INSIGHT

This Issue's Feature:

Resilience of Complex Systems

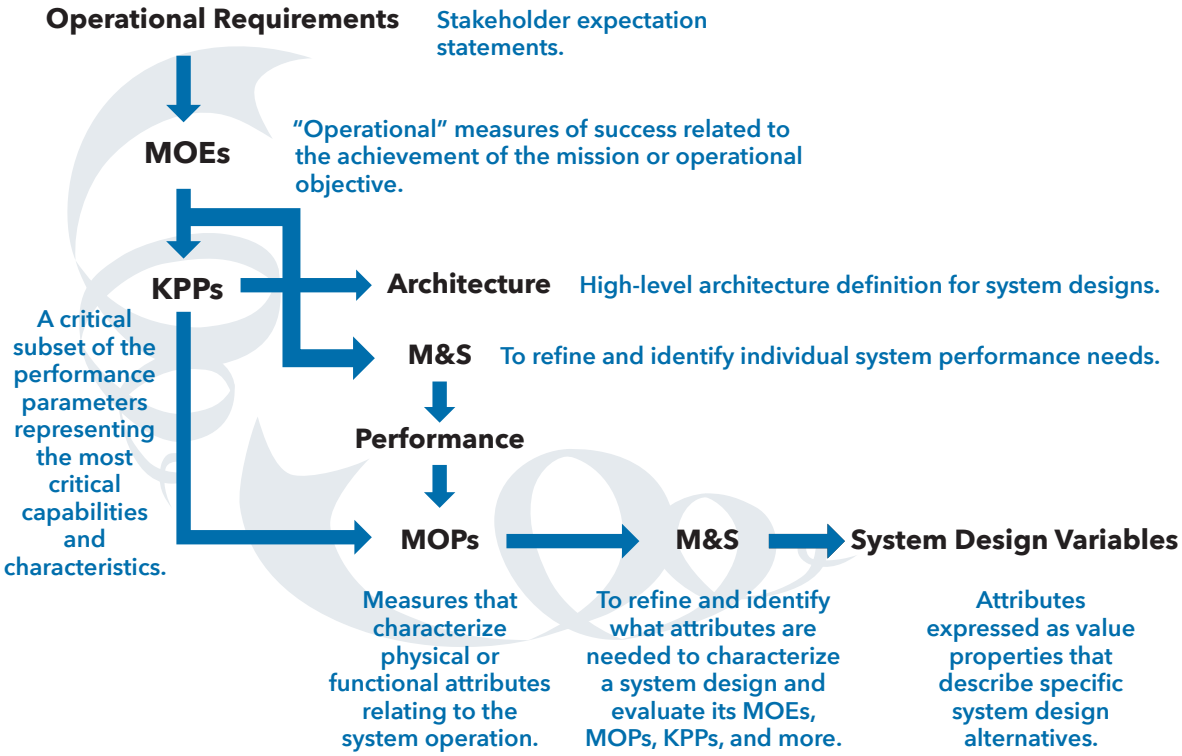

High-Level systems engineering process identifying evaluation measures and design variables

Illustration credit: from the article *Bringing Operational Perspectives into the Analysis of Engineered Resilient Systems* by Valerie B. Sitterle, Erika L. Brimhall, Dane F. Freeman, Santiago Balestrini-Robinson, Tommer R. Ender, and Simon R. Goerger (page 30)

MOEs	Measures of Effectiveness
KPPs	Key Performance Parameters
M&S	Modeling and Simulation
MOPs	Measures of Performance

APRIL 2025
VOLUME 28 / ISSUE 1

A PUBLICATION OF THE INTERNATIONAL COUNCIL ON SYSTEMS ENGINEERING

INCOSE®

WELCOME TO CATIA MBSE USER DAYS FRANCE

MAY 20TH-25, 2025

VELIZY FRANCE

This event is held in French

DASSAULT SYSTEMES

INSIGHT

A PUBLICATION OF THE INTERNATIONAL COUNCIL ON SYSTEMS ENGINEERING

APRIL 2025 VOLUME 28/ ISSUE 1

Inside this issue

APRIL 2025
VOLUME 28/ISSUE 1

About This Publication

INFORMATION ABOUT INCOSE

INCOSE's membership extends to over 25,000 members and CAB associates and more than 200 corporations, government entities, and academic institutions. Its mission is to share, promote, and advance the best of systems engineering from across the globe for the benefit of humanity and the planet. INCOSE charters chapters worldwide, includes a corporate advisory board, and is led by elected officers and directors.

For more information, click here:
The International Council on Systems Engineering
(www.incose.org)
INSIGHT is the magazine of the International Council on Systems Engineering. It is published six times per year and

OVERVIEW

features informative articles dedicated to advancing the state of practice in systems engineering and to close the gap with the state of the art. ***INSIGHT*** delivers practical information on current hot topics, implementations, and best practices, written in applications-driven style. There is an emphasis on practical applications, tutorials, guides, and case studies that result in successful outcomes. Explicitly identified opinion pieces, book reviews, and technology roadmapping complement articles to stimulate advancing the state of practice.
INSIGHT is dedicated to advancing the INCOSE objectives of impactful products and accelerating the transformation of systems engineering to a model-based discipline.
Topics to be covered include resilient systems, model-based systems engineering, commercial-driven transformational systems engineering, natural systems, agile security, systems of systems, and cyber-physical systems across disciplines and domains of interest to the constituent groups in the systems engineering community: industry, government, and academia. Advances in practice often come from lateral connections of information dissemination across disciplines and domains. ***INSIGHT*** will track advances in the state of the art with follow-up, practically written articles to more rapidly disseminate knowledge to stimulate practice throughout the community.

Editor-In-Chief William Miller
insight@incose.net +1 908-759-7110

Layout and Design Chuck Eng
chuck.eng@comcast.net

Member Services INCOSE Administrative Office
info@incose.net +1 858 541-1725

PERMISSIONS

*** PLEASE NOTE: If the links highlighted here do not take you to those web sites, please copy and paste address in your browser.**

Permission to reproduce Wiley journal Content:
Requests to reproduce material from John Wiley & Sons publications are being handled through the RightsLink® automated permissions service.

Simply follow the steps below to obtain permission via the Rightslink® system:

- Locate the article you wish to reproduce on Wiley Online Library (http://onlinelibrary.wiley.com)
- Click on the 'Request Permissions' link, under the ‹ARTICLE TOOLS› menu on the abstract page (also available from Table of Contents or Search Results)
- Follow the online instructions and select your requirements from the drop down options and click on 'quick price' to get a quote
- Create a RightsLink® account to complete your transaction (and pay, where applicable)
- Read and accept our Terms and Conditions and download your license
- For any technical queries please contact customercare@copyright.com
- For further information and to view a Rightslink® demo please visit www.wiley.com and select Rights and Permissions.

AUTHORS – If you wish to reuse your own article (or an amended version of it) in a new publication of which you are the author, editor or co-editor, prior permission is not required (with the usual acknowledgements). However, a formal grant of license can be downloaded free of charge from RightsLink if required.

Photocopying
Teaching institutions with a current paid subscription to the journal may make multiple copies for teaching purposes without charge, provided such copies are not resold or copied. In all other cases, permission should be obtained from a reproduction rights organisation (see below) or directly from RightsLink®.

Copyright Licensing Agency (CLA)
Institutions based in the UK with a valid photocopying and/or digital license with the Copyright Licensing Agency may copy excerpts from Wiley books and journals under the terms of their license. For further information go to CLA.

Copyright Clearance Center (CCC)
Institutions based in the US with a valid photocopying and/or digital license with the Copyright Clearance Center may copy excerpts from Wiley books and journals under the terms of their license, please go to CCC.

Other Territories: Please contact your local reproduction rights organisation. For further information please visit www.wiley.com and select Rights and Permissions.
If you have any questions about the permitted uses of a specific article, please contact us.

Permissions Department – UK
John Wiley & Sons Ltd.
The Atrium,
Southern Gate,
Chichester
West Sussex, PO19 8SQ
UK
Email: Permissions@wiley.com
Fax: 44 (0) 1243 770620
or

Permissions Department – US
John Wiley & Sons Inc.
111 River Street MS 4-02
Hoboken, NJ 07030-5774
USA
Email: Permissions@wiley.com
Fax: (201) 748-6008

ARTICLE SUBMISSION insight@incose.net

Publication Schedule. ***INSIGHT*** is published six times per year. Issue and article submission deadlines are as follows:

- April 2025 issue – 2 January 2025
- May 2025 issue – 1 February 2025
- June 2025 issue – 1 March 2025
- August 2025 issue – 1 May 2025
- October 2025 – 1 July 2025
- December 2025 – 1 September 2025
- February 2026 issue – 1 November 2026

For further information on submissions and issue themes, visit the INCOSE website: www.incose.org

 ISSN 2156-485X; (print)
ISSN 2156-4868 (online)

ADVERTISE

Readership
INSIGHT reaches over 25,000 members and CAB associates and uncounted employees and students of more than 130 CAB organizations worldwide. Readership includes engineers, manufacturers/purchasers, scientists, research and development professionals, presidents and chief executive officers, students, and other professionals in systems engineering.

Issuance	Circulation
2025, Vol 28, 6 Issues	100% Paid

Contact us for Advertising and Corporate Sales Services
We have a complete range of advertising and publishing solutions professionally managed within our global team. From traditional print-based solutions to cutting-edge online technology the Wiley-Blackwell corporate sales service is your connection to minds that matter. For an overview of all our services please browse our site which is located under the Resources section. Contact our corporate sales team today to discuss the range of services available:

- Print advertising for non-US journals
- Email Table of Contents Sponsorship
- Reprints
- Supplement and sponsorship opportunities
- Books
- Custom Projects
- Online advertising

Click on the option below to email your enquiry to your nearest office:

- Asia and Australia corporatesalesaustralia@wiley.com
- Europe, Middle East and Africa (EMEA) corporatesaleseurope@wiley.com
- Japan corporatesalesjapan@wiley.com
- Korea corporatesaleskorea@wiley.com

USA (also Canada, and South/Central America):

- Healthcare Advertising corporatesalesusa@wiley.com
- Science Advertising Ads_sciences@wiley.com
- Reprints Commercialreprints@wiley.com
- Supplements, Sponsorship, Books and Custom Projects busdev@wiley.com

Or please contact: Marcom@incose.net

CONTACT

Questions or comments concerning:

Submissions, Editorial Policy, or Publication Management
Please contact: William Miller, Editor-in-Chief
insight@incose.net

