A STAMP-Based Approach to Developing Quantifiable Measures of Resilience

Paul M. Beach[1], Robert F. Mills[2], Brandon C. Burfeind[2], Brent T. Langhals[1,] Logan O. Mailloux[1]
[1]Dept. of Systems Engineering, Air Force Institute of Technology, Wright-Patterson AFB, OH
[2]Dept. of Electrical & Computer Engineering, Air Force Institute of Technology, Wright-Patterson AFB, OH

Abstract - *The quality of resilience is a desirable attribute in today's complex cyber-physical systems, but there is little consensus on what constitutes a suitable metric for resiliency. This work seeks to build upon an existing method for developing suitable resiliency metrics for complex cyber-physical systems. Specifically, several definitions of resilience are presented and their applicability to quantifiable measures of resilience is discussed. Next, methods for identifying and evaluating the impact of disruptive events on a system of interest and the development of resilience strategies is discussed. Finally, a detailed case study demonstrating a systems-based approach for the development and analysis of quantifiable measures of resiliency is presented.*

Keywords: Cyber-physical systems, resilience, STAMP, STPA-Sec, systems engineering

1 Introduction

Today's society is increasingly reliant upon complex cyber-physical systems to support everything from critical infrastructure to national defense, and our dependence on these systems makes them attractive targets for our adversaries. Despite their criticality, the U.S. Department of Defense (DoD) has concluded that many of their vital systems lack the resilience to operate through attacks from sophisticated adversaries [1]. This challenge is faced across critical infrastructure and industry sectors as well. Making today's systems and designing tomorrow's systems to be more resilient requires a systematic approach for developing and evaluating quantitative measures of resiliency to support engineering design choices.

Cyber-physical systems (CPSs) are computers and networks which control "physical processes… with feedback loops where physical processes affect computations and vice versa" [2]. Given these interdependences, a systems approach is well suited to the analysis and development of resiliency metrics for these systems. Thus, existing approaches such as the Systems-Theoretic Accident Model and Process (STAMP) and Systems Theoretic Processes Analysis for Security (STPA-Sec) are evaluated as to their appropriateness for use in analyzing causal events that may disrupt these systems and thus, negatively impact their resiliency.

This paper seeks to provide the Systems Engineering (SE) practitioner with a method for developing quantifiable metrics of resilience for various Systems of Interest (SoI). These metrics can be used during design phases to inform system architectural and engineering decisions, or to propose changes to or develop resilience strategies for existing systems. More concretely, we aim to expand upon an existing method for computing resiliency metrics by applying a systems approach to identifying cases which may degrade resilience in order to support the development of a quantitative evaluation of resiliency in complex CPSs.

The remainder of the paper is organized as follows. In Section 2, existing approaches for defining resilience are compared. Section 3 presents an existing approach for computing resilience metrics and an extension based on STPA-Sec is suggested. Section 4 presents a case study demonstrating this approach, and the paper is concluded in Section 5.

2 Defining Resilience

While frequently desired, it is observed in the literature that the lack of a clear definition of "resilience" (or "resiliency") has hampered efforts to create quantifiable measures of resilience. Furthermore, many attempts to describe resilience offer definitions that overlap with related (but distinctly different) concepts including robustness, fault-tolerance, flexibility, survivability and agility [3-9]. For example, it is possible for a system to be robust but not necessarily resilient, and vice versa (see [6] for an excellent discussion). Note, this is hardly a new debate, as Holling's seminal article "Resilience and Stability of Ecological Systems" (1973) was the first to provide a systems-focused definition of resiliency [10]. Several definitions of resilience are provided in Table 1, and some include this overlap. While these are all valuable traits, they lie outside a more specifically scoped definition of resilience, and efforts to merge broader understandings of resilience under a unified concept increases the difficulty of creating a single "resilience" metric.

Instead of creating an unnecessarily complex "one-size-fits-few" measurement, we believe each of these aspects should be assessed and leveraged as individual metrics for the analysis of complex systems [6],[11]. More precisely, we suggest that attempts to define and quantify resiliency should be restricted to the narrower definition of "bouncing back" [3], [12], which is the focus of this paper.

Additionally, there are engineering trade-offs between availability, integrity and confidentiality. These priorities are often competing and can complicate the definition of resilience in a system [6], [13]. For example, consider a military command and control system that has a secure, ground-based communications link disabled in an enemy attack. Decision makers might choose to rely on a less secure wireless communications link to ensure uninterrupted dissemination of

Table 1: A Sample of the Diversity of Perspectives Found in the Literature for Defining Resilience.

