AFB: Air University, 2002.

Lazarevic, A., Ertoz, L., Kumar, V., Ozgur, A., and Srivastava, J. "A comparative Study of Anomaly Detection schemes in Network Intrustion Detection." *SIAM International Conference on Data Mining*, 2003: 25-36.

Metz, S., Johnson III, D. *Asymmetry and U.S. Military Strategy L Definition, Background, and Straetegic Concepts.* U. S. Army Strategic Studies Institute, 2001.

Mingo, A., "Size, uneven competition and resource availability: A factorial experiment on seedling establishment of three Mediterranean species." *Plant Biosystems - An International Journal Dealing with all Aspects of Plant Biology: Official Journal of the Societa Botanica Italiana*, 2009: 181-189.

National Institute of Standards and Technology. *Guide for Conducting Risk Assessments.* U.S. Department of Commerce, September 2012.

Rahme, J., Widmer, A., Karrenberg, S. "Polen competition as an asymmetric reproductive barrier between two closely related Silene species." *Journal of Evolutionary Biology*, 2009: 1937-1973.

Reward, B., Leuschner, C. "Does root competition asymmetry increase with water availability?" *Plant Ecology & Diversity*, 2009: 255-264.

Rushby, J. *Formal methods and their role in the certification of critical systems.* London: Springer, 1997.

Van Leeuwen, J. ed. *Handbook of theoretical computer science: Algorithms and complexity.* Cambridge: Elsevier Science Publishers, 1990.

Appendix A: Meeting Agenda

Day 1:

08:00 – 08:30	**Check in**
08:30 – 08:45	**Welcome and introductions**
08:45 – 10:20	**Level setting presentations** Dr. Chris Oehmen (PNNL) Mr. Charles Nelson (White House OSTP) Dr. Rich Colbaugh (Periander, Inc.) Dr. Greg Shannon (CMU/CERT) Dr. Tyler Moore (SMU)
10:20 – 10:30	**Break**
10:30 – 13:00	**Breakout session 1** Goal: Definition, attributes, state-of-the-art
12:00 – 13:00	**Working lunch**
13:00 – 13:30	**Group reports** 5 to 10 minute presentations summarizing morning discussion
13:30 – 16:00	**Breakout session 2** Goal: Gaps prioritized; begin discussing research challenges
16:00 – 16:30	**Group reports**

Day 2:

08:00 – 08:30	**Check in**
08:30 – 11:00	**Breakout session 3** Goal: Complete challenges discussion; key directions for the research agenda
11:00 – 11:15	**Break**
11:15 – 12:00	**Group reports** 15-minute presentations summarizing morning discussion
12:00 – 13:00	**Working lunch**
13:00 – 16:00	**Group strategy** Goal: Distill results into an agreed-upon research agenda; flesh out agenda with key domains and directions to address each gap
16:00 – 16:30	**Wrap up**

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by **Battelle** *Since 1965*

U.S. DEPARTMENT OF
ENERGY

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99352
1-888-375-PNNL (7665)
www.pnnl.gov