Advertising—*please contact:*
Marcom@incose.net

Member Services – *please contact:* info@incose.org

ADVERTISER INDEX

April Volume 28-1

CORPORATE ADVISORY BOARD — MEMBER COMPANIES

Advanced Systems Engineering, LLC
Aerospace Corporation, The
Airbus
AM General LLC
Analog Devices, Inc.
Arcfield
Australian National University
AVIAGE SYSTEMS
Aviation Industry Corporation of China, LTD
BAE Systems
Bechtel
Becton Dickinson
Belcan Engineering Group LLC
BMT Canada
Boeing Company, The
Booz Allen Hamilton Inc.
Boston Scientific Corporation
BTS Software Solutions
California State University Dominguez Hills
Carnegie Mellon Univ. Software Engineering Institute
Change Vision, Inc.
Colorado State University Systems Engineering Programs
Cornell University
Cranfield University
C.S. Draper Laboratory, Inc.
Cubic Corporation
Cummins, Inc.
Dassault Systèmes
Defense Acquisition University
Deloitte Consulting, LLC
Denso Create Inc
DENTSU SOKEN INC
Drexel University
Eaton
Eindhoven University of Technology
EMBRAER
FAMU-FSU College of Engineering
Federal Aviation Administration (U.S.)
Ford Motor Company
GE Aerospace
General Dynamics
General Motors
George Mason University
Georgia Institute of Technology
Hitachi Energy
Honeywell Aerospace Technologies
Idaho National Laboratory
ISAE - Supaero
ISDEFE
IVECO Group
Jama Software
Jet Propulsion Laboratory
John Deere & Company
Johns Hopkins University
KBR, Inc.
KEIO University
L3Harris Technologies
Lawrence Livermore National Laboratory
Leidos
LEONARDO
Lockheed Martin Corporation
Los Alamos National Laboratory
Loyola Marymount University
Magna
ManTech International Corporation
Marquette University
Massachusetts Institute of Technology
MBDA (UK) Ltd
Medtronic
MetaTech Consulting Inc.
Missouri University of Science & Technology
MITRE Corporation, The
Mitsubishi Electric Corporation
Mitsubishi Heavy Industries, Ltd
Modern Technology Solutions Inc
National Aeronautics and Space Administration (NASA)
National Reconnaissance Office (NRO)
National Security Agency Enterprise Systems
Naval Postgraduate School
Nissan Motor Co, Ltd
Northrop Grumman Corporation
Pacific Northwest National Laboratory
Pennsylvania State University
Petronas International Corporation Limited
Prime Solutions Group, Inc
Project Performance International (PPI)
Purdue University
RealmOne
Rolls-Royce
RTX
Saab AB
SAFRAN
SAIC
Sandia National Laboratories
Saudi Railway Company
SENSEONICS
Shanghai Formal–Tech Information Technology Co., Ltd
Shell
Siemens
Sierra Nevada Corporation
Singapore Institute of Technology
Southern Methodist University
SPEC Innovations
Stevens Institute of Technology
Strategic Technical Services LLC
Swedish Defence Materiel Administration (FMV)
Systems Planning and Analysis
Taiwan Space Agency
Tata Consultancy Services
Thales
The George Washington University
The University of Arizona
The University of Utah
Torch Technologies
TOSHIBA Corporation
Trane Technologies
Tsinghua University
UK MoD
UNCOMN
Universidade Federal De Minas Gerais
University of Alabama in Huntsville
University of Arkansas
University of California San Diego
University of Connecticut
University of Maryland
University of Maryland, Baltimore County
University of Maryland Global Campus
University of Michigan, Ann Arbor
University of New South Wales, The, Canberra
University of South Alabama
University of South-Eastern Norway (USN)
University of Southern California
University of Texas at El Paso (UTEP)
US Department of Defense
Veoneer US Safety Systems, LLC
Virginia Tech
Volvo Cars Corporation
Volvo Construction Equipment
Wabtec Corporation
Wayne State University
Weber State University
Wichita State University College of Engineering
Woodward Inc
Worcester Polytechnic Institute (WPI)
Woven by Toyota, Inc.
Zuken, Inc

FROM THE EDITOR-IN-CHIEF

William Miller, insight@incose.net

We are pleased to publish the April 2025 *INSIGHT* published cooperatively with John Wiley & Sons as the systems engineering practitioners' magazine. The *INSIGHT* mission is to provide informative articles on advancing the practice of systems engineering as the state-of-the-art advances as evidenced in *Systems Engineering*, the Journal of INCOSE also published by Wiley, as well as papers presented at symposia and conferences by INCOSE and in the broader systems community.

The focus of this April issue of *INSIGHT* is themed on resilience of complex systems. The imperative to address 'resilience' is a priority of the future of systems engineering (FuSE) to realize the *System Engineering Vision 2035*. FuSE is charged by the INCOSE Strategic Plan v1.0 (17 June 2024) Objective O.1 Advance systems engineering as the world's trusted authority and Key Result KR1.1 Satisfaction of/progress against future of systems engineering roadmap. INCOSE Technical Operations has recently chartered the loss-driven systems engineering (LDSE) project to achieve unification of the loss-driven quality characteristics such as security, safety, resilience, operational risk, environmental protection, availability, etc. Resilience is not a "blank space" in systems engineering as there is substantial body of work to build on.

We lead the April 2025 *INSIGHT* with "A Generic State-Machine Model of System Resilience" by Scott Jackson, Stephen Cook, and Timothy Ferris from the April 2015 *INSIGHT* themed on resilience with Scott as the theme editor. This article has to do with the states a system must pass through from an operational state to a restoration of full functionality, to partial functionality, or to a final decommissioning. The article shows the system must pass between 7 defined states in 28 defined transitions.

"An SoS Analytical Workbench Approach to Architectural Analysis and Evolution" by Daniel DeLaurentis, Navindran Davendralingam, Karen Marais, Cesare Guariniello, Zhemei Fang, and Payuna Uda from the October 2016 *INSIGHT* themed on systems of systems with Stephen Cook the theme editor was awarded best paper in 2016. The article summarizes the development of a System of Systems Analytic Workbench (SoS AWB) that provides a set of computational tools to facilitate better-informed decision-making on evolving SoS architectures. The workbench motif is adopted since SoS practitioners typically generate archetypal technical queries that can be mapped to appropriate analysis methods best suited to provide outputs and insights directly relevant to posed questions. After an overview of the workbench framework, four distinct methods currently available for use are presented along with their distinctive aspects in the concept of use.

"How Infrastructure Can Become Reborn by Becoming Born Robust" by Josh Sparber from the December 2016 *INSIGHT* themed on critical infrastructure protection and recovery (CIPR) with theme editors Loren (Mark) Walker, Mike DeLamare, and John Juhasz illustrates how using risk-based profiles can build bottom up constructions of structures resistant to electromagnetic pulses (EMPs) within a micro-grid. With an assemblage of components tested for risk, simulated as SysML EMP threat use cases, an engineer can test and redesign selected portions of the power grid against EMP vulnerability and find paths enhancing survivability or improving the reliability of the enclosing system.

"Effective and Efficient Preparation for the Unforeseeable" by Steve Hinsley, Michael Henshaw, and Carys Siemieniuch from the June 2017 *INSIGHT* themed on systems that are fit for purpose with guest editor the late Jack Ring states that a system of systems (SoS) is fit for purpose when it implements the correct, timely, and complete transfers of material, energy, and/or information (MEI) between its constituents and with its external environment that are necessary to achieve a particular result. The article then addresses the challenge in maintaining the SoS fit for purpose after unpredictable changes in operation, composition, or external factors.

"Bringing Operational Perspectives into the Analysis of Engineered Resilient Systems" by Valerie Sitterle, Erika Brimhall, Dane Freeman, Santiago Balestrini-Robinson, Tommer Ender, and Simon Goerger from the September 2017 *INSIGHT* themed on the Systems Engineering Research Center (SERC) with the support of Executive Director Dinesh Verma, former Chief Technology Officer (CTO) Jon Wade, and the late Barry Boehm, chief scientist, focuses on the evaluation of early-stage design alternatives regarding their modeled operational performance and characteristics. The work in this article ties together differentiated operational needs with requirements specification and maturation of previous analytical constructs toward a more operationally relevant viewpoint.

"Extending Formal Modeling for Resilient Systems Design" by Azad Madni, Michael Sievers, Ayesha Madni, Edwin

Ordoukhanian, and Parisa Pouya from the October 2018 *INSIGHT* themed on Industry 4.0 and opportunities for systems engineering presents a flexible contract-based approach that employs a combination of formal methods for verification and testing and flexible assertions and probabilistic modeling to handle uncertainty during mission execution. A flexible contract (FC) is a hybrid-modeling construct that facilitates system verification and testing while offering the requisite flexibility to cope with non-determinism.

"Systems Theory Principles and Complex Systems Engineering Concepts for Protection and Resilience in Critical Infrastructure: Lessons from the Nuclear Sector" by Adam Williams from the June 2020 *INSIGHT* themed on critical infrastructure protection and recovery (CIPR) with theme editor Mitchell Kerman presents US Sandia National Labs research exploring the safety, safeguards, and security risks and their mitigation for three different nuclear sector-related activities — spent nuclear fuel transportation, small modular reactors, and portable nuclear power reactors. The research shows that a systems-theoretic approach can better identify interdependencies, conflicts, gaps, and leverage points across traditional safety, security, and safeguards hazard mitigation strategies in the nuclear reactors, materials, and waste sector.

"Engineering a Cyber Resilient Product Line" by Patrice Williams, Paula Moss, Susan Bataller, and Suzanne Hassell from the September 2020 *INSIGHT* themed on cyber secure and resilient approaches with feature-based product line engineering with theme editors Beth Wilson and Bobbi Young describes how to apply cyber resiliency analysis to product line architectures, introducing the "cyber resiliency wheel" technique.

"Harmonizing the Domains of Loss-Driven Systems Engineering" by Keith Willett from the December 2020 *INSIGHT* themed on loss-driven systems engineering (LDSE) with theme editor John Brtis establishes interrelationships among the LDSE domains (reliability, sustainability, survivability, risk management, resistance, resilience, agility, safety, and security) to harmonize role, fit, function, and impact among the domains focusing on sustaining value-delivery. Traditional systems engineering treats these as separate domains with varying degrees of detail, rigor, and results. LDSE proposes consolidating these domains for a comprehensive, cohesive, and consistent approach to address system loss. System characteristics include what it is (structure, state), what it does (function, behavior), where it resides (environment, containing whole), what it uses (resources, energy source, raw material), what it contains (content), and why it exists (value delivery). An adversity produces a disturbance that can induce stress in a system so it may suffer some loss within one or more of these characteristics.

"Versatile Test Reactor Open Digital Engineering Ecosystem" by Christopher Ritter, Jeren Browning, Peter Suyderhoud, Ross Hays, AnnMarie Marshall, Kevin Han, Taylor Ashbocker, John Darrington, and Lee Nelson awarded a best paper in 2022 from the March 2022 *INSIGHT* themed on digital engineering with theme editors Frank Salvatore and Tracee Gilbert hypothesize using digital engineering principals to reduce risk and cost and gain schedule efficiencies in the design of a 300-MWt sodium-cooled fast reactor. This ecosystem was deployed to over 200 engineers and used to deliver the conceptual design of the virtual test reactor (VTR).Initial results show significant reductions in user latency (1000x at peak use), the possibility of direct finite-element-analysis (FEA) integrations to computer-aided design (CAD) tools, and nuclear reactor system design descriptions (SDDs) that one can fully link throughout design in data-driven requirements-management software. Early results have led to the VTR program maintaining milestone performance during the COVID-19 pandemic.

"Systematic Identification and Analysis of Hazards for Automated Systems" by Lina Putze and Eckard Böde from the December 2022 *INSIGHT* themed on the Archimedes Initiative, a global systems engineering research network, with theme editors Wouter Leibbrandt and Dinesh Verma addresses the problems of finding common sources of criticality for specific application classes and identifying and quantitatively assessing new sources of harm within particular automated driving systems. The introduction of automation into technical systems promises many benefits, including performance increase, improved resource economy, and fewer harmful accidents. In the automotive sector, automated driving is seen as one key element in Vision Zero by eliminating common accident causes such as driving under the influence, reckless behavior, or distracted drivers. However, this is contrasted by new failure modes and hazards from the latest technologies.

"Enhancing Early Systems R&D Capabilities with Systems — Theoretic Process Analysis" by Adam Williams from the September 2023 *INSIGHT* themed on systems engineering in early-stage research and development with theme editors Michaael DiMario and Ann Hodges demonstrates the benefit of systems–theoretic process analysis (STPA) for early system R&D strategy and development. The article describes diverse use cases for cyber security, nuclear fuel transportation, and US electric grid performance. The traceability, rigor, and comprehensiveness of STPA serves to improve R&D strategy and development. Leveraging STPA as well as related systems engineering techniques can be helpful in early R&D planning and strategy development to better triangulate deeper theoretical meaning or evaluate empirical results to better inform systems engineering solutions.

We hope you find *INSIGHT*, the practitioners' magazine for systems engineers, informative and relevant. Feedback from readers is critical to *INSIGHT*'s quality. We encourage letters to the editor at insight@incose.net. Please include "letter to the editor" in the subject line. *INSIGHT* also continues to solicit special features, standalone articles, book reviews, and op-eds. Please contact us at FuSE@incose.net if you are interested in contributing to our body of knowledge accounting for uncertainty in the engineering of systems. For information about *INSIGHT*, including upcoming issues, see https://www.incose.org/publications/insight. For information about sponsoring *INSIGHT*, please contact the INCOSE marketing and communications director at marcom@incose.net. ■

APRIL 2025
VOLUME 28/ISSUE 1

INSIGHT Special Feature

A Generic State-Machine Model of System Resilience

Scott Jackson, jackson@burnhamsystems.net, **Stephen Cook,** stephen.cook@incose.org, and **Timothy L. J. Ferris,** timothy.ferris@incose.org

INTRODUCTION

System resilience means different things to different people and different things across different industries and system contexts. For example, in some contexts, the need is for the system performance to be unaffected after the occurrence of a defined threat event, whereas in others, system impairment or even loss may be acceptable, particularly for severe threats, providing certain constraints, such as system safety, remain uncompromised. When wishing to specify resilience or to describe or analyse the performance and behaviour of a system in response to a threat situation, it is useful to have a conceptual model that can support these activities.

