19th ICCRTS

Title of Paper:

Decision Making for Resilience within the Context of Network Centric Operations

---

Authors:

Zachary A. Collier
US Army Engineer Research & Development Center
3909 Halls Ferry Road, Vicksburg, MS 39180
601-634-7570
Zachary.A.Collier@usace.army.mil

Igor Linkov
US Army Engineer Research & Development Center
696 Virginia Road, Concord, MA 01742
978-318-8197
Igor.Linkov@usace.army.mil

Point of Contact: Igor Linkov

Topics:

Primary (Topic 1: Concepts, Theory, and Policy)
Alternate (Topic 10: Cyberspace Management)
Alternate (Topic 5: Experimentation, Metrics, and Analysis)

ABSTRACT: Recent calls from the US White House for enhanced resilience of our critical infrastructure in the face of persistent threats, (both natural and manmade), underscores the importance of developing and adopting a resilience-focused approach within individual communities, organizations, the DOD, and the Nation. However, the concept of resilience is still not well understood and varies across disciplines. In this paper, we study two proposed definitions of resilience, one from the National Academy of Sciences and one from the literature on Command and Control and Network Centric Operations. The convergences and divergences are explored between these two approaches to resilience (and by extension, related concepts such as agility). This paper proposes a decision making framework that integrates the event management cycle defined by the National Academy of Sciences into a resilience matrix that accounts for the physical, information, cognitive, and social domains in which these systems exist, as defined by Network Centric Operations. This systems-based approach can be used to comparatively assess the relative resilience of different systems and the contributions of individual responses or safeguards to overall system resilience.

1. INTRODUCTION

Individuals, organizations, communities, and nations are becoming increasingly dependent upon cyber infrastructure. This infrastructure, spanning hardware and software, cloud-based systems, and other information technology systems, supports practically all of the critical functions of our global society (e.g., finance, health care, defense). However, this cyber infrastructure is vulnerable to attacks and natural hazards, and can lead to failure of critical infrastructure, loss of sensitive information, and infringement of intellectual property (US White House, 2009). Internal flaws such as bugs, poor design, testing, quality assurance, and maintenance, can also lead to losses. Moreover, due to "ubiquitous connectivity" (Alberts, 2010), the highly networked cyber systems can result in losses that may cascade throughout multiple economic sectors and geographic scales (Rinaldi et al., 2001). For example, a cyber attack like the one against Aramco, Saudi Arabia's national oil company, had the potential to disrupt oil production for nations around the world (New York Times, 2012), causing downstream shocks throughout many other industry sectors (Kelic et al., 2013).

Cybersecurity is a critical national security concern, yet the US Department of Defense (DoD) is currently not poised at the required level of readiness against cyber threats (Defense Science Board, 2013). In response, calls for *resilience* against cyber threats have been made from the highest levels of government, for example in documents such as Executive Order 13636 (2013), Presidential Policy Directive 21 (2013), and the White House Cyberspace Policy Review (2009).

However, the concept of resilience is still widely debated among practitioners and theorists in different fields, and is often confused with related but distinct concepts such as risk, robustness, and vulnerability. Generally speaking, one can define resilience as "*an ability to recover from or adjust easily to misfortune or change*" (Merriam-Webster Dictionary, 2013), but this definition may be nuanced depending on the particular application area. For instance, scholars have distinguished between "engineering" resilience and "ecological" resilience (Holling, 1996; Walker et al., 2004; Gallopín 2006; Park et al., 2012), each being applicable in different areas of research. Being distinct concepts, risk and resilience thus require different management approaches (Linkov et al., 2014).

Apart from difficulties in defining resilience, another difficulty has emerged in how resilience may be measured. NAS (2012) stresses the importance of metrics in their report, stating that a numerical basis for assessing resilience is required for monitoring changes and charting improvement. Some have propose quantitative methods for measuring resilience based on time to recovery and loss of performance (e.g., Schultz et al., 2012), however the difficulty tends to lie in translating the loss in performance to the value that one puts on that performance within the context of some higher level mission or objective. Linkov et al. (2013a) take a more general approach and base the generation of metrics on the National Academy of Sciences (NAS) definition that incorporates the event management cycle (NAS, 2012) coupled with the recognition stemming from Network Centric Operations (NCO) doctrine that cyber systems (and socio-technical systems in general) span multiple, interconnected domains (Alberts, 2002). Linkov et al. (2013b) use this approach to generate a process for identifying cybersecurity metrics.

In this paper, we will explore, in more detail, the NAS and NCO approaches to resilience and outline areas in which they converge and diverge. In particular, the NAS definition will be examined within the context of the broader concept of *agility* as defined in the command and control literature (Alberts, 2011), and how concepts from the field of decision analysis can be applied to bridge some of these gaps.