Author(s)	**Definition provided**
Henry & Ramirez-Marquez	The ratio of recovery at time t to loss suffered by the system at some previous point in time t_d, as indicated by Я(t) = Recovery(t)/Loss(t_d)$ (2012) [3].
Francis & Bekera	An endowed or enriched property of a system that is capable of effectively combating (absorbing, adapting to or rapidly recovery [sic] from) disruptive events [4].
Griffor et al.	Concerns related to the ability of the CPS to withstand instability, unexpected conditions, and gracefully return to predictable, but possibly degraded, performance [34].
Holling	A measure of the persistence of systems and of their ability to absorb change and disturbance and still maintain the same relationships between populations or state variables [10].
Nan & Sansavini	The ability of a system to resist the effects of a disruptive force and to reduce performance deviations [35].
Jackson	The ability of organizational, hardware and software systems to mitigate the severity and likelihood of failures or losses, to adapt to changing conditions, and to respond appropriately after the fact [36].

instructions to subordinates, but the adversary now has the opportunity to monitor these transmissions. The ability to communicate has been restored, but is this considered resilience? Practitioners looking to improve "resilience" in their systems should weigh and appropriately prioritize the importance of each of these components in their analysis. Note, this is not necessarily an easy undertaking, as we also lack clear metrics for comparing aspects of confidentiality, integrity and availability amongst one another [6].

Most of the approaches surveyed explicitly recognized there is an inherently temporal aspect to resilience—something must occur to the system and the system must then respond for the notion of resilience to ever come into play. It follows that the results of this process are a variable function of time, and any useful measurement of resilience should include time as an input parameter [3], [4], [8]. Fig. 1 provides a notional example of an engineering resilience curve for a time-dependent system performance metric $F(t)$. This metric could represent any relevant aspect of system performance, such as average response times for a web server. At $F(t_0)$, the SoI is operating at its nominal, pre-disruption state. At t_e, the system experiences a disruptive event and the system's performance begins to degrade until it reaches its final disrupted state $F(t_d)$ at time t_d. The SoI exists in the disrupted (or "vulnerable") state until time t_s, when one or more resilience actions are applied to increase the figure-of-merit until it reaches its stable recovered state (or "recovered steady state") value of $F(t_f)$ at time t_f. A few items to note are that it is possible that the time between each state is small or even zero (e.g., the degradation occurs immediately following the onset of the disruptive event, with no observable time difference between t_e and t_d). It is also possible that $F(t_f)$ is not the same as $F(t_0)$—i.e., after applying the resilience action(s), the SoI is performing at a lower or a higher level than it was prior to the disruptive event [3], [6] [14]. That said, for the purposes of quantifying resiliency, any recovery should reach some sustainable steady state that meets some minimum level of performance [15].

Having discussed several considerations for what constitutes a suitable resilience metric, we now discuss an approach for developing quantifiable measures of resilience.

3 Aspects of Resilience

Of the methods surveyed, an approach proposed by [3] provides an excellent groundwork upon which a model of resiliency metrics can be built. In this work, the authors formulate a process to compute quantifiable measures of resilience, which is comprised of the following four elements: (1) identify figure(s)-of-merit, (2) enumerate disruptive events, (3) identify resilience actions, and (4) compute resilience analysis.

This method incorporates a definition of resilience that is predicated upon measuring the deviation of one or more system performance metrics for any given disruptive event. This framework establishes a generic yet implementable methodology for measuring resilience which can be applied broadly to a wide range of systems. The precise methods for completing each of these tasks will naturally vary depending on the context in which they are applied, however a primary goal of our work is to extend [3] with a practical methodology of how this can be accomplished for a given complex cyber-physical SoI.

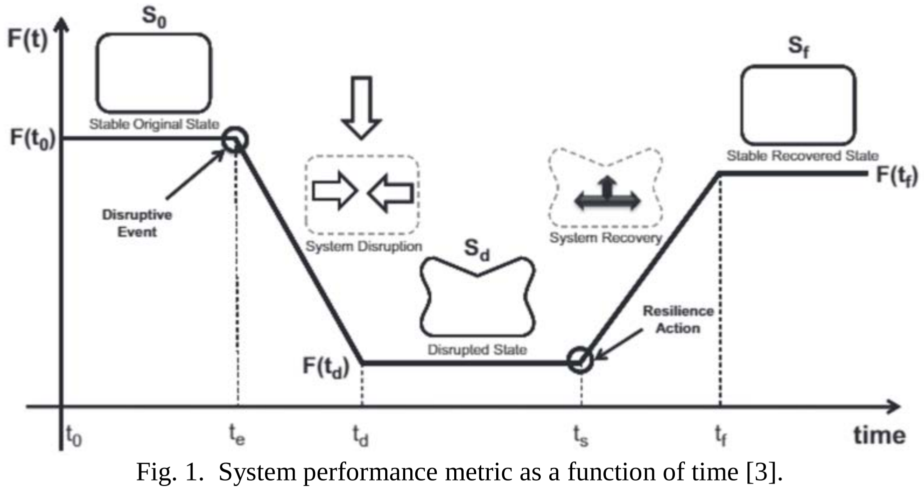

Fig. 1. System performance metric as a function of time [3].