In this article, we present a state-based conceptual model of the variety of states that a system may experience when encountering and resolving a resilience-related situation. We contend that it has promise for framing discussion on resilience objectives of a particular system during the design process by imbuing a common understanding of the expected resilience characteristics of the system to all stakeholders. Furthermore, we advocate that during operations the model informs decisions on how best to deal with multiple resilience-related issues, such as an impending threat and an impaired system.

A resilient system could pass through a number of states when it encounters a threat event beyond its design limits. The system would start in the nominal operating state and transition to partially functional states and sometimes to completely non-functional states depending on the event and the resilience of the system. If the system can undergo repair, it can also transition back to the nominal state. The transitions between these states are dependent on the type and magnitude of the threat event, the resilience of the system as determined by its design, and the decisions and actions undertaken by its operators, and the conditions including thresholds that trigger transitions.

The development of the state-machine model arose from the desire among the authors to capture the essence or definition of what comprises system resilience in a usable way for practice. The model evolved to include the potential of the system to sense and anticipate future threat events whether they be within design limits or otherwise as well as states of impaired functionality and agreed outcome states. Certain engineers may well have different conceptions of what should be included in resilience; adjusting the states and state transitions to suit their preferences easily accommodates this.

The adapted state-machine or automata model chosen comes from Aleksander and Hanna (1976) and is attributable to Mealy (1955) because of its simplicity and suitability. A Mealy model is sequential automata that transitions between states and produces outputs based on its current state, the current input, and the state transitions that are available for the current state. In this initial formulation, the model is asynchronous and remains in its current state until it receives a pre-defined state-transition trigger. Upon such an event, the model will undergo the defined state transition and system capability sets to the appropriate pre-defined value. Although many alternative automata representations are possible and varying numbers of states could meet the need in hand, what we describe below is what we consider a reasonably comprehensive model with general utility. Information provided later in the article specifies how to customise the model and use it for a specific system.

The generic state-machine model of resilience shown in Figure 1 derives from the examination of a set of resilience case studies and resilience principles (Jackson and Ferris 2013). It encompasses a sufficient number of states and state transitions for a wide variety of purposes and the descriptions of the states and the state transitions follow.

RESILIENCE STATES

State A – Nominal Operational State

In this state, the system can operate normally or operate in some expected rest condition (such as those experienced by emergency services when not deployed). One would expect the system to operate for long periods in this state and the inputs of note would be either the knowledge of an approaching threat or the arrival of the threat event itself. A threat could be an external threat such as another train on a collision course as it was in the Metrolink case (NTSB 2010) or an internal threat. The latter could include design and construction flaws, for example, the use of undersized reinforcing rods in the Minneapolis Bridge case (Wald and Chang 2007) or age-related flaws such as the aging fuel line seals in the Nimrod case (Haddon-Cave 2009).

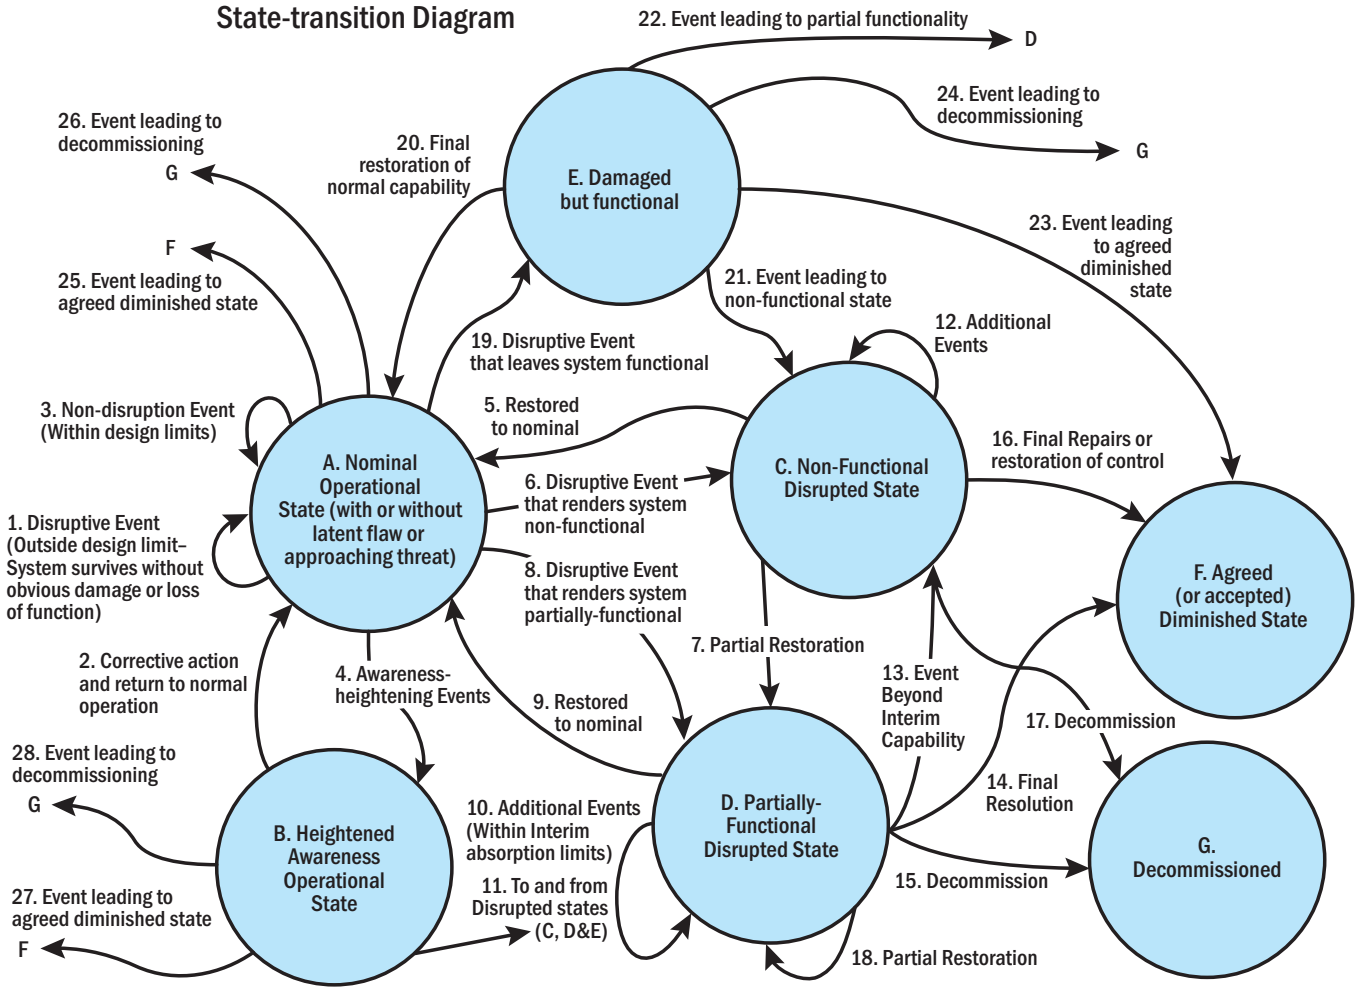

Figure 1. A Generic State-Machine Model of Resilience

State B – Heightened Awareness Operational State

The system enters this state in response to an awareness-heightening event warning of impending external or internal threats. An example of the first type of threat is an approaching train detected using positive train control. In the case of latent flaws, this state could represent a latent structural flaw detected by sensing or by other means. The system is fully functional in this state.

State C – Non-Functional Disrupted State

In this state, the system is completely non-functional following a disruptive event. This is probably the most common state associated with major disasters.

State D – Partially-Functional Disrupted State

In this state, the system is subject to a disruptive event that reduces its functionality. This is a commonly encountered state for complex systems designed to be resilient and perform a large number of functions in the face of internal and external threats. This state can also be used to describe systems that are yet to reach their nominal capability, the initial operating capability of the Hubble space craft (NASA 2009) is an example. The system suffered from reduced functionality because of a latent flaw even though it remained undamaged. An electricity distribution network that encounters an overload event that causes certain circuit breakers to trip is another example that one can model using this state.

State E – Damaged but Functional State

In this state, the system has suffered some physical damage but has been able to retain its functionality (at least in the near term). This may happen when there is sufficient redundancy in the system to retain functionality even with damage to one of its redundant branches. In this state, the system may have a diminished ability to handle any further threats but it would be able to function as long as there is no more degradation.

State F – Agreed Diminished State

In this state, there is recognition that the system has diminished functionality from which it cannot recover to normal except through a complete rebuilding of the system. This can happen say in a power station when one or more generators or nuclear reactors is damaged and the damaged plant is permanently decommissioned but the remaining parts of the plant continue operating offering reduced capability.

State G – Decommissioned State

To reach this state, the system suffered a reduction in functionality from a threat event and someone makes the decision to decommission and scrap it. Once in this state, there is no pathway to return the system to a functional state.

EXAMPLE TRANSITIONS BETWEEN STATES

Transition 1 – Disruptive Event outside Design Limit (State A to State A)

This transition represents an event that has not substantially damaged the system

nor affected its functionality. This can occur when the design margin is considerably greater than the nominal design limit and the threat level is below the margin level. The reason this transition was included is to recognise that events of this type while not affecting functionality could well cause subtle damage that reduces operational life of system components and potentially lowers the margin level. Such information could be captured in future conditional and stochastic model expansions.

Transition 2 – Return to Normal Alertness (State B to State A)

This transition occurs upon mitigation of the detected threat or if the threat dissipated. The events that could cause this transition could include the end of a threat time window or the rectification of latent system flaws.

Transition 3 – Non-Disruptive Event Inside Design Limit (State A to State A)

This transition is similar to Transition 1 in that the system damage is not overt and the functionality remains unaffected since the threat is inside the design limit. This transition has been included in the model in recognition that it is common for systems to push beyond their nominal operating conditions and although this does not cause immediate impact, it may reduce service life and mandate additional maintenance.

Transition 4 –Awareness Heightening Event (State A to State B)

This transition occurs by one of two types of events. Firstly, it can result from the detection of an approaching external threat, for example, the detection of an approaching train by positive train control as recommended after the Metrolink case (NTSB 2010). It can also arise from the detection of a latent flaw, for example, when an aircraft structural integrity system detects microscopic cracks in the structure.

Transition 5 – Restored to Nominal (State C to State A)

The USS Cole (DoD 2001) is a useful example of this transition. The ship became completely non-functional and then later restored to a completely nominal condition. The expectation is that the transition from non-functional to fully functional occurs via a comprehensive repair and restoration operation.

Transition 6 – Disruptive Event that Renders the System Non-Functional (State A to State C)

This transition occurs when a disruptive event occurs without the system entering the Heightened Awareness Operational State. This can happen in response to an unexpected external event such as a tsunami or an internal event such as the failure without warning of the Minneapolis Bridge (Wald and Chang 2007).

Transition 7 – Partial Restoration (State C to State D)

It is common in large, complex systems to achieve partial restoration of functionality from the Non-Functional Disrupted State as an intermediate condition before either returning the system to the Nominal State, the Agreed Diminished State or the Decommissioned State.

Transition 8 –Disruptive Event that Renders the System Partially Functional (State A to State D)

Like Transition 6, this transition occurs when a disruptive event occurs without the system entering the Heightened Awareness Operational State but in this case, the system retains some degree of functionality. This transition represents a common occurrence for complex systems designed to retain partial functionality after disruptive events.

Transition 9 – Restored to Nominal (State D to State A)

This transition occurs when a partially functional system restores to full functionality.

Transition 10 – Additional Events (within interim absorption capability) (State D to State D)

This transition occurs when the system encounters additional potentially disruptive events after which the system retains a degree of functionality. The functionality decreases or conceivably increases, but the degree of functionality is not explicit in this model but easily occurs with an expansion of the number of partially functional disrupted states.