2. DEFINITIONS OF RESILIENCE

In their report on resilience of communities against disasters, NAS defines the notion of resilience as "*The ability to prepare and plan for, absorb, recover from, or more successfully adapt to actual or potential adverse events"* (NAS, 2012). This definition stresses the four actions that a community (or a system in general) may take to enhance resilience: plan/prepare, absorb, recover, and adapt. Whereas absorption, recovery, and adaptation take place after an adverse event has taken place, the planning and preparation stage is anticipatory, occurring before an adverse event.

The definition of resilience stemming from NCO is: "*Resilience provides an entity with the ability to repair, replace, patch, or otherwise reconstitute lost capability or performance (and hence effectiveness), at least in part and over time, from misfortune, damage, or a destabilizing perturbation in the environment*" (Alberts, 2011). This definition has many commonalities with the one above, for instance, an emphasis on actions that a system may take (repair, replace, patch, reconstitute) in response to an adverse event. These four responses mainly relate to the recovery of a system after an adverse event.

One aspect that the NCO definition captures that the NAS definition does not is the focus on *why* we care about resilience – namely the capability, performance, or effectiveness of the system that we want to be resilient. Implicit in this is that there is a valued function that the system provides, and thus someone applying their values to the system – the observer of the system is an inextricable component of the system itself. In addition, this definition acknowledges that resilience is meaningless without a consideration of recovery time to reach the minimum acceptable operating threshold.

However, Alberts notes that resilience is only one component of a larger, more broadly encompassing theme termed "agility", which is "*the ability to successfully effect, cope with, and/or exploit changes in circumstances*" (Alberts, 2011). In this view, changes in circumstances need not be adverse to the system, but instead can include circumstances that are beneficial and provide opportunities. Thus, while traditional concepts of resilience are focused solely on returning performance back to "normal", agility stresses that non-adverse changes can be exploited to improve performance. To fully realize agility, an entity or system must possess resilience, as well as responsiveness, versatility, flexibility, innovativeness, and adaptability. Table 1 highlights some of the differences in definitions of resilience between these two paradigms.

Table 1: Comparison of Approaches to Resilience

Comparison	**National Academy of Sciences**	**Network Centric Operations**
Resilience *of what?*	Communities at various scales	Mainly military organizations, but generalizable to other systems
Resilience *to what?*	Adverse events, especially disasters. Emphasizes an "all-hazards" approach.	Destruction, interruption or degradation of a capability
Goal of resilience	As an end in of itself	As a means to enable Agility
Anticipatory?	Yes	Resilience coupled with Responsiveness can be anticipatory
Accounting for performance thresholds?	Unclear	Yes, an acceptable range of performance exists
Accounting for time to recovery?	Unclear	Yes
Stakeholder values incorporated?	Yes, in the risk management process along with goal and objective identification	Yes, in the definition of an entity's measure of value/desired state
How much resilience is needed?	Benefits in resilience investments must balance or outweigh the costs (costs and benefits may be non-monetary and/or qualitative)	For agility, the goal is to achieve "requisite agility", which is optimized based on the probability and cost/benefits of adverse/positive events.
Metrics for resilience	No comprehensive set of metrics explicitly proposed. Recommends the development of a national resilience scorecard.	*"Estimated proportion of problem space in which adequate C2 capability remains after degradation and in which timely and relevant restoration is possible"* (McEver et al., 2008)

In particular, Alberts (2011) notes that for active agility, resilience (or versatility, flexibility, innovativeness, adaptability) must be coupled with responsiveness, which explicitly takes into account the time it takes to respond to an adverse event or to anticipate and take a proactive measure. The process of responsiveness begins with a change in circumstances (e.g., a cyber attack), followed by detection of the change, a decision regarding a course of action, the execution of that action, and a time lag to reach the desired effect. This event management cycle is much different than the one proposed by NAS (Figure 1). In Figure 1, time spans from left to right, and the red delta symbol represents the adverse change in circumstances.