3.1 Figure(s)-of-Merit

Termed “Figures-of-Merit” (FoM), the foundation of any resiliency analysis must be built upon the unambiguous identification of one or more “quantifiable and time-dependent system-level delivery functions” [3]. Alternative approaches refer to this aspect as the critical function, performance, or quality of the system [8]. These metrics serve as the baseline from which the deviation due to disruptive events is measured. In [3], a network flow problem is presented using shortest-path, maximum flow and overall network health as three possible metrics, however there is little elaboration or guidance on the selection of a suitable FoM beyond the requirement that it must be quantifiable. Therefore, we offer the following commentary to assist in the selection of such metrics.

When considering the kinds of metrics available there are three broad classes of metrics: technical, cost and organizational [16]. Examples of technical metrics often include the amount of available bandwidth or latency between systems or the percentage of functioning nodes in a system [8]. Cost metrics may include the cost of building redundancies into a system or repairing damage and can be used in conjunction with the technical metrics to evaluate engineering trade-off considerations when designing resilient systems. Likewise, organizational metrics consider processes that affect resilience, and may also be considered depending on where the system boundary lies and the scope of the analysis desired. Examples of typical organizational resilience metrics include employee training and awareness, processes and procedures, staffing levels or stakeholder involvement [17]. Given their nature, organizational resilience metrics may be difficult to quantify or build into a model-based simulation, but this alone does not preclude their use.

For existing systems, a survey of available metrics may be a good starting point since utilizing existing data is typically cheaper and easier to collect than gathering new data. However, one pitfall to avoid is limiting analysis to only the metrics available at hand (due to ease of use/collection) at the expense of mindfully identifying and seeking out appropriate metrics that would best satisfy the questions at hand. System operators and subject matter experts should be consulted to identify both useful existing metrics as well as which new metrics need to be collected. Identifying suitable metrics is not a trivial task, and there is growing evidence to suggest that the choice of metrics reflect the underlying priorities of an organization. Accordingly, these metrics should be deliberately defined to align organizational needs (or objectives) with desired system functionality [18].

A final consideration when selecting metrics is that using too many metrics may become expensive and/or challenging to gather, interpret and/or maintain. Thus, in the final analysis the practitioner must balance the trade-offs between the costs of using these measures and their ultimate benefit to the analysis.

3.2 Disruptive Events

In order to address the issue of resilience, an understanding of potential disruptive events to which the SoI may be subject must be developed. In this context, “events” may include (but are not limited to) component failures, accidents caused by human error, malicious attacks, and natural hazards [7].

Traditionally, efforts to analyze disruptive events in a system focus on performing risk analyses by identifying hazards, quantifying risks and establishing appropriate mitigations (examples may include Failure Mode and Effects Analysis (FMEA) or Fault Tree Analysis (FTA)). A significant drawback of these approaches is that they consider only the individual components and do not address the interactions or emergent properties of complex systems. Furthermore, as system complexity increases, attempts to enumerate the probabilistic losses of the different kinds of possible hazards also tends to result in an increasingly difficult analysis [19].

Thus, an analytical reduction approach is not typically feasible for this kind of effort. The authors in [3] acknowledge this challenge, but do not offer any recommendations on how the practitioner should go about identifying these events. Rather, they simply restrict the definition of a disruptive event to those that affect the SoI such that the value of one or more FoM(s) is reduced. To this end, an extension of an existing safety-based accident model process is proposed for the systematic identification of disruptive events using a systems theory approach. Once these events are identified, the expected (i.e., average or most likely) or worst-case scenarios can be considered in the resiliency analysis [3], [15], [16], [20].

A significant difference in analyzing resilience apart from the more well-established field of reliability engineering is that reliability engineering does not consider component interactions, emergent behaviors, and perhaps most importantly, actions by calculating, intelligent adversaries [21], [22]. Adversaries do not attack systems in a statistically predictable fashion, so attempting to answer the question “What is probability that my system is attacked in fashion XYZ?” is extremely difficult and of questionable utility. A better question is “What functions must occur (or not occur) for the SoI to continue to operate correctly?” This question is best approached from a systems theory perspective.