Transition 11 – The set of transitions between the Heightened Awareness State and Damaged States (State B to State C, D or E)

These are a set of transitions that mirror transitions from State A to States C, D, or E. These transitions represent the detection of a threat to the situation, where disruption to the system already exists. This transition is common in military scenarios, for example, when there is detection of incoming missiles and defences ready, but some missiles penetrate the defences nonetheless.

Transition 12 – Additional Events (State C to State C)

This transition occurs when the system encounters additional disruptive events that leave the system non-functional with or without additional damage. This transition could also describe unsuccessful attempts to restore the system.

Transition 13 – Event Beyond Interim Capability (State D to State C)

This transition occurs when partially functional systems become non-functional by a subsequent event beyond the interim resilience of the impaired system.

Transition 14 – Final Resolution (State D to State F)

This transition occurs when the system moves from the partially disrupted state to one of acceptance as the final state of the system following the threat event (or threat sequence) albeit not the fully functional state. The transition occurs upon restoration of additional capability to meet an agreed-to level of functionality, or the sealing of an agreement to accept the current impaired level of functionality as the final state.

Transition 15 – Decommission (State D to State G)

This transition occurs when the system undergoes decommissioning from the partially functional state.

Transition 16 – Agreed Repairs or Restoration of Control (State C to State F)

This transition occurs when a non-functional system restores to an agreed level of functionality or control of the system returns to the operators. An example of this transition occurred in the Apollo 11 case (National Geographic News 2010) when the pilot regained control of the space craft after it was rendered partially functional by software errors.

Transition 17 – Decommission (State C to State F)

This transition occurs when a non-functional system becomes permanently decommissioned.

Transition 18 – Partial Restoration (State D to State D)

This transition occurs when the system functionality of a partially functional system is improved, but to neither its fully functional state, nor its final diminished state. If there were additional partially functional states then there would be a chain of such transitions between such states.

Transition 19 – Disruptive Event that Leaves System Functional (State A to State E)

This transition occurs when the

Table 1. Example subset of a state transition table N for an electricity generation and distribution system

	e_1 Single Generator Failure	e_2 Start Back-up Generator 1	e_3 Start Back-up Generator 2	e_4 Repair Generator(s)	e_5 Reduce Demand
q_1 Nominal	q_2 75% capacity	q_1 100% capacity	q_1 100% capacity	q_1 100% capacity	q_1 100% capacity
q_2 Partially-Functional Disrupted	q_2 additional 25% capacity reduction per generator lost	q_2 12.5% additional capacity	q_1 100% capacity for single generator failure	q_1 100% capacity	q_3 unchanged capacity level ft
q_3 Damaged but Functional	q_2 additional 25% capacity reduction per generator lost	q_3 12.5% additional capacity	q_1 100% capacity for single generator failure	q_1 100% capacity	q_3 unchanged capacity level ft
q_4 Non-Functional Disrupted State	q_4 0% capacity	q_2 12.5% additional capacity	q_2 12.5% additional capacity	q_1 100% capacity	q_2 unchanged capacity level ft
q_5 Agreed Diminished State	q_2 additional 25% capacity reduction per generator lost	q_5 12.5% additional capacity	q_5 12.5% additional capacity	q_5 100% of original capacity	q_5 unchanged capacity level ft

system suffers damage by an event but its functionality is unimpaired. This can occur when there is damage to redundant equipment or when the system sustains damage that does not affect its functionality. A system in this state may well be in need of urgent repairs or have a reduced operational life and is thus is not in its Nominal Operational State.

Transition 20 – Final Restoration (State E to State A)

This transition occurs when repairs on the system are complete and the system returns to its nominal state.

Transition 21 – Event Leading to Non-Functional State (State E to State C)

This transition occurs when a damaged but functional system encounters another event that leaves it non-functional.

Transition 22 – Event Leading to Partial Functionality (State E to State D)

This transition occurs when a damaged but functional system encounters another event that leaves it partially functional. For example, a system may have redundant branches, one of which suffered damage in a previous event but the other branch lets the system function normally as long as that branch is undamaged. A subsequent threat event that damages the remaining branch will result in this transition or Transitions 23 or 24.

Transition 23 – Event Leading to Agreed Diminished State (State E to State F)

This transition occurs when a damaged but functional system migrates to one in the Agreed Diminished State. This could happen when a decision to not to repair a redundant branch of the system or a damaged but functional major component is decommissioned to deal with damage with concomitant reduction in system capability.

Transition 24 – Event Leading to Decommissioning (State E to State G)

This transition occurs when a damaged but functional system encounters another event that leads to decommissioning such as an event that damages the second branch of the system leaving the system entirely non-functional and irreparable.

Transition 25 – Event Leading to Agreed Diminished State (State A to State F)

This transition occurs when an event occurs which leaves a previously normal operating system in an agreed diminished state.

Transition 26 – Event Leading to Decommissioning (State A to State G)

This transition occurs when an event occurs which leaves a previously normal operating system in a decommissioned state.

Transition 27 – Event Leading to Agreed Diminished State (State B to State F)

This transition occurs when an event happens which leaves a previously normal operating in heightened awareness system in an agreed diminished state.

Transition 28 – Event Leading to Decommissioning (State B to State G)

This transition occurs when an event occurs which leaves a previously normal operating system in a decommissioned state.

USING THE STATE MODEL TO SPECIFY RESILIENCE

In order to use a model of this type to specify resilience, it is necessary to define the elements the 5-tuple <*Q*, *E*, *F*, *N*, *W*>, where:

$Q = \{q_1, q_2, \ldots, q|Q|\}$ is the set of discrete internal states selected to define the possible states of the system.

$E = \{e_1, e_2, \ldots, e|E|\}$ is the set of events that can cause state transitions in the system;

$F = \{f_1, f_2, \ldots, f|F|\}$ is the set of functional capability levels possibly exhibited by the system;

N is the function that relates every pair of elements from *E* and *Q* to the next state q_{t+1}, i.e., the state transitions;

W is the function that relates every pair of elements to an element in *F*, i.e., f_{t+1}; and $|Q|$, $|E|$, $|F|$ are the number of elements in *Q*, *E*, and *F*, respectively.

The first step in specifying a system using this class of model would be to enumerate the events that could cause state transitions *E*. The next step would be to decide the number of system functional capability levels *F* and then determine the number of system states *Q*. As the model is a discrete automaton, the functions *N* and *W* fit in a table in textual form. An illustrative subset

Table 2. The US Airways Flight 1549 Case

Transition Event (e_t)	Current State (q_t)	New State (q_{t+1})
8. Ingest flock of geese	*A*, Nominal	*C*, Partially-Functional Disrupted
10. Stabilize situation – work out what still functions and how to land	*C*, Partially-Functional Disrupted	*F*, Diminished State
7. Control aircraft and land (ditch)	*F*, Diminished State	*C*, Non-functional Disrupted
17. Salvage	*C*, Non-functional Disrupted	*G*, Decommissioned

of an example state-transition table for an electricity generation and distribution system is in Table 1.

In this example, the Nominal Operational State comprises four operational generators of equal capacity. Thus if one fails, 25% of the capacity is lost. The states on the left hand side represent the state of the system at the time of the event and the columns represent the incoming event. Each transition cell in the table indicates the next state q_{t+1} and then the system capacity f_{t+1}. As would be expected, many events do not cause a state transition or a change in capacity because there is usually only a small subset of the transition applicable to any one state. In completing this example table, it became apparent that it would be useful to add additional richness to the model. For example, given the small number of generators it would be appropriate to introduce additional states or another state variable to reflect the number of generators fully functional. Further, it would be useful to introduce additional state variables, specific to the system context, to reflect current power demand and percentage outage. Such additions are simple to accommodate in a model of this type.

USING THE STATE MODEL TO UNDERTAKE CASE STUDY ANALYSIS

US Airways flight 1549 (Pariès 2011) collided with a flock of geese which totally rendered the engines powerless. The aircraft was able to remain functional due to redundant power and hydraulic systems enabling the pilot to ditch the aircraft in the Hudson River in a controlled way. All passengers survived, but US Airways decommissioned the aircraft. Table 2 shows the analysis of this case using the generic state-machine model described above.

CONCLUSIONS

A generic resilience model has been described that can form a useful basis for practitioners; one that can be extended using the formalism provided. Above all, this class of model provides a roadmap that assists in identifying possible threat events and system consequences and paths to recovery from identified threat events. The model is useful in analysing system failure case studies in a systematic way. Work will continue in developing the model and evaluating its utility. ■

REFERENCES

- Aleksander, I. and H., F. Keith. 1976. *Automata Theory: An Engineering Approach*. London, GB: Edward Arnold. ISBN 0-7131-2547-9.
- DoD (Department of Defense). 2001. USS Cole Commission Report.
- Haddon-Cave, C. 2009. The Nimrod Review. London, GB: The House of Commons.
- Jackson, S., and T. Ferris. 2013. "Resilience Principles for Engineered Systems." *Systems Engineering* 16 (2):152-164.
- NASA (National Aeronautics and Space Administration). 2009. "STS-125: The Final Visit [Hubble]." Last Modified 19 June 2009 Accessed 15 February. http://www.nasa.gov/mission_pages/shuttle/shuttlemissions/sts125/main/overview.html.
- National Geographic News. 2010. "Moon Landing Facts: Apollo 11 at 40." Accessed 29 January. http://news.nationalgeographic.com/news/2009/07/090715-moon-landing-apollo-facts.html.
- NTSB (National Transportation Safety Board). 2010. Collision of Metrolink Train 111 With Union Pacific Train LOF65-12 Chatsworth, California September 12, 2008. edited by D. A. P. Hersman. Washington, US-DC.
- Pariès, J. 2011. "Lessons from the Hudson." In *Resilience Engineering in Practice: A Guidebook*, edited by E. Hollnagel, J. Pariès, D. D. Woods, and J. Wreathhall. Farnham, Surrey, GB: Ashgate Publishing Limited.
- Wald, M. L., and K. Chang. 2007. "Minneapolis Bridge Had Passed Inspection." New York Times Accessed 29 April. http://www.nytimes.com/2007/08/03/us/03safety.html.

An SoS Analytical Workbench Approach to Architectural Analysis and Evolution

Daniel DeLaurentis, ddelaure@purdue.edu; **Navindran Davendralingam,** davendra@purdue.edu; **Karen Marais,** kmarais@purdue.edu; **Cesare Guariniello,** cguarini@purdue.edu; **Zhemei Fang,** fang59@purdue.edu; and **Payuna Uday,** payuna@gmail.com

ABSTRACT
This article summarizes the development of a System of Systems Analytic Workbench (SoS AWB) that provides a set of computational tools to facilitate better-informed decision-making on evolving SoS architectures. The workbench motif is adopted since SoS practitioners typically generate archetypal technical queries that can be mapped to appropriate analysis methods best suited to provide outputs and insights directly relevant to posed questions. After an overview of the workbench framework, four distinct methods currently available for use are presented along with their distinctive aspects in the concept of use.

INTRODUCTION

The importance of systems of systems (SoS)-derived capabilities documented in this edition of *INSIGHT* implies the associated importance of sound analysis tools with which to reason about development and implementation options for SoS architecture evolution. Evolving and refining a SoS presents significant decision-making challenges across both technical and programmatic domains. SoS generally involve integrating multiple independently managed systems to achieve a unique capability, therefore involving needs for collaboration and negotiation as well as control. In such complex systems, human behavioral and social phenomena in collaboration are critical as are cascading impacts from interdependencies; altogether, emergent outcomes are the norm. Handling such situations goes well beyond the immediate mental faculties of decision-makers and even capabilities of existing system-level decision-support tools. The current "cutting edge" in analysis for SoS seeks a collection of methods, processes and tools that provides the SoS practitioner with meaningful quantitative insights into projected SoS behavior and the possibilities for evolving the SoS, the set of options on system addition, deletion, reorganization required to meet the capability objective. Current policies set forth in the acquisition guidance documents, emerging SoS standards, and informal guidance, such as US Department of Defense (DoD) Systems Engineering Guide for Systems of Systems (U.S. DoD 2008a) and Defense Acquisition Guidebook (U.S. DoD 2008b), provide useful guidance but are in need of a supporting analytic perspective to complete the picture for more informed decision-making.