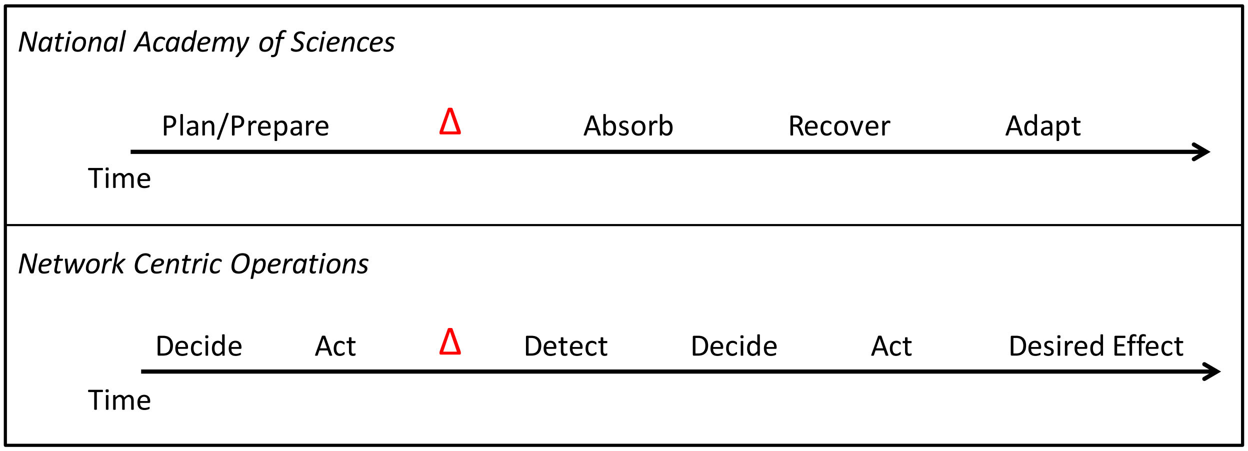

Figure 1: Comparison of Event Management Cycles

3. DECISION MAKING: THE COMMON THREAD

How then can these two differing definitions of resilience be reconciled? First, one must more closely examine the natures of these differing event management cycles. Table 2 lists each of the "steps" associated with both of the event management cycles shown in Figure 1, with an associated definition in terms of what type of action that step represents.

Table 2: Typology of Event Management Cycle Steps

Source	Event Management Cycle	Type of Action
NAS	Plan/Prepare	Response (Proactive)
NAS	Absorb	Response
NAS	Recover	Response
NAS	Adapt	Response
Both	Δ (Adverse Event)	State of Nature
NCO	Detect	Perception
NCO	Decide	Selection of Response
NCO	Act	Response
NCO	Desired Effect	State of Nature + Perception

Once can see that, with the exception of the adverse event, which is common to both approaches and necessary for the definition of resilience, the NAS definition is comprised entirely of steps that are *responses* to the event. Absorb, Recover, and Adapt are strictly post-event responses, while Plan/Prepare is a response (consisting of the Decide and Act steps) one can take to a perceived or anticipated future event. In terms of NCO, the event management cycle encompasses a broader picture, in which an event is first perceived, then a response is selected, and the response is executed. The final step, Desired Effect, is a combination of the response having made a change (or possibly having made no change) in the world (e.g., system performance), and the perception or detection of the effects of the response. The relationship between these two approaches is that NCO uses the generic term of Act to describe the post-event response to an adverse event, while NAS enumerates *specific mechanisms by which one can act or respond* (Figure 2). Similarly, NAS combines the pre-event steps of Decide and Act into the Plan/Prepare step in which one anticipates and executes an anticipatory course of action.