3.2.1 STAMP and STPA

Our approach to analyzing disruptive events uses the Systems-Theoretic Accident Model and Process (STAMP), which is a systems-based accident model that focuses on enforcing behavioral safety constraints rather than preventing failures [23]. In this approach, safety is a control problem rather than a reliability problem. Constraints are exerted in a hierarchical safety control structure by higher level elements (e.g., people, organizations, engineering activities, etc.) on lower level elements [24]. Constraints are the atomic blocks in STAMP, because events leading to losses are only able to occur when safety constraints from a higher level in the safety control structure are not successfully enforced [25]. This is contrasted with traditional safety analysis methods that focus on a chain-of-events model (such as FMEA and FTA) [26] and is a primary reason to favor a top-down approach when evaluating complex systems.

Fig. 2 illustrates a typical control loop. Based on inputs provided from sensors about a given process, the controller adjusts control variables via actuators in order to maintain the desired operating condition, and a number of control loops may

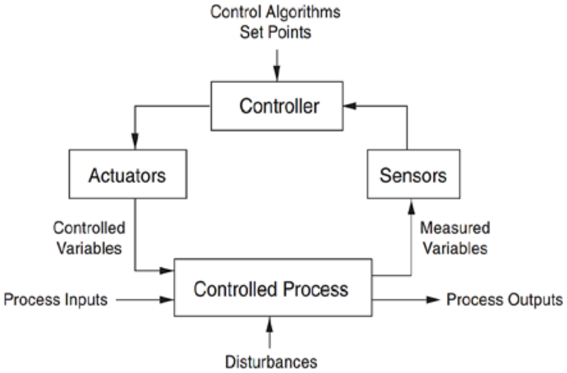

Fig. 2. Typical control loop [23].

be assembled in a hierarchical fashion to control a larger system. In Section 4, a hierarchical control structure modeled from a fictitious missile defense system is presented to further demonstrate the concept.

While it may seem unusual to consider a process designed to analyze safety failures to derive resiliency metrics, the methodical, top-down approach used by STAMP is extremely conducive. Considering the various ways how a control failure can occur produces a comprehensive, high-level list of hazards in which disruptive events might arise. This is inclusive of emergent behaviors caused by component interactions, allowing it to handle complex scenarios where analytical reduction approaches would fail [23]. Additionally, when considering the threats to resilience from a thinking adversary (i.e., system security) the prevailing perimeter defense mindset has been demonstrably ineffective at completely preventing intrusions and warrants the inclusion of a complementary, top-down, SE approach to address current shortcomings to these challenges [27].

Systems-Theoretic Process Analysis (STPA) is the hazard analysis technique built upon the foundation provided by STAMP, which follows a three-step process of: creating basic systems engineering information, identifying unsafe control actions, and identifying causal factors of unsafe control actions [25]. After gathering the basic SE information in step 1 (system purpose and goals), the second step (identify unsafe control actions) takes the controlled processes and control actions and identifies under which conditions an unsafe action could occur. This is accomplished by considering under what circumstances (a) a control action required for safety is not provided, (b) an unsafe control action is provided that leads to a hazard, (c) a potentially safe control action is provided too late, too early, or out of sequence, or (d) a safe control action is stopped too soon or applied too long.

The final step (identify causes of unsafe control actions) uses "guidewords" (such as "measurement inaccuracies" or "conflicting control actions") to assist the practitioner in enumerating the possible ways an unsafe control action might occur [25]. A useful set of security-focused guidewords was added to the existing safety-focused guidewords in [28], including "intentional congestion of feedback path" or "overriding legitimate control actions", to name a few.

It is during this step a fifth circumstance is also considered, in which an appropriate control action is provided but the controlled process fails to follow it. When implemented, the STPA process will generate a comprehensive, traceable list of hazards (i.e., disruptive events) that can be used during the development process to generate high-level safety requirements and constraints [25]. For further information and an exemplary case study on the application of STPA, see [26].

3.2.2 STPA-Sec

It can be argued that there is little distinction between safety and security other than the intentions of the actors; both are ultimately focused on preventing losses [29]. Safety is focused on preventing inadvertent losses by benevolent actors while security is focused on preventing deliberate losses by malicious actors. In other words, both are trying to provide mission assurance and ensure that the system continues to provide utility in the face of disruption. With this in mind, the STPA for Security (STPA-Sec) extension was created, and the process it follows is similar to STPA with additional considerations for the uniqueness of cyber related issues [29]. Despite its relative infancy, STPA-Sec is being actively investigated and adapted by practitioners both within the DoD and elsewhere [28], [30], [31].