A number of research groups are working on advancements in this important area. Ongoing research is focusing on 'situational awareness' products for both SoS and constituent system-level decision-support as well as strategic approaches for modeling SoS architectures and their ability to restructure quickly to respond to failures, new needs and missions. In this short article, we exemplify this activity via overview of work in the area of SoS analysis methods funded by the DoD Systems Engineering Research Center (SERC). It is important to note, however, that analysis methods for SoS should be (and most are) applicable to civil/commercial applications as well, an especially relevant approach with emergence of 'smart, connected' cyber-physics system networks, Internet-of-things, and more.

One analysis framework developed and demonstrated via the SERC is the Flexible and Intelligent Learning Architectures for SoS, FILA-SoS, (Dagli 2015) developed to provide a decision making aid for SoS managers based on the 'wave model' (Dahmann et al. 2011) described in earlier articles of this issue. FILA-SoS adopts a complex system approach, for example, fuzzy inference systems and genetic programming, together with the 'wave model' processes to address four of the most challenging aspects of system-of-system architecting:

1. Dealing with the uncertainty and variability of the capabilities and availability of potential component systems

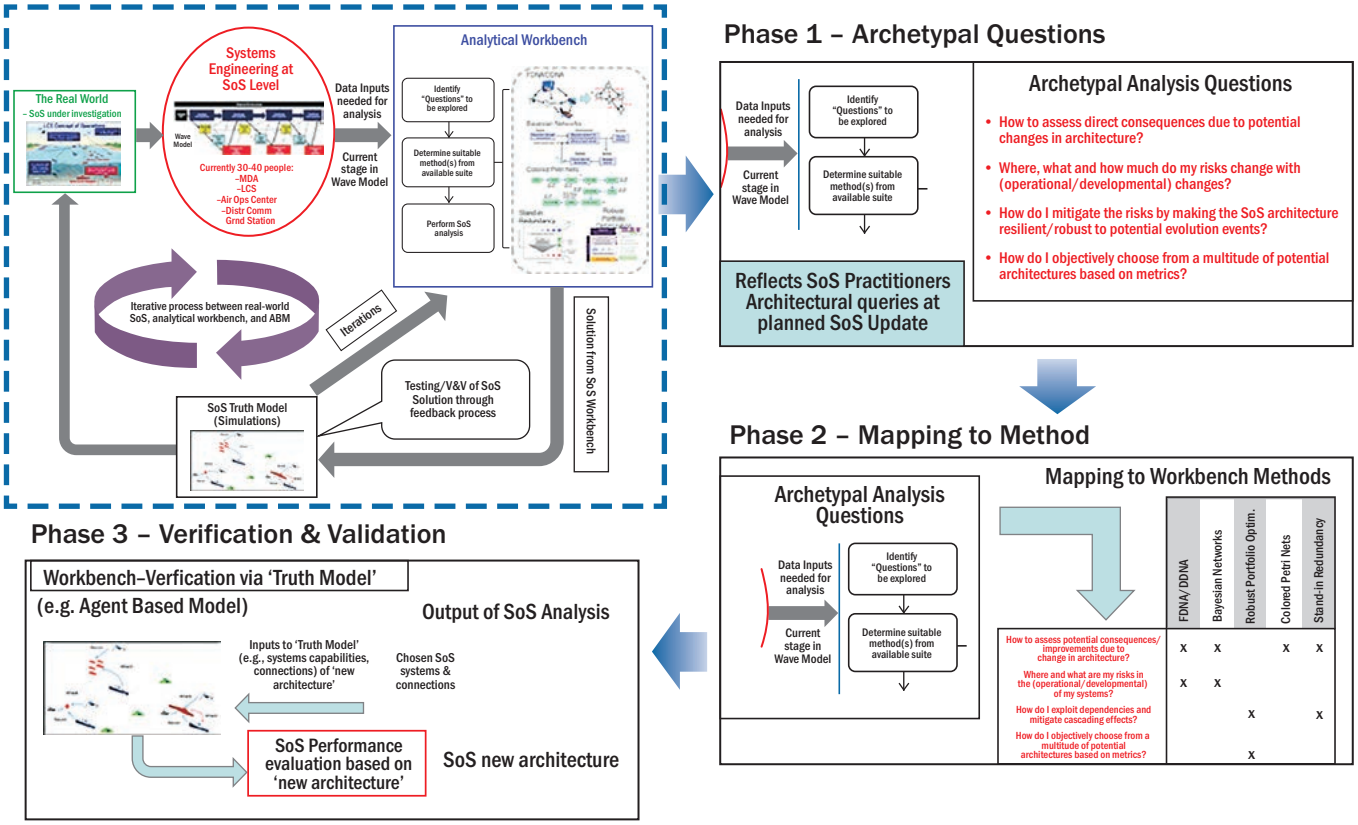

Figure 1. SoS Analytic Workbench overview

2. Providing for the evolution of the systems-of-systems needs, resources and environment over time
3. Accounting for the differing approaches and motivations of the autonomous component system managers
4. Optimizing systems-of-systems characteristics in an uncertain and dynamic environment with a fixed budget and resources.

The remainder of the article dives a bit deeper into a second example from the SERC SoS analysis portfolio. The Systems of Systems Analytic Workbench (SoS AWB) takes the approach of a set of computational tools to facilitate better-informed decision-making on SoS architectures. The work is motivated by the idea that SoS practitioners typically possess information and archetypal technical queries that can be mapped to appropriate analysis methods best suited to provide outputs and insights directly relevant to posed questions.

OVERVIEW AND CONCEPT OF USE: THE SOS ANALYTIC WORKBENCH

While there are many available methods, work to date on the SoS AWB has identified a set of theories and methodologies that have been adapted and expanded to support archetypal SoS decision-making. Method selection basis draws from evidence of prior case studies and subject matter expert views on the most pressing needs for analysis-based decision support. These methods are: System Operability Dependency Analysis (SODA)/System Development Dependency Analysis (SODA/SDDA), Systems Importance Measures (SIMs), Robust Portfolio Optimization (RPO) and Multi-Stakeholder Dynamic Optimization (MuSTDO). The philosophy of the approach is guided by similar 'workbench' type paradigms such as Lean Six Sigma where a collection of methods are employed to deal with quantitative aspects of lean manufacturing. However, the features of a SoS prompt the need for a collection of tools/methods to translate various technical complexities of the SoS tradespace into meaningful and actionable information for subsequent decision-making. Figure 1 illustrates the envisioned workbench and its primary phases of use in SoS level analysis and decision-making.

The iterative process of the workbench in Figure 1 (within dotted box) starts with an SoS practitioner's desire to explore a SoS tradespace for subsequent evolution; the evolution can involve the addition, removal or reconfiguration of the SoS architecture, based on desired objectives on achieving target performance (capabilities). The practitioner possesses data, such as from the DoD Architecture Framework (DoDAF 2.0) architecture description, a set of SysML models, or other sources, that describe the state of the current architecture and also of potential, yet-to-be introduced systems. The first phase involves the identification of 'archetypal questions' that typically arise from SoS practitioners' technically motivated queries on assessing the connection between objective metrics (SoS performance) against constraints such as cost and schedule. While not exhaustive, the questions that typically arise include: How to assess direct consequences due to changes in architecture? Where, what and how much do risks change with operational changes? How to mitigate risks? Which systems and connections should be added/removed? These typical questions also reflect the desire to examine the coupled behaviors that SoS exhibit, and their consequence on aforementioned metrics. Phase 2 involves the mapping of these archetypal questions, with data sets that are available to the SoS practitioner, to the relevant tool(s) in the workbench. The mapping may involve employment of multiple tools in the workbench due to overlapping analysis and decision-making requirements. The iterative process proceeds with Phase 3 in which the archetypal analyses of SoS practitioners are executed in concert with available 'truth

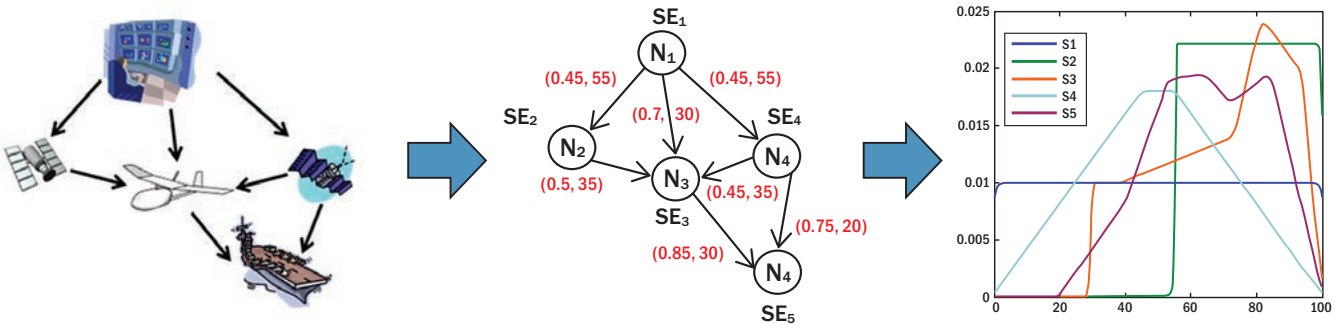

Figure 2. (a) Operational network (b) SDDA network representation (c) Operability analysis of nodes

models' (computational simulations, field testing) to provide preliminary verification of the next SoS evolution solution. The solution in this case refers to suggested architectural changes (addition/removal of systems and/or links) towards fulfilling target SoS capabilities, while preserving acceptable levels of risk (operational or developmental) and cost. The data generating capabilities of simulation (or field data) are paired with the higher-level architectural capabilities of the individual methods (as appropriate) to provide candidate solutions for subsequent implementation in the real world SoS environment.

SOS AWB ANALYTIC METHODS

System Operational Dependency Analysis (SODA)/System Dependency Development Analysis (SDDA)

The SODA methods assess operability, reliability, and resilience (ability to recover operability) in both operational and development contexts of SoS architectures (Guariniello and DeLaurentis 2016, 1). The architecture modeling is a directed network where nodes represent either the acquisition component systems or the capabilities. Figure 2 shows a notional example. Links on the network represent various kinds of dependencies between the constituent systems. The link types fall into two classifications: functional dependency in an operational network, or sequential development dependency in a development network. Each dependency is characterized by strength and criticality. Figure 2(b) shows the translation of the notional network in Figure 2(a) into a functional dependency network where the numbers in red are the strength and criticality of dependency. The ultimate goal of the technique is to analyze effects of such dependencies – and of their strength and criticality – on operability, and to identify valid operating and developing strategies and architectures. For operational networks, SODA is used to assess the effect of topology and of possible degraded functioning of one or more systems on the operability of the network. Figure 2(c) shows the stochastic SODA application on analyzing the effect of operational disturbances on the example network of Figures 2(a, b); here, the effects that systems 2, 3, 5 would experience given input disturbances for systems 1 and 4 are shown. For development networks, SDDA is used to assess how development time or capabilities are affected by the network topology and by delays in the development of component systems.

System Importance Measures (SIMs) Resilience Design

The Systems Importance Measures Resilience Design methodology is a four-phase method that highlights the relative importance of different disruptions via a prescribed set of measures (systems importance measures), and provides design guidance on how to improve the overall SoS resilience (Uday and Marais 2014). The four phases of the process are:

- Phase 1 – Identify potential disruptions ('What can go wrong?')
- Phase 2 – Determine impacts of disruptions ('What are consequences of unmitigated disruptions?')
- Phase 3 – Determine current SoS resilience ('How well is the SoS able to handle the disruptions?')
- Phase 4 – Improve SoS resilience using Design Principles ('What can be done to improve SoS resilience?')

The method provides a platform for multiple analysts and decision-makers to study, modify, discuss, and document options on implementing SoS resilience, in a fashion that scales with the SoS size. The visual nature of the resilience map provides a useful, highly intuitive, and immediate way for summarizing key points of concern in iteratively building resilience into an SoS. Figure 3 below illustrates the SIMs process.