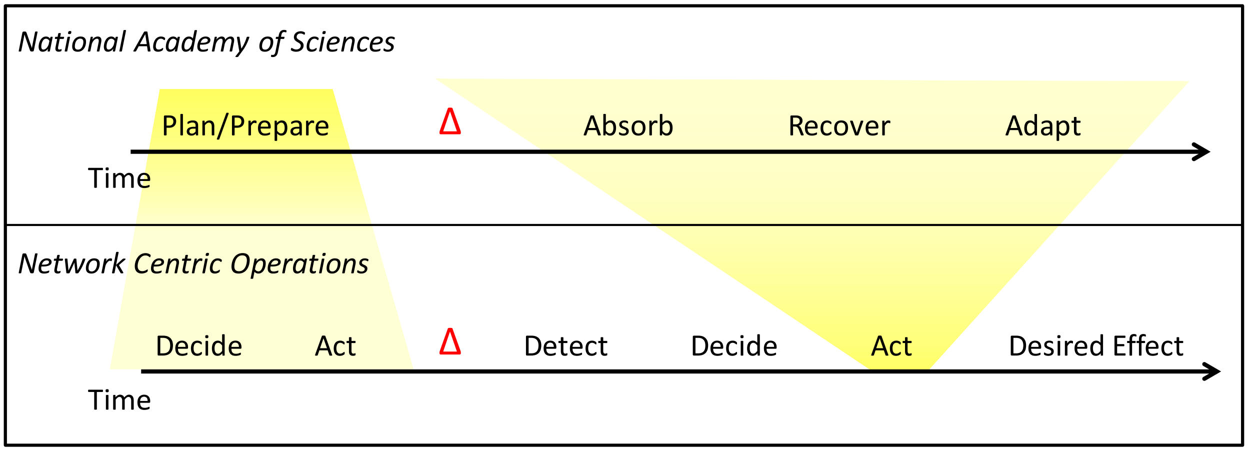

Figure 2: Relationship between Event Management Cycles

NCO uses the generic term of Act to denote a response to an adverse event but, especially given the multiple types of responses (Plan/Prepare, Absorb, Recover, and Adapt), there is not clear guidance on how one should respond to most effectively reach a Desired Effect. After all, the response taken may have a negative effect or no effect at all. Moreover, there are multiple ways in which one can, for instance, recover from a cyber attack. There is a need for a clear methodology to select from multiple anticipatory and reactive response alternatives that best ensure a posture that will minimize negative consequences and maximize desired effects.

The answer lies within the Decide step. In particular, insights from the field of Decision Analysis (Raiffa, 1968; Clemen, 1996) can improve the probability of achieving desired effects. In Decision Analysis, there is a distinction between a *good decision* and a *good outcome*, namely that one cannot guarantee a good outcome (i.e., desired effects), but through careful and insightful problem framing, data collection, and sound analysis, one can increase the odds of achieving a good outcome by making a good decision (Howard, 1988). Specifically, tools within the field of Multi-Criteria Decision Analysis (MCDA) (Belton & Stewart, 2002; Linkov & Moberg, 2012) can provide decision makers with the necessary guidance to select from among multiple alternative courses of action based on physical data (e.g., monitoring data, simulations, costs) and subjective value judgments (e.g., risk tolerance, priority of performance criteria). In particular, MCDA approaches follow a sequence where the decision maker must 1) Identify goals and objectives, 2) Identify the available alternatives (in this case responses), 3) Identify criteria and sub-criteria relevant to the level of achievement of the goal, 4) Assign relative weights to each of the criteria and sub-criteria in terms of their importance, 5) Score each alternative in terms of its performance along each of the criteria and sub-criteria, 6) Synthesize the scores and weights to select the preferred alternative and perform a sensitivity analysis.

Figure 3 depicts how a particular response can be chosen from among several alternatives after an adverse event has been detected. In this case, there exists a numerous set of potential alternative responses one can select, including various ways to absorb, recover from, and adapt to a disturbance. Assuming only one can be chosen, the second recover option ("Recover 2") is selected in the Decide step, and carried out in the Act step. Following the execution of the selected alternative, one can check whether the desired effect (i.e., acceptable level of restored performance) has been achieved. If not, a new alternative may be chosen iteratively until that acceptable level has been reached.

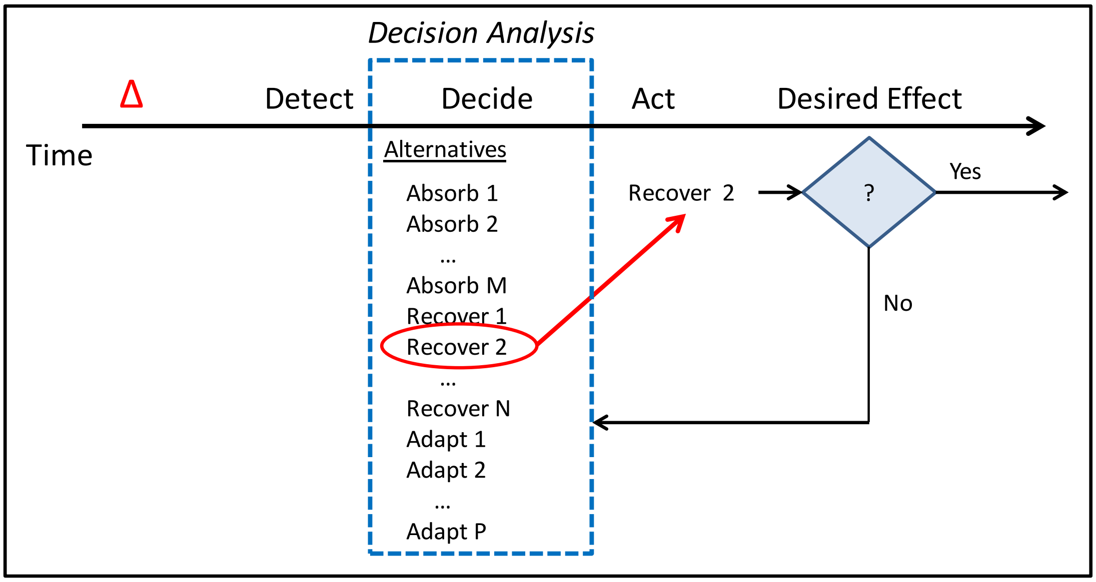

Figure 3: Schematic of Reactive Decision

As Alberts (2011) notes, reactive decisions alone are insufficient to ensure agility (and resilience), and that anticipatory decisions can be made before an adverse event that may mitigate damage, buy extra time, or even preempt the event altogether. This is shown in Figure 4, where an initial decision is made selecting from among several Plan/Prepare responses. When an adverse event occurs, and the desired effect is achieved, then the process ends. If however, a desired effect is not achieved, then a new round of decision making can occur.