The application of STPA-Sec can be structured into three phases (Concept Analysis, Architectural Analysis, and Design Analysis) which increase in both difficulty (from a technical expertise perspective) and duration [32]. As such, this analysis should be conducted by leveraging groups of engineers/subject matter experts. In the Concept Analysis phase, the system's purpose is analyzed to enumerate unacceptable losses and hazards. In the Architectural Analysis phase, the SOI's functional form is considered as the particular model elements and their responsibilities are identified, control relationships are modeled, and control actions are captured and mapped against the appropriate hazard(s). Finally, Design Analysis delves into particular scenarios where a given control action might be issued to cause a hazardous scenario. A more thorough treatment of the process can be found in [32], and the case study in Section 4 leverages this tailored STPA-Sec approach.

3.3 Identify Resilience Actions

At this point, the relevant FoM(s) have been identified and a list of disruptive events have been characterized via STPA-Sec, and potential resilience actions (in the form of component recovery mechanisms) can be identified. Each potential resilience solution or mechanism and its associated costs (in terms of time and/or resources) are used in the resilience analysis. Naturally, these are highly dependent on the design of the system under consideration (and other related factors such as its criticality), but general options may include approaches such as the repair or replacement of affected components [3], adjustments to training or procedures, or resource allocation. Ultimately, the generation of resilience options provides flexibility to determine the best approach for responding to disruptive events. Due to space constraints, we refer the reader to [3] for additional details on the topic.

3.4 Conduct Resilience Analysis

The final step involves conducting a resilience analysis which evaluates (using time and cost) how well the identified resilience actions perform while addressing the identified disruptive events. The insights generated by this process can be used to identify which strategies should be implemented in the event of a loss. Alternatively, they may identify areas of unacceptable risk to the system which may drive additional risk management decisions (e.g., transfer of risk through insurance). These results inform the SE during systems analysis in understanding how differing designs and architectures can affect system resilience, or during operations to identify critical functionality that needs to be protected from disruptions [3].

4 Case study

In this section, an illustrative example is used to demonstrate the proposed methodology for identifying, understanding, and defining resilience metrics. For this, we consider a fictitious missile defense system (henceforth called the Fictional Missile Defense System, or FMDS) originally conceived in [33]. This example was selected because it is a non-trivial representation of a realistic CPS that readily demonstrates the applicability of the proposed resilience approach to complex systems.

4.1 Step 1: Identify Figure(s)-of-Merit

For our case study, one FoM is evaluated: interceptor availability is the number of interceptor missiles whose Built-In-Test (BIT) reports indicate they are operating correctly and available for tasking. The related resilience metric is computed as a ratio of available interceptors to the nominal number in the FMDS, however if the FoM represents a quantity where a smaller number is desired (e.g., response times), the reciprocal can be considered [3]. While additional FoM's are typically considered, one is sufficient to demonstrate the approach.

4.2 Step 2: Enumerate Disruptive Events

Using the STPA-Sec Concept Analysis phase process outlined in 3.2.2, we consider the SoI's high-level purpose and generate a list of unacceptable losses and contributing hazards [32]. The FMDS is employed to destroy inbound missiles by means of detecting enemy missiles and launching interceptor missiles. Unacceptable losses are typically strategic in nature and may be specified by an organization's leadership. Leveraging one of the author's backgrounds in ballistic missile defense systems to simulate a review, two unacceptable losses are identified. The first is a successful attack causing loss of life or damage to property (i.e., mission failure). The other is loss or significant damage to any components of the FMDS. Three possible hazards that could contribute to these losses (given a worst-case scenario) are listed in Fig. 3, along with their mappings to the unacceptable losses.

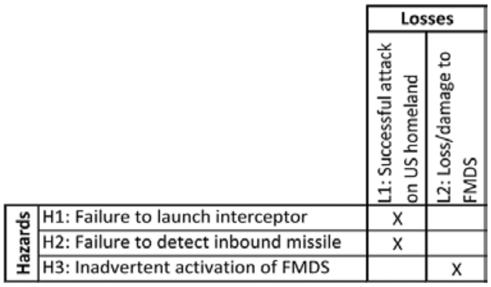

		Losses	
		L1: Successful attack on US homeland	L2: Loss/damage to FMDS
Hazards	H1: Failure to launch interceptor	X	
	H2: Failure to detect inbound missile	X	
	H3: Inadvertent activation of FMDS		X

Fig. 3. Hazards and losses mapping.

The Architectural Analysis phase considers the high-level control structure from a functional perspective and enables reasoning about critical component interactions within the system via the various control actions. Fig. 4 presents the functional control structure for the FMDS. A thorough analysis would consider each of the listed control actions, but we limit this discussion to only a single control action: ignite (issued by the flight computer to the interceptor hardware). In practice, the selected FoM(s) may be subject to interactions from additional control actions and thus a more exhaustive analysis would be required. After studying the system and building a representative model of the functional control structure and associated control actions, the practitioner creates a table that considers which (if any) hazardous conditions might be caused by the given control actions. Table 2 shows the results of the control action analysis where each of the four unsafe control action situations is considered.