Robust Portfolio Optimization (RPO)

The robust portfolio method adopts an 'investment-like' perspective in system of systems engineering, where the objective is to balance the risks in holding financial assets against the expected return on investment (Davendralingam and DeLaurentis 2015, 269). Here, we treat the SoS as a monolithic portfolio of systems that can be 'acquired' and/or 'connected' following feasible rules. In the context of a SoS, the expected returns correspond to an expected 'capability' due to investing in

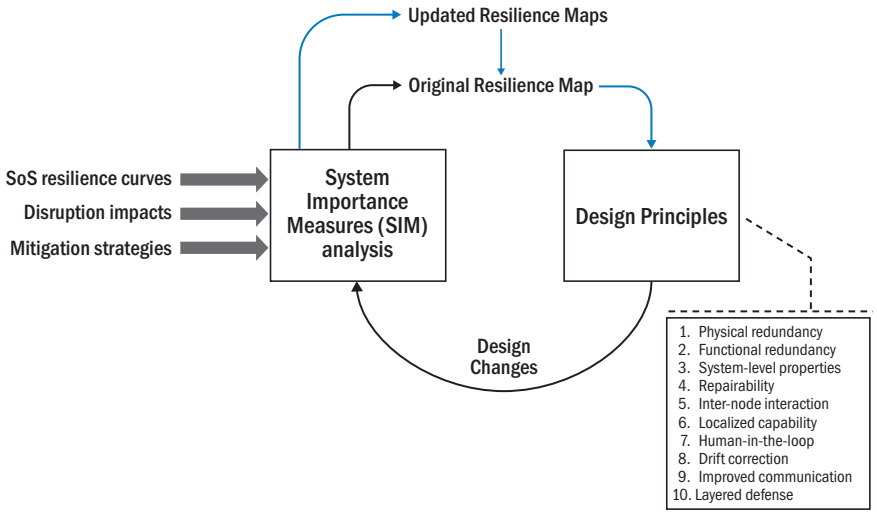

Figure 3. Systems Importance Measures (SIMs) design methodology

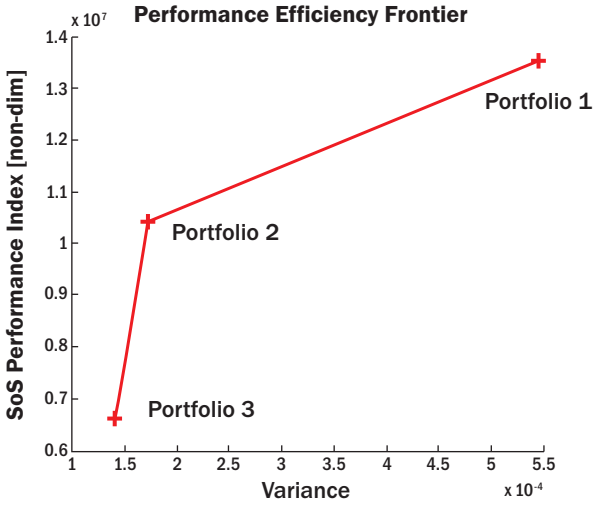

Mission Package	Portfolio 1	Portfolio 2	Portfolio 3
ASW Variable Depth	–	–	–
Multi Fcn Tow	X	X	–
Lightweight tow	–	–	X
RAMCS II	–	–	–
ALMDS (MH-60)	X	X	X
SUW N-LOS Missiles	X	–	–
Griffin Missiles	–	X	X
Seaframe Package System 1	X	–	–
& Combat Package System 2	–	X	–
Management Package System 3	–	–	X

Figure 4. (a) Efficient Portfolio Frontier (b) Portfolio constituent system description

a system, and the risks can be attributed to developmental or operational risks of the individual system. The method adopts state-of-the art developments in operations research with the objective of identifying optimal collections of systems that can give rise to a desired SoS level capability, given defined acceptable levels of risk. The resulting optimization problem accounts for uncertainties in the estimated uncertainties, for problems that can be represented in the method's abstraction of using nodes and specific rules of connectivity. The robust optimization formulation can ensure that identified 'portfolios' remain near optimal, by explicitly including the impact of uncertainty in the problem formulation. Figure 4 (a) exhibits illustrative results for the case of selecting appropriate 'portfolios' of systems from a candidate collection for the case of naval warfare scenario acquisitions. The graph shows the reward-to-risk efficiency frontier that shows the optimal SoS-level performance (here, a notional index) achieved, given a prescribed acceptable level of system development risk. The table shown in Figure 4(b) identifies the individual systems for each corresponding point on the graph; selection of the systems is also subject to various connectivity rules that ensure feasible collections of systems for a given architecture.

Multi-Stakeholder Dynamic Optimization (MuSTDO)

The dynamic nature of SoS evolutions and, decoupled nature of decision-making due to localized authority within an SoS (such as seen in an 'acknowledged' SoS) means that the interplay of tactical and strategic decisions can result in increased risks in an SoS evolution. This temporally coupled nature of decision-making, combined with ubiquitous uncertainty, further exacerbates the already existing complexities in SoS architectural decision-making. The MuSTDO framework uses concepts from operations research and portfolio optimization to provide objective, multi-stage solutions that balance impacts of near-term and long-term SoS architectural decisions (Fang et al. 2013). More specifically, the method uses algorithmic innovations from approximate dynamic programming and a coordination mechanism based on the idea of transfer contracts to enable decentralized, multi-stage decision-making between stakeholders. The idea here is to relegate the quantitative complexities of decision-making coordination to the algorithm, while delegating the decision-making and tradeoff assessment to the SoS decision-maker, within a coordinated quantitative framework.

SUMMARY

The distinctive features of decision-making problems in a SoS evolution context present unique needs for analysis methods and tools. In addition, the ability to quickly recognize the most applicable tools to a particular SoS analysis problem is important. This article highlighted some specific desired aspects of analysis tools in this context and provided a snapshot of one stream of recent work (the SoS AWB) that exemplifies a broader set of emerging capabilities that enable the benefits of SoS elaborated throughout this issue. The range of analysis tools contained in the present SoS AWB seek to address the most important archetypal analysis problems that an SoS practitioner, or a system owner seeking to thrive in a SoS, may encounter. ■

ACKNOWLEDGEMENTS

This material is based upon work supported, in whole or in part, by the US Department of Defense through the Systems Engineering Research Center (SERC). The SERC is a federally funded University Affiliated Research Center managed by Stevens Institute of Technology. The authors would also like to acknowledge Mr. Scott Lucero, SERC Technical Monitor, for valuable feedback and support of this research effort.

REFERENCES

- US DoD (Department of Defense). 2008a. "Systems Engineering Guide for System-of-Systems." http://www.acq.osd.mil/se/docs/SE-Guide-for-SoS.pdf.
- ———. 2008b. "Defense Acquisition Guidebook." https://dag.dau.mil/Pages/Default.aspx.
- Dahmann, J., G. Rebovich, and R. Lowry. 2011. "An 'Implementers' View of System Engineering for System of Systems." Paper presented at Institute of Electrical and Electronics Engineers Systems Conference, Vancouver, CA, 10-12 May.
- Guariniello, C. and D. DeLaurentis. 2016. "Supporting design via the System Operational Dependency Analysis Methodology." *Research in Engineering Design* 163:1-17.
- Davendralingam, N., and D. DeLaurentis. 2015. "A Robust Portfolio Optimization Approach to System of System Architectures." *Systems Engineering* 18 (2): 269–283.
- Uday, P., and K. Marais. 2014. "Resilience-based System Importance Measures for System-of-Systems." Paper presented at Conference on Systems Engineering Research (CSER), Redondo Beach, US-CA, 21-22 March.

- Fang, Z., D. A. DeLaurentis. 2015. "Multi-Stakeholder Dynamic Planning of System of Systems Development and Evolution." Paper presented at Conference on Systems Engineering Research, Hoboken, US-NJ, 17-19 March.
- Dagli, C. 2015. "Flexible and Intelligent Learning Architectures for SoS (FILA-SoS): Architectural evolution in Systems-of-Systems." Paper presented at Conference on Systems Engineering Research, Hoboken, US-NJ, 17-19 March.

ABOUT THE AUTHORS

[Editor: Author biographies were current when the paper was initially published in 2016.]

Dr. Dan DeLaurentis is Professor in Purdue's School of Aeronautics & Astronautics and Director of Purdue's Center for Integrated Systems in Aerospace (CISA). He is the co-lead of the Enterprises as Systems and System of Systems Thrust Area in the DoD's Systems Engineering Research Center (SERC) where he also is Principal Investigator (PI) on a project developing an analytical workbench for systems of systems (SoS) architecture analysis, design and evolution. He also is a member of the SERC's Research Council.

Navindran Davendralingam is a research scientist at Purdue's School of Aeronautics & Astronautics. He works as part of the Center for Integrated Systems in Aerospace led by Dr. DeLaurentis and received his Ph.D. in aerospace engineering from the Purdue University in 2011. Dr. Davendralingam currently works on projects that are/ have been funded by the DoD Systems Engineering Research Center (SERC) and the Naval Postgraduate School (NPS).

Karen Marais is an associate professor in the School of Aeronautics and Astronautics at Purdue. Her research interests include safety and risk assessment, systems of systems, and environmental impacts of technology.

Cesare Guariniello received a master's degree in automation and robotics engineering and a master degree in astronautical engineering at the University of Rome "La Sapienza." He then moved to Purdue University, where he earned a PhD in aeronautics and astronautics, under the supervision of Dr. Daniel DeLaurentis. Cesare is currently a research associate in the department of Aeronautics and Astronautics at Purdue University. His research expertise spans system-of-systems modeling and analysis, model-based systems engineering, space systems design and architecture evaluation, robotics, and remote sensing.

Zhemei Fang is a PhD candidate in the School of Aeronautics and Astronautics Engineering, Purdue University. She received her master degree from Huazhong University of Science and Technology in China in 2011. She currently works in the System-of-Systems Laboratory led by Dr. DeLaurentis. Her research interests revolve around system-of-systems architecture development, evolution analysis, conflict coordination mechanism design, and so on.

Payuna Uday is an aviation planning consultant at Landrum & Brown, responsible for tasks involving airspace and airfield modeling, gate requirement and scheduling, airport operations analysis, and aviation-related research. Dr. Uday received her PhD and Master's degrees in aeronautical systems engineering from Purdue University. She also holds a B. Tech degree in electronics and communication engineering from the National Institute of Technology, Trichy (India).

How Infrastructure Can Become Reborn by Becoming Born Robust

Josh Sparber, jsparbear5@gmail.com

■ **ABSTRACT**
Systems Modeling Language (SysML) is a tool for guiding engineers in designing power grid circuits sufficiently robust to withstand known electromagnetic pulses (EMPs). Careful examination of existing data shows that EMPs, and sometimes geomagnetically induced currents (GICs) that accompany EMPs are truly a powerful threat to power grid survival. Systems engineers, employing SysML can isolate power grid failure susceptibilities and areas for necessary power grid design improvements with selected SysML packages defined as enclaves associated with risk. These enclaves can be decomposable into stereotyped components available for risk categorization, building simulation libraries, or follow–on tests. As an example, a stereotype Source, instantiated as a Photovoltaic (PV) Inverter, increasingly important in microgrid renewable energy, is linked to a high frequency alternating current (HFAC) microgrid risk enclave package. Simulation allows evaluation of SysML use cases with EMP Actors. Real world test, construction, and strategic grid readjustment can then segue quickly.

■ **KEYWORDS:** electromagnetic pulses, microgrid, photovoltaic, satellite, simulation, SysML, use cases

INTRODUCTION: CONSTRUCTING THE ROBUST POWER GRID

Why This Paper Was Written

This article illustrates how using risk based profiles can build bottom up constructions of structures resistant to EMPs within a microgrid. With an assemblage of components tested for risk, simulated as SysML EMP threat use cases, an engineer can test and redesign selected portions of the power grid against EMP vulnerability and find paths to enhancing survivability or improving the reliability of the enclosing system.