Within the Decide step, there must be some internal mechanisms that allow for the comparison between available alternatives. Figure 5 illustrates the comparative assessment of alternatives using the resilience matrix approach developed by Linkov et al. (2013a,b). In this approach, the columns represent the responses found in the NAS report. The rows represent another concept from NCO - the four operational domains in which cyber systems (or other types of systems) exist, namely within the physical, information, cognitive, and social domains (Alberts, 2002). Together, these two aspects construct a 4x4 matrix. Each cell thus represents the system's ability within that domain to execute the particular response, and thus manage adverse events. For example, in Figure 5, Alternative 1 might represent an enhancement to a system's ability within the physical domain to recover from a cyber attack (e.g., a backup electrical generator to ensure the continued functionality of critical electronic equipment in the case of a power outage). Alternatives may then be compared by the degree in which they enhance resilience of a system, but also by their costs. In Figure 5, Alternative 1 has much greater benefits to resilience than the combination of Alternatives 2 and 3, but is also more costly. Thus, alternatives can be assessed and compared in terms of costs and benefits so that resources may be allocated in an efficient manner.

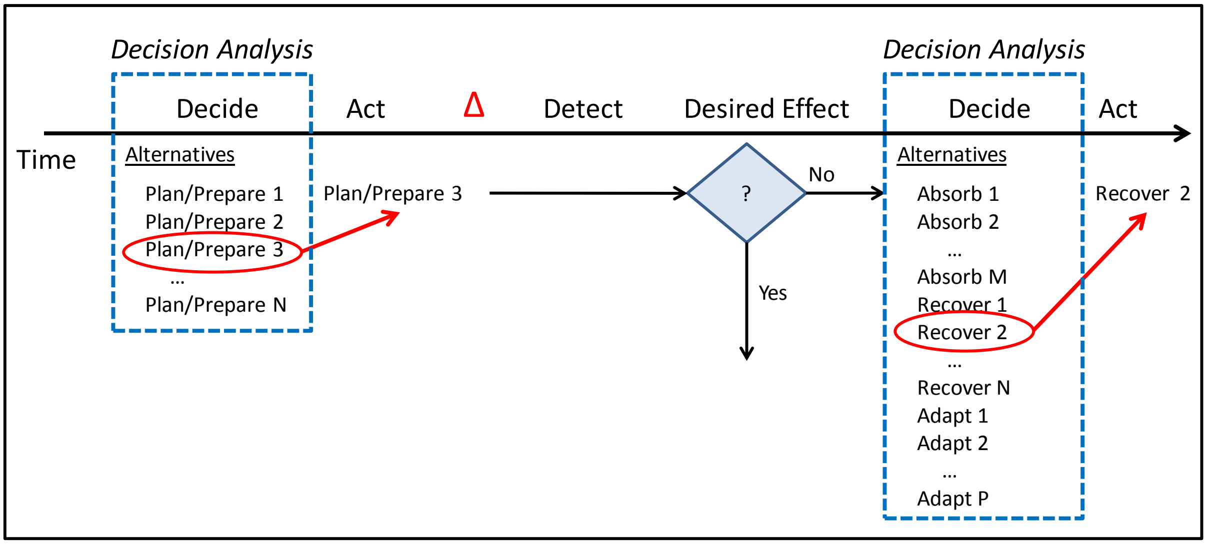

Figure 4: Schematic of an Anticipatory Decision

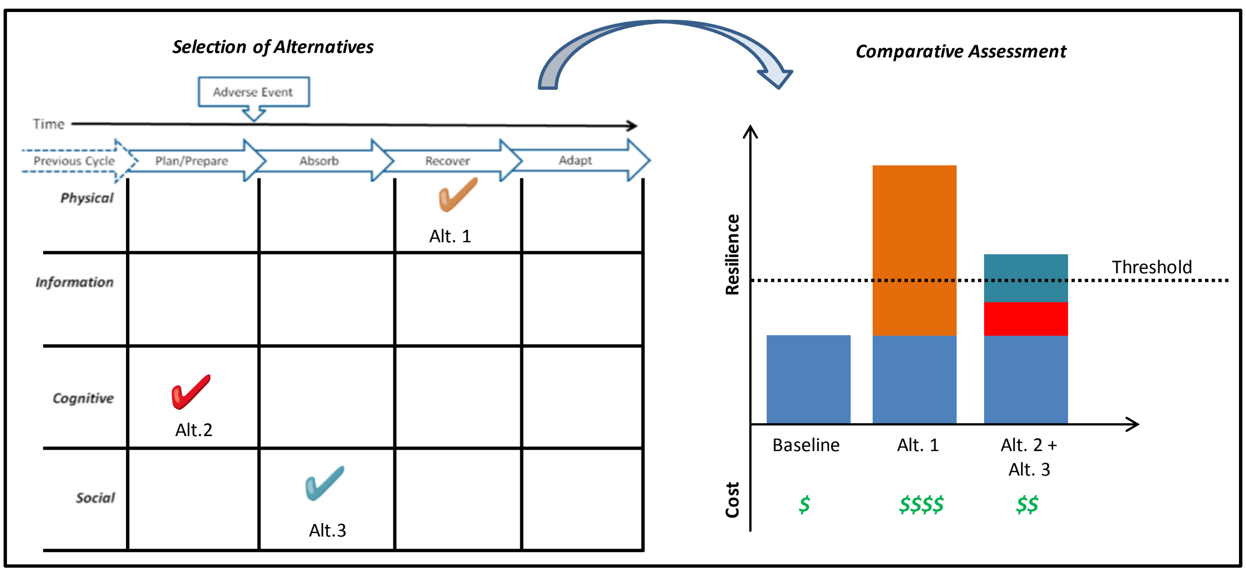

Figure 5: Comparative Assessment of Resilience-Enhancing Alternatives