The final phase of STPA-Sec, Design Analysis, defines the disruptive events of interest that are considered in the resiliency analysis. This is accomplished by studying the specifics of each control action using the STPA-Sec process model with associated guidewords as discussed earlier. By considering the various aspects of how an unsafe or insecure control action can occur, a list of potentially disruptive events is generated. While the approach presented in [3] identifies the FoM(s) before disruptive events, it is possible that creating the list of disruptive events might also highlight the necessity of one or more previously unconsidered FoM(s).

In this example, when analyzing the ignite control action there are many guidewords that could generate high-level causal scenarios of interest but for demonstration purposes we will leverage the STPA-Sec extension guideword "overriding legitimate control actions". When applied to the control action analysis element CA-Ops-1a (not providing causes hazard), it highlights the possibility that this action could create a scenario analogous to that of the Stuxnet worm on the Iranian's Natanz nuclear facility [19]. Thus, the disruptive event to be considered is a cyber-attack that degrades the operational capability of the FMDS, and more specifically, the availability of the interceptor missiles by preventing the ignite command from being issued.

Table 2: STPA-Sec Control Action Analysis.

Control Action	Not Providing Causes Hazard	Providing Causes Hazard	Providing Too Early/Too Late/Wrong Order Causes Hazard	Stopping Too Soon/Applying Too Long Causes Hazard
IGNITE	CA-Ops-1a: Not providing IGNITE command is hazardous if engagement is required to defeat missile attack [H-1]	CA-Ops-1b: Providing IGNITE command is hazardous if there is no attack underway [H-3]	CA-Ops-1c: Providing IGNITE command before ARM is hazardous if an engagement is required to defeat missile attack [H-1]	CA-Ops-1d: Stopping IGNITE too soon is hazardous if one or more inbound missiles have not yet been defeated [H-1]

4.3 Step 3: Identify Resilience Actions

For each of the disruptive events identified in the previous step, one or more resilience actions need to be identified as part of the overall resilience strategy. In this simplified example, two resilience strategies are considered for the hypothetical cyber-attack on the FMDS. The first is to replace the affected flight computer hardware at a cost of $1M per device (which includes the cost of installation and certification). This process is expected to take two months per interceptor and there are enough installation teams to install two at a time. The second strategy is to repair the affected hardware at a cost of $250K per device, however the process takes three months per interceptor (due to the additional investigation and analysis that must be given to a previously compromised system). Furthermore, there is only one team with the specialized level of knowledge to conduct the appropriate cyber assessment and validation on this system, so only one system can be repaired at a time in this manner.

4.4 Step 4: Compute Resilience Analysis

With the FoM(s), disruptive events and resilience actions identified, the final step involves conducting a resilience analysis for the given scenarios. For the purpose of this analysis, we contemplate a cyber-attack at time t_e that successfully affects 25% of the 44 Interceptors in the FMDS before defenders identify and mitigate the attack from propagating further at time t_d. Therefore, from time t_0 until the attack is halted, the FoM for the system decreases from $F(t_0)$ = 44 interceptors to $F(t_d)$ = 33 interceptors, and the associated resilience metric Я decreases from Я(t_0) = 1 to Я(t_d) = 0.75. Later at time t_s, one or more of the resilience actions identified in Step 3 are executed to begin restoring the resilience of the SoI. Here we consider three approaches; the replacement strategy (approach 1), the repair strategy (approach 2), and a hybrid strategy that applies both simultaneously (approach 1 & 2). This allows us to evaluate various resilience solutions as an optimization problem while considering the cost/time trade-offs between the approaches.

Based on the results of the analysis, several insights about the resilience of the system and methods of responding to the disruptive event studied are uncovered. The replacement strategy (approach 1) is able to restore the resilience of the system in 12 months (11 missiles, 2 months each, 2 teams) for $11M. The repair strategy (approach 2) takes much longer, requiring 33 months to fully restore the system but a much more economical cost of $2.75M (11 missiles, 3 months each, 1 team). Finally, the hybrid strategy leverages both approaches to expedite restoration, achieving its pre-disruption performance in only 9 months at a total cost of $8.75M (8 by approach 1, 3 by approach 2). By varying the utilization of the two approaches, a trade-off analysis can be performed to identify the most desirable approach for a given organization. The system's support of national defense would likely merit the most expedient restoration option; however, this example represents only a single (and arguably simplified) threat in an otherwise complex operational environment.