The Power Law Nature of Power Outage Events

The Failure of Risk Management describes how a power outage is in a class of events that best follow the power law distribution rather than the normal distribution (Hubbard 2009, 184). The class members are "volcanic eruptions, forest fires, earthquakes, *power outages* [emphasis mine], asteroid impacts, and pandemic viruses (Hubbard 2009, 184)." These events can seriously affect "stressed systems that allow for both common mode failures and [a cascade of failures] (Hubbard 2009, 185)." Increasing complexity now exposes the grid to new vulnerabilities. EMP power outages have and probably will cause high–intensity failures more regularly now. EMP events have the strong potential to cause a more severe impact on the power grid than expected.

The Study of Solar Events

Solar flares and coronal mass ejections (CMEs) cause radio blackouts, radiation storms, and geomagnetic storms as effects (US National Oceanic and Atmospheric Administration. n.d.). Solar events can severely impact grid structures, as solar flares striking these earthbound structures or geomagnetically induced currents (GICs) from CMEs burning through these structures. Intensive Kepler satellite data shows that stars like our sun, G main sequence stars, are capable of flares much larger than previously recorded (Shibata 2016). According to data taken from 80,000 stars, even "superflares," flares 100 to 1,000 times anything previously recorded, can occur in sun-like stars as frequently as every 800 to 5,000 years (Shibata 2016).

Since World War II, scientists take solar dynamics into account, contributing to Earth weather. The science of helioseismology enabled scientists to study the internal solar "architecture" that leads to solar effects (Morrow 2015). While many satellites now fielded (SOHO, STEREO, SDO) observe both the nature of the sun and similar stars for gathering data for more effective predictions, scientists have a poor understanding of space weather dynamics and cannot predict it reliably (Thompson 2015).

An insufficiently prepared grid system could collapse even due to a relatively minor weather event, intentional mischief, or even decay or neglect. Increasing human and natural threats means that the probability of an EMP event causing damage may happen with greater intensity and greater regularity than in previous times. It makes sense to examine fortifications that will efficiently protect the existing power infrastructure. How can we build a power grid system that will provide this protection?

Table 1. A few microgrid component bdd stereotypes and possible specific ibds

bdd stereotype	device ibds	bdd stereotype	device ibds
Main Grid Components	Power Transformers, VVR–VVO circuits, High Power Caps	Sources	Solar, Wind, Small Hydro, Diesel, Geothermal, Co-Generators
Converters/Inverters	AC/AC, AC/DC, DC/AC, DC/DC	Protectors	Fuses, Ground Fault Detectors, Spark Gaps, Circuit Breakers
Storage/Recovery	Ultra/Supercapacitor, Flywheel/ Battery Farms, Compressed Air Energy Storage, Hydrogen Storage, or Pumped Hydro	Communication	Risk Applications, SCADA, Telemetry, GPS

BUILDING A MODEL: ANALYSIS

In the following paragraphs, the author will discuss an approach to organizing portions of the power grid for understanding risk through a model–based system model, and simulated for test based on that system model or an improved "should be" design with enough robustness to withstand impacting EMPs or GICs. The depicted power grid system is an HFAC microgrid. The author intentionally chose this grid utility—it is a hybrid between a microgrid fitted to conventional AC power and the currently popular and nonconventional DC microgrid. A solar voltage source serves as an example of a vulnerable component. Microgrids and the renewable energy of solar power will be having greater and greater impacts in constructing a robust grid in the coming years.

Starting with Block Diagrams

Friedenthal, Moore, and Steiner (2009) state that a generic type for a part can be "modeled in SysML with a block definition diagram, bdd (page 34)." While a bdd usage can be for a generic functional unit, specific standard functional units derive from the generic functional unit. We demonstrate the use of these derived block diagrams, called internal block diagrams (ibds) (Friedenthal, Moore and Steiner 2009, 44), later.

Stereotypes and Block Definition Diagrams

According to Friedenthal, Moore, and Steiner (2009), "Stereotypes are used…to customize the language for specific domains (page 542)." For example, functionally distinct portions of microgrids could be framed as microgrid bdd stereotypes: sources, transducers, converters and inverters, protectors, storage/recovery devices, communication devices, assets, interfaces, and the connected main grid power components. There is no display of management and control elements—only a few sample functional breakdowns.

Internal Block Diagrams

For recasting power infrastructure into a more robust system, SysML "provides a mechanism for dealing with legacy system elements that have not been developed using rigorous modeling techniques;" that is a 'profile (Friedenthal, Moore, and Steiner 2009, 307). ' This author is interested in building profiles within grid portions specifically oriented to risk. Groups of bdds characterize a particular enclave or super groups of bdds associated with particular types of risk information. Once fully characterized, this profile information could pass on to any grid portions of similar composition for both risk classification and risk comparison.

Ibds would be the actual realizable components for simulation and later test, within those bdds. Ibds would convert into SysML state machine diagrams for checking changes to internal parameters under EMP or GIC stresses. Ibds within this discussion are blocks within other blocks; however, most fully developed ibds show interconnections within and between each other and to other external bdds.

Look at the inverter below. The Table 1 bdd stereotype 'Sources' could be the bdd for the instantiation of a solar source ibd as the PV inverter in Figure 1. Another ibd instantiation could be a concentrated solar power receiver, another type of solar energy source. The American solar industry planned building 7 Gigawatts of solar stations of over 1 Megawatt capacity in 2016 (Greentech Media 2015). Photovoltaic (PV)inverters will play an important role. In Figure 4, the solar source ibd is a 'part' within the sources stereotype block.

Use Cases and Simulation

Manmade nuclear high-altitude electromagnetic pulses (HEMP) contain E1 (extremely short duration), E2 (moderate

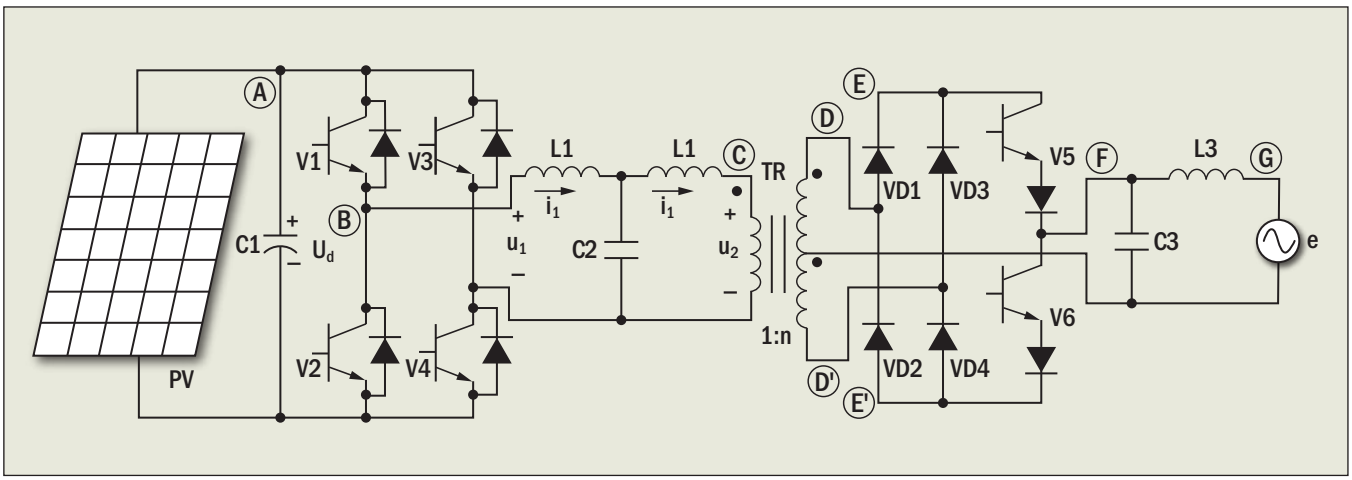

Figure 1. PV Inverter Circuit Source (Next Electronics 2016)

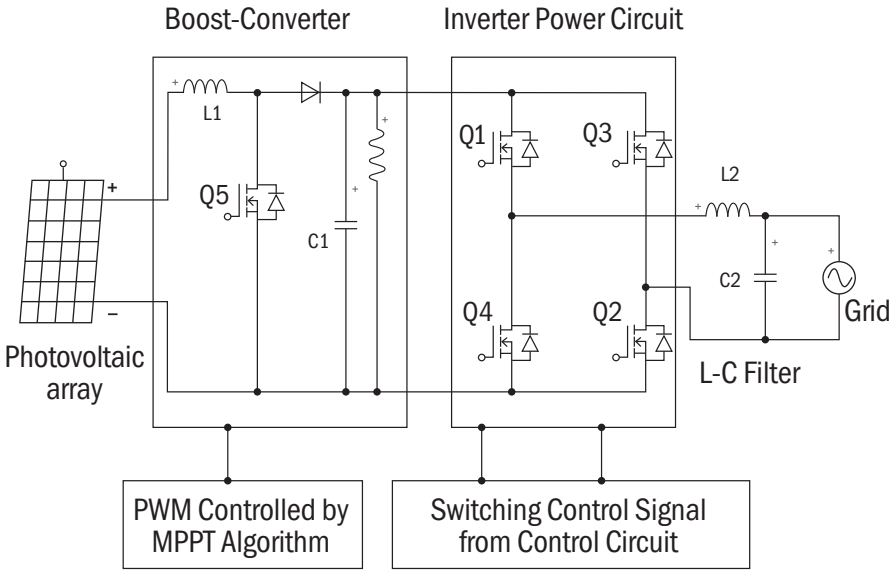

Figure 2. PV inverter source functional blocks leading to a detailed Matlab simulation (Mollah, Panda, and Saha 2016, Figure 1)

duration), and E3 (long duration) portions (Pry 2013, 5). E1 and E3 are the gravest danger to equipment (Pry 2013, 5). Naturally occurring EMPs; such as CMEs, have effects similar to E3; that is, they give rise to powerful GICs that travel through the ground—there is a strong claim for emphasizing testing power transformer neutrals to GIC exposure (Emanuelson 2013a, 5). Conspicuous interconnection length between parts or an interconnection length subject to harmonic resonance is vulnerable to E1 type pulses (Emanuelson 2013b). A simulation can highlight where circuit vulnerabilities exist to either an EMP pulse or a GIC moving through the ground, connectors, or vulnerable components (Hathaway and Byers 2015).

There are many ways of preparing a PV inverter circuit for simulation. Simulink can be used to build a simulation model of the PV inverter by connecting subcomponents (Osorio 2014). Full detail modeling of system parameters: voltage, current, thermal, or power properties, can also be done in a "Sim Power Systems" application (Osorio 2014), or in Matlab, see Figure 2.

Figure 3 shows Figure 2 as a possible ibd precursor to a state machine diagram ready for simulation. PTC Integrity Modeler 8.2 allows sufficiently detailed lower level use cases to be connected to simulations by allocating activities to state machine diagrams or adding 'constraint blocks' to ibds (PTC 2013b, 77–79, 86–89). In PTC Integrity Modeler 8.2, a 'SySim Profile' can also apply to an enclave of "at risk" components in preparation for simulation of a detailed lower level use case (PTC 2016a, 3). In PTC Integrity Modeler 8.2, ibds 'owned' by risk enclaves, with a 'SySim ControlBlock Stereotype' assigned, are ready for front–end analyses runnable in Visual Studio (PTC 2016a, 4, 11). Access to Simulink is also available in PTC Integrity Modeler 8.2 (PTC 2016a, 3).

BUILDING A MODEL: SYNTHESIS

"A package is a container for other model elements" (Friedenthal, Moore, and Steiner 2009, 81). In this discourse, generic functions are stereotyped bdds. Let packages contain these stereotyped bdds in an enclave of "at risk" components in a "risk profile." These enclaves would be microgrid portions available for EMP/GIC impact evaluation. Figure 4 is an HFAC microgrid. Create a stereotyped Sources bdd for Solar Power and Wind Power inputs. Then, finally instantiate the Solar Power bdd as an ibd of the PV Inverter circled below in Figure 4. The Main Grid component stereotype bdd could contain the ibds for the AC Line, Linear Load, Motor, and the Distributed Network. The Converter/Inverter stereotype could contain the AC/AC converter ibd.

A risk package enclave for a microgrid under review could ultimately contain the whole set of stereotyped bdds. See Figure 5. A larger package could form a "risk profile" of such components to be tested for robustness and reintegration that could be used to characterize all similar microgrid portions. Friedenthal, Moore, and Steiner (2009) state, "A profile is a kind of package used as a container for a set of stereotypes and supporting definitions (page 346)."