In addition, it is important to note that in practice, typically more than one alternative is selected. For example, the combination of Alternatives 2 and 3 were chosen in Figure 5. This implies the need for a portfolio-based decision process (Salo et al., 2011). Taking a portfolio-focused approach, a decision maker can “mix and match” combinations of alternatives that allow for the maximum enhancements in resilience and/or agility for a given cost. It also allows for a deeper understanding of what combination of alternatives may dominate or be dominated by other, more efficient combinations.

Finally, a structured decision making methodology provides a platform for adaptive management. Adaptive management, first proposed in the environmental field (Holling, 1978; Walters, 1986), is a way to make downstream decisions towards achieving a desired goal under uncertainty and when confronted with new information. Traditional adaptive management approaches allow decision makers to adapt project activities in light of new information or changing conditions within the operational environment. Over time, as new information becomes available through monitoring, new actions may be taken to more effectively course-correct towards the desired goal state. Convertino et al., (2013) extend traditional adaptive management to explicitly incorporate structured decision making models. This extension, termed enhanced adaptive management (EAM), updates the inputs of a multi-criteria decision model as new monitoring information becomes available, and re-ranks alternative courses of action given this new information. If the new information causes the change in the optimal course of action, this new alternative is selected, and the process is repeated. As it relates to resilience, EAM can be explicitly linked to the conceptual models for resilience decisions, shown in Figure 6.

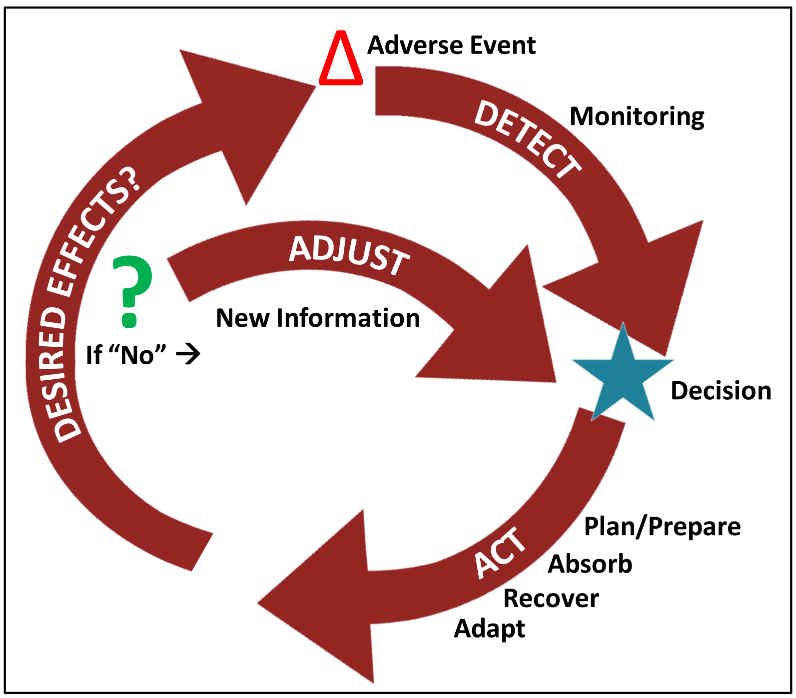

Figure 6: Enhanced Adaptive Management for Resilience (adapted from Jones, 2009)