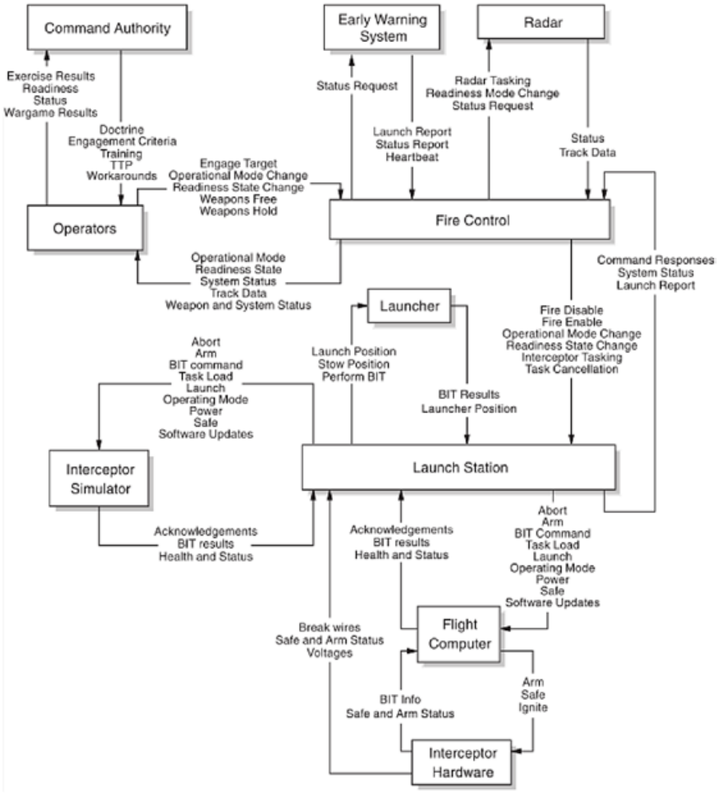

Fig. 4. Fictional Missile Defense System [33]

5 Conclusions and Future Work

This work demonstrates an approach for computing quantifiable metrics of resiliency for complex cyber-physical systems, and further it highlights the appropriateness and utility of a systems-based approach such as STPA-Sec for the critical evaluation and development of these metrics. This claim is supported through a practical case study applying the technique to a notional missile defense system. Systems-based approaches such as STPA-Sec help the SE practitioner manage the analytical complexity of system design and analysis and are a useful addition to the toolkit. Future work includes building representative stochastic models for Monte Carlo simulations to demonstrate an optimization-based approach for modeling resilience strategies.

Acknowledgment

Funding for this research was provided as part of the Air Force Office of Scientific Research grant 18RT0095 in support of Dynamic Data Driven Application Systems, PI: Dr. Eric Blasch.

Disclaimer

The views expressed in this paper are those of the authors and do not reflect the official policy or position of the U.S. Air Force, the Department of Defense, or the U.S. Government.

6 References

[1] J. R. Gosler and L. Von Thaer, "Task force report: Resilient military systems and the advanced cyber threat," *Washington, DC: Department of Defense, Defense Science Board,* p. 41, 2013.

[2] E. A. Lee, "Cyber Physical Systems: Design Challenges," in *2008 11th IEEE International Symposium on Object and Component-Oriented Real-Time Distributed Computing (ISORC)*, 2008.

[3] D. Henry and J. E. Ramirez-Marquez, "Generic metrics and quantitative approaches for system resilience as a function of time," *Reliability Engineering \& System Safety,* vol. 99, pp. 114-122, 2012.

[4] R. Francis and B. Bekera, "A metric and frameworks for resilience analysis of engineered and infrastructure systems," *Reliability Engineering \& System Safety,* vol. 121, pp. 90-103, 2014.

[5] S. Musman and S. Agbolosu-Amison, "A measurable definition of resiliency using mission risk as a metric," 2014.

[6] M. Bishop, M. Carvalho, R. Ford and L. M. Mayron, "Resilience is more than availability," in *Proceedings of the 2011 New Security Paradigms Workshop*, 2011.

[7] S. Sheard, "A Framework for System Resilience Discussions," in *INCOSE International Symposium*, 2008.

[8] A. A. Ganin, E. Massaro, A. Gutfraind, N. Steen, J. M. Keisler, A. Kott, R. Mangoubi and I. Linkov, "Operational resilience: Concepts, design and analysis," *Sci. Rep.,* vol. 6, pp. 1-12, 2016.

[9] S. Hosseini, K. Barker and J. E. Ramirez-Marquez, "A review of definitions and measures of system resilience," *Reliability Engineering \& System Safety,* vol. 145, pp. 47-61, 2016.

[10] C. S. Holling, "Resilience and Stability of Ecological Systems," *Annual review of ecology and systematics,* vol. 4, pp. 1-23, 1973.