Once a decision is made to select a particular response, the particular response alternative is executed. If the results do not produce the desired effect (as evidenced through monitoring), a feedback loop back to the decision stage exists. Given the new state of the world in which a response was unsuccessfully made, other new monitoring information, and possibly new preference information (e.g., increased urgency to recover functionality), a new response alternative can be selected. Thus, the MCDA decision model and EAM work in parallel to update selected courses of action as new events occur and conditions change.

4. CONCLUSIONS

Given the complexity of cyber (and other socio-technical) systems, designing agile (and thus resilient) systems is a daunting task. As new safeguards are developed and implemented, adversaries continue to develop novel ways to breach information technology systems, steal sensitive data, and disrupt critical infrastructure. While significant advances in the field of cybersecurity have been achieved, solutions tend to focus on the technical issues at component levels such as threat detection, encryption, and other mitigation procedures and technologies and not on how to manage cyber risk and make decisions at system level. Confusion over the meanings of terms like resilience and risk has further hindered this progress.

Ultimately the ability of cyber systems to be resilient (i.e., maintain an acceptable level of performance) in an uncertain and risky environment rests upon the ability of decision makers and planners to make good decisions. Resilience thus cannot be ensured based on ad hoc decision making alone – structured tools are necessary to aid in making sense of relevant information, uncertainties, and preferences. There is a critical need to approach cybersecurity risks from a systems perspective, recognizing the complex interactions between cyber, physical, and human systems. Decision aiding tools, such as the ones offered by the field of decision analysis, can hold the key to future success and superiority in cyber operations within the physical, information, cognitive, and social domains. More generally, these ideas of resilience and agility transcend many types of systems and threats, and are therefore not exclusive to just cyber systems. Indeed, a broad array of systems and threats can be considered using these principles.

Moving forward, several parallel efforts must be pursued. First, diverse communities of researchers and practitioners need to work across discipline boundaries to develop a coherent and consistent terminology with which to discuss systems. This will reduce confusion in definitions and provide a common language to explore ideas about risk, resilience, robustness, agility, etc. Second, the development of specific resilience metrics is necessary. NAS (2012) acknowledges the difficulties in generating metrics based on issues of geographic scale, time frame, and community priorities. In addition, careful selection of metrics is critical, since choosing the wrong metrics may likely lead to unintended suboptimal results (Williamson, 2006). Third, effort must be made towards the development of decision support tools that utilize the previously mentioned language and metrics. Fourth, refinements in theory, metrics, and decision aids must be made iteratively through testing and experimentation on real world systems. Finally, resilience-based thinking must be institutionalized within organizations through education, training, and outreach.

5. ACKNOWLEDGEMENTS

The authors would like to thank David S. Alberts for his helpful comments on the manuscript. Permission was granted by the USACE Chief of Engineers to publish this material. The views and opinions expressed in this paper are those of the individual authors and not those of the US Army, or other sponsor organizations.

6. REFERENCES

Alberts DS (2002) *Information Age Transformation: Getting to a 21st Century Military*. DOD Command and Control Research Program: Washington, DC. http://www.dtic.mil/get-tr-doc/pdf?AD=ADA457904

Alberts DS (2010) The Agility Imperative: Précis. DOD Command and Control Research Program: Washington, DC. http://www.dodccrp.org/files/Alberts_Agility_Imperative_Precis.pdf

Alberts DS (2011) *The Agility Advantage: A Survival Guide for Complex Enterprises and Endeavors*. DOD Command and Control Research Program: Washington, DC.

Belton V, Stewart T (2002) *Multiple Criteria Decision Analysis: An Integrated Approach*. Springer: Boston.

Clemen RT (1996) *Making Hard Decisions: An Introduction to Decision Analysis, 2nd Ed.* Duxbury Press: Boston.

Convertino M, Foran CM, Keisler JM, Scarlett L, LoSchiavo A, Kiker GA, Linkov I (2013) Enhanced Adaptive Management: Integrating Decision Analysis, Scenario Analysis and Environmental Modeling for the Everglades. *Scientific Reports* 3:2922, DOI: 10.1038/srep02922.

Defense Science Board (2013) Task Force Report: Resilient Military Systems and the Advanced Cyber Threat.

Executive Order 13636 (2013)—Improving Critical Infrastructure Cybersecurity. http://www.gpo.gov/fdsys/pkg/FR-2013-02-19/pdf/2013-03915.pdf

Gallopín GC (2006) Linkages between vulnerability, resilience, and adaptive capacity. *Global Environmental Change* 16(3): 293-303.

Holling CS (1978) *Adaptive Environmental Assessment and Management*. John Wiley and Sons: New York, NY.

Holling CS (1996) Engineering Resilience versus Ecological Resilience. In: Schulze P (Ed.), *Engineering Within Ecological Constraints*. National Academy Press: Washington DC, pp. 31–44.

Howard RA (1988) Decision analysis: practice and promise. *Management Science* 34(6): 679-695.

Jones G (2009) The adaptive management system for the Tasmanian Wilderness World Heritage Area — linking management planning with effectiveness evaluation. In: Allan C, Stankey G (Eds.), *Adaptive Environmental Management. A Practitioners Guide*. Springer: Dordrecht, The Netherlands.

Kelic A, Collier ZA, Brown C, Beyeler WE, Outkin AV, Vargas VN, Ehlen MA, Judson C, Zaidi A, Leung B, Linkov I (2013) Decision Framework for Evaluating the Macroeconomic Risks and Policy Impacts of Cyber Attacks. *Environment Systems & Decisions* 33(4): 544-560.

Linkov I, Moberg E (2012) *Multi-Criteria Decision Analysis: Environmental Applications and Case Studies*: CRC Press: Boca Raton, Florida.

Linkov I, Eisenberg DA, Bates ME, Chang D, Convertino M, Allen JH, Flynn SE, Seager TP (2013a) Measurable resilience for actionable policy. *Environmental Science & Technology* 47(18):10108–10110.

Linkov I, Eisenberg DA, Plourde K, Seager TP, Allen J, Kott A (2013b) Resilience metrics for cyber systems. *Environment Systems & Decisions* 33(4): 471-476.

Linkov I, Bridges T, Creutzig F, Decker J, Fox-Lent C, Kröger W, Lambert JH, Levermann A, Montreuil B, Nathwani J, Nyer R, Renn O, Scharte B, Scheffler A, Schreurs M, Thiel-Clemen T (2014) Changing the Resilience Paradigm. *Nature Climate Change,* forthcoming.

McEver JG, Martin DM, Hayes RE (2008) Operationalizing C2 Agility: Approaches to Measuring Agility in Command and Control Contexts. Presentation at the 13th ICCRTS. http://www.dodccrp.org/events/13th_iccrts_2008/presentations/255.pdf

Merriam-Webster (2014) *Resilience*. http://www.merriam-webster.com/dictionary/resilience

National Academy of Sciences (2012) Disaster Resilience: a National Imperative. National Academic Press: Washington, DC. http://www.nap.edu/catalog.php?record_id=13457

New York Times (2012) Aramco says cyberattack was aimed at production. http://www.nytimes.com/2012/12/10/business/global/saudi-aramco-says-hackers-took-aim-at-its-production.html

Park J, Seager TP, Rao PSC, Convertino M, Linkov I (2012) Integrating risk and resilience approaches to catastrophe management in engineering systems. *Risk Analysis* 33(3): 356-367.

Presidential Policy Directive 21 (2013) —Critical Infrastructure Security and Resilience. http://www.whitehouse.gov/the-pressoffice/2013/02/12/presidential-policy-directive-critical-infrastructuresecurity-and-resil.

Raiffa H (1968) *Decision Analysis*. Addison-Wesley: Reading, MA.

Rinaldi S, Peerenboom J, Kelly T (2001) Identifying, understanding, and analyzing critical infrastructure interdependencies. *IEEE Control Systems Magazine* 21(6):11–25.

Salo A, Keisler J, Morton A (eds.) *Portfolio Decision Analysis*. Springer.

Schultz MT, McKay SK, Hales LZ (2012) The Quantification and Evolution of Resilience in Integrated Coastal Systems. Technical Report ERDC TR-12-7 U.S. Army Engineer Research & Development Center: Vicksburg, MS.

Walker B, Holling CS, Carpenter SR, Kinzig A (2004) Resilience, adaptability and transformability in social-ecological systems. *Ecology and Society* 9(2):5.

Walters CJ (1986) *Adaptive Management of Renewable Resources*. Blackburn Press: Caldwell, NJ.

Williamson RM (2006) What Gets Measured Gets Done: Are You Measuring what Really Matters? Strategic Work Systems, Inc.: Columbus, NC. https://www.swspitcrew.com/articles/What%20Gets%20Measured%201106.pdf

US White House (2009) Cyberspace policy review: assuring a trusted and resilient information and communications infrastructure. http://www.whitehouse.gov/assets/documents/Cyberspace_Policy_Review_final.pdf.