[11] M. W. Maier, The Art of Systems Architecting, CRC press, 2009.

[12] J. M. Kendra and T. Wachtendorf, "Elements of resilience after the world trade center disaster: reconstituting New York City's Emergency Operations Centre," *Disasters,* vol. 27, pp. 37-53, 2003.

[13] J. M. Willis, R. F. Mills, L. O. Mailloux and S. R. Graham, "Considerations for secure and resilient satellite architectures," in *Cyber Conflict (CyCon US), 2017 International Conference on*, 2017.

[14] N. Yodo and P. Wang, "Engineering resilience quantification and system design implications: A literature survey," *Journal of Mechanical Design,* vol. 138, p. 111408, 2016.

[15] M. Bishop, *Resilience and Security*, 2017.

[16] D. Bodeau, R. Graubart, L. LaPadula, P. Kertzner, A. Rosenthal and J. Brennan, "Cyber resiliency metrics, version 1.0, rev. 1," *The MITRE Corp, Bedford, MA, MP120053, Rev,* vol. 1, 2012.

[17] R. Caralli, J. Allen, D. White, L. Young, N. Mehravari and P. Curtis, "CERT Resilience Management Model, version 1.2," 2016.

[18] S. A. Melnyk, D. M. Steward and M. Swink, "Metrics and performance measurement in operations management: dealing with the metrics maze," *Journal of Operations Management,* vol. 22, pp. 209-217, 2004.

[19] A. Nourian and S. Madnick, "A systems theoretic approach to the security threats in cyber physical systems applied to stuxnet," *IEEE Transactions on Dependable and Secure Computing,* 2015.

[20] A. M. Madni and S. Jackson, "Towards a conceptual framework for resilience engineering," *IEEE Systems Journal,* vol. 3, pp. 181-191, 2009.

[21] Y. Lin, D. Li, C. Liu and R. Kang, "Framework design for reliability engineering of complex systems," in *The 4th Annual IEEE International Conference on Cyber Technology in Automation, Control and Intelligent*, 2014.

[22] R. Magwood, "System Theoretic Process Analysis for Security CyberSecurity," 2017.

[23] N. Leveson, Engineering a safer world: Systems thinking applied to safety, 2011.

[24] N. Dulac, B. Owens, N. Leveson and others, "Demonstration of a new dynamic approach to risk analysis for NASA's constellation program," *Complex Systems Research Laboratory, Massachusetts Institute of Technology, Cambridge, MA,* 2007.

[25] N. Leveson, *An STPA Primer, Version 1*, 2013.

[26] S. Kawakami, "Application of a Systems-Theoretic Approach to Risk Analysis of High-speed Rail Project Management in the US," Cambridge, 2014.

[27] J. L. Bayuk and B. M. Horowitz, "An architectural systems engineering methodology for addressing cyber security," *Systems Engineering,* vol. 14, pp. 294-304, 2011.

[28] C. Schmittner, Z. Ma and P. Puschner, "Limitation and improvement of STPA-Sec for safety and security co-analysis," in *International Conference on Computer Safety, Reliability, and Security*, 2016.

[29] W. Young and N. G. Leveson, "An integrated approach to safety and security based on systems theory," *Communications of the ACM,* vol. 57, pp. 31-35, 2014.

[30] Air Force Cyber College, "Top-down Purpose-based Cybersecurity," [Online]. Available: https://www.sans.org/summit-archives/file/summit-archive-1492176717.pdf.

[31] B. T. Carter, G. Bakirtzis, C. R. Elks and C. H. Fleming, "A Systems Approach for Eliciting Mission-Centric Security Requirements," *CoRR,* vol. abs/1711.00838, 2017.

[32] M. Span, L. O. Mailloux, R. F. Mills and W. Young, "Conceptual Systems Security Requirements Analysis: Aerial Refueling Case Study," Submitted 2018.

[33] S. J. Pereira, G. Lee and J. Howard, "A system-theoretic hazard analysis methodology for a non-advocate safety assessment of the ballistic missile defense system," 2006.

[34] E. R. Griffor, C. Greer, D. A. Wollman and M. J. Burns, "SP 1500-201 Framework for Cyber-Physical Systems: Vol 1," National Institute of Standards and Technology, 2017.

[35] C. Nan and G. Sansavini, "A quantitative method for assessing resilience of interdependent infrastructures.," *Reliability Engineering & System Safety,* vol. 157, pp. 35-53, 2016.

[36] S. Jackson, "System Resilience: Capabilities, Culture and Infrastructure," *INCOSE International Symposium,* vol. 17, no. 1, pp. 885-899, 2017.