A microgrid level risk profile could contain all the bdds, use cases for the simulation, and any standards that would apply. See Figure 6. In PTC Integrity Modeler 8.2, a "risk profile" could be customizable as an 'ergonomic profile' (PTC 2016b). To create such a custom profile takes "a good working knowledge of the VBScript language, the modeler meta model and the modeler automation interface (PTC 2016b)."

CONCLUSION

To this author's mind, model–based system engineering can achieve a selective, strategic deployment of architectures

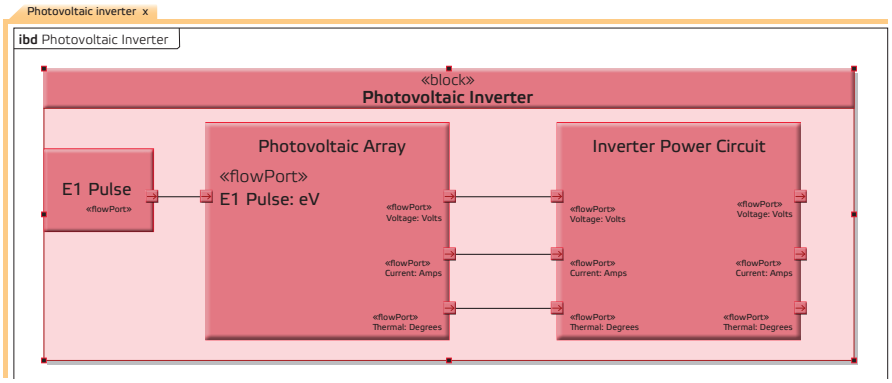

Figure 3. Ibd with state inputs shown

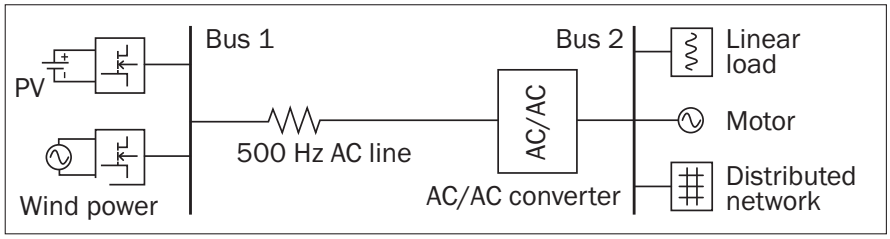

Figure 4. HFAC microgrid (Mariam, Basu, and Conlon 2013, Figure 6)

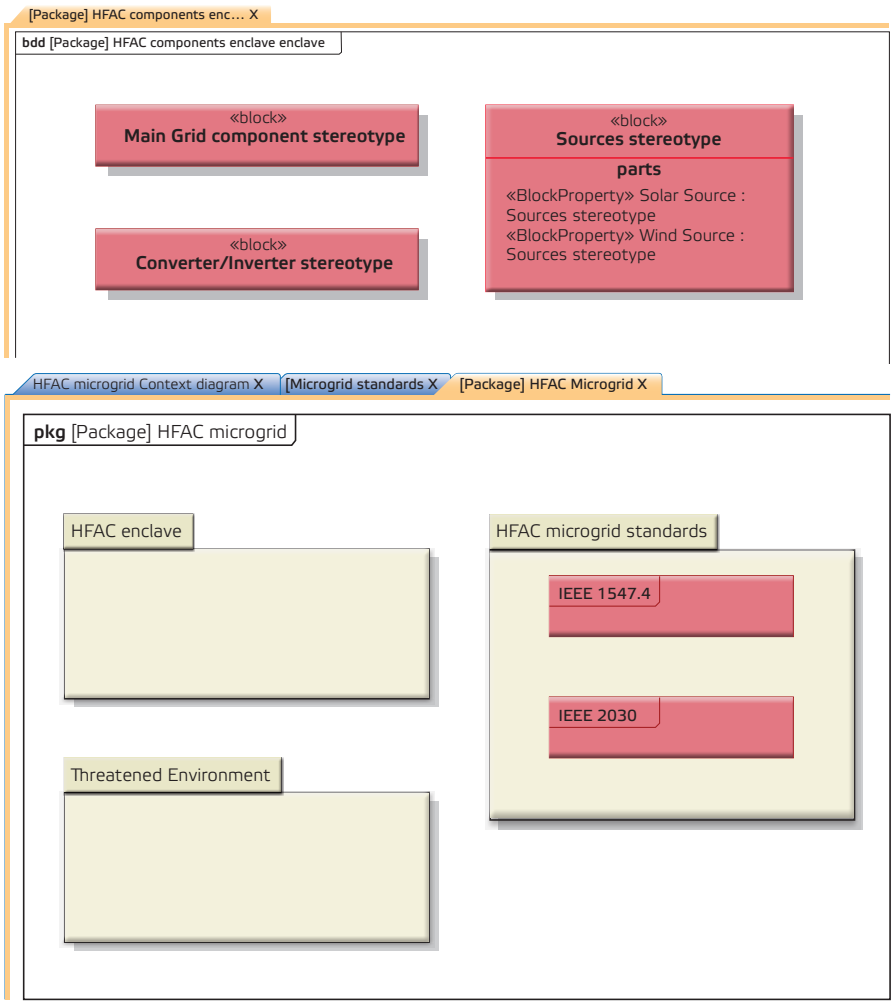

Figure 5 and Figure 6. Packages for HFAC microgrid (PTC 2016b; Etemadi 2013, 22; Basso 2014)

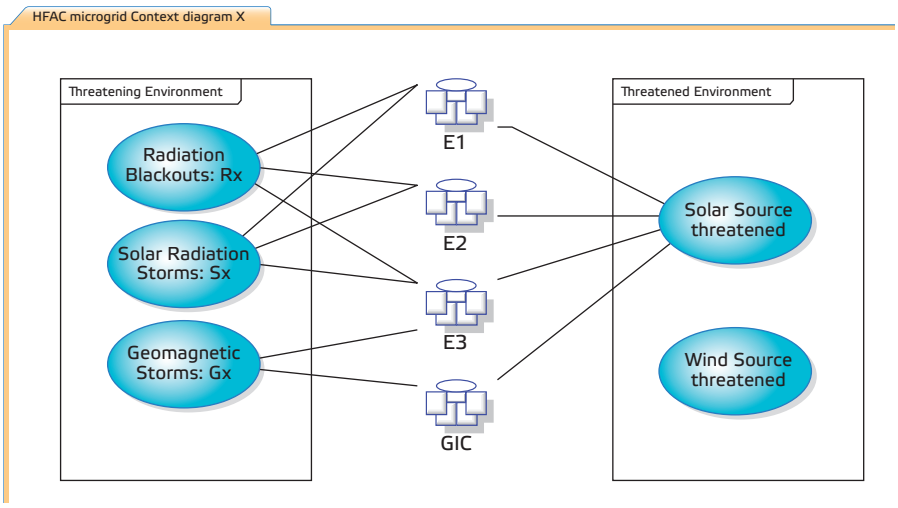

Figure 7. Context diagram with upper level use cases (US National Oceanic and Atmospheric Administration. n.d.). N.B.: HFAC threatened environment use cases are referenced in the HFAC microgrid package shown in Figure 6.

composed of robust elements through SysML. The use of groups of risk profiles, common to many power grid architectures, such as microgrids, could determine categories of components needed to be systematically simulated, prepared for test, and for possible later redesign and follow–on retest. Stereotyped components would help classify exactly what portions of a particular grid structure had needs to address vulnerabilities to various EMPs or GICS. A vast library of common simulations could be aligned to risk classifications available to any engineer.

Power grid structures, recombined with robust components, will possess high survivability after thorough test and needed redesign. Constructions based on intensively tested components would be field ready for incorporation into any portion of the power grid needing survivability enhancement.

Simulation will be key to facilitating more robust designs. Any new constructions will have a need for thorough test. Simulation allows for a great variety of architectures to be quickly and thoroughly tested before anyone commits to a bill of materials or a single rivet goes into a support structure. ■

REFERENCES

- Basso, T. 2014. "IEEE 1547 and 2030 Standards for Distributed Energy Resources Interconnection and Interoperability with the Electricity Grid." NREL. http://www.nrel.gov/docs/fy15osti/63157.pdf.
- Emanuelson, J. 2013 a. "E1, E2, and E3." http://www.futurescience.com/emp/e1-e2-e3.html.
- ———. J. 2013b. "EMP Myths." http://www.futurescience.com/emp/EMP-myths.html.
- Etemadi, A. 2014. *Microgrids: Operation, Control, and Protection*. Saarbrüken, Germany: Lambert Academic Publishing.
- Friedenthal, S., A. Moore, and R. Steiner. 2009. *A Practical Guide to SysML: The Systems Modeling Language*. Burlington, US–MA: Morgan Kaufmann OMG Press.
- Hathaway, K., and B. Byers. 2015. "LT Spice IV: Building and Simulating a Buck Converter," video. https://www.youtube.com/pGo8XlZ7U4s?si=IcTZggjGTwmC9YUJ.
- Mariam, L., M. Basu, and M. F. Conlon. 2013. "A Review of Existing Microgrid Architectures." *Journal of Engineering* 2013, no. 637614: 1–10. Dublin Institute of Technology. http://dx.doi.org/10.1155/2013/937614.

- Mollah, A. H., G. K. Panda, and P. K. Saha. 2016. "Single Phase Grid Connected Inverter for Photovoltaic System with Maximum Power Point Tracking." *International Journal of Advanced Research in Electrical, Electronics, and Instrumentation Energy*. http://www.rroij.com/open-access/single-phase-gridconnected-inverter-forphotovoltaic-system-with-maximum-powerpoint-tracking.php?aid=43213.
- Morrow, A. 2015. "A Look Back at NASA Solar Missions." US NASA Goddard Space Flight Center. Last updated Dec. 1, 2015. http://www.nasa.gov/feature/goddard/nasas-soho-celebrates-20-years-of-space-based-science.
- Next Electronics. 2016. http://www.next.gr/uploads/34/807PET21-photovoltaic-systems-Figure01.jpg.
- Osorio, C. 2014. "Modeling and Simulation of PV Solar Power Inverters," (video). https://www.youtube.com/watch?v=GnZFi9CzF9Q.
- Pry, P. V., Dr. 2013. "Electric Armageddon." Createspace: World Wide Web.
- PTC. 2016a. PTC Integrity Modeler SySim Tutorial. Version 8.2. http://www.ptc.com/model-based-systems-engineering/integrity-modeler.
- ———. 2016b. PTC Model–Based Systems Engineering Tutorial. Version 8.2. http://www.ptc.com/model-based-systems-engineering/integrity-modeler.
- Shibata, Kazunari. 2016. "Superflares on Solar type Stars and Their Implications on the Possibility of Superflares on the Sun." Lecture at the Space Weather Workshop, Broomfield, US-CO, 25-29, April. www.swpc.noaa.gov/sites/default/files/images/u33/final_shibata_SWW_2015.pdf.
- Thompson, M. J. 2014. "Grand Challenges in the Physics of the Sun and Sun–Like Stars." *Frontiers in Astronomy and Space Sciences* http://dx.doi.org/10.3389/fspas.2014.00001.
- US National Oceanic and Atmospheric Administration. n.d. "NOAA Space Weather Scales." Space Weather Prediction Center. http://www.swpc.noaa.gov/noaa-scales-explanation.

ABOUT THE AUTHOR

[Editor: Author biography was current when the paper was initially published in 2016.]

Josh Sparber, receiving air radar training in the USMC 1975, spent 20 years in the electronics industry, 1980 – 2000. After earning an MSEE from Cal State Fullerton 1999, he familiarized himself with the practice of systems engineering for 16 years with the US Department of Defense. He has been an INCOSE member since 2005 and received his Certified System Engineering Professional certification in 2007. He has also persisted in a lifelong interest in environmental issues. He received a Master's Degree in Environmental Policy and Management from the University of Denver in 2015, in which his thesis focus was the search for possible system solutions to rationally and robustly reconstruct the US National Power Grid against the impact of naturally occurring electromagnetic pulses. The focus on protecting the power grid from natural EMPs is only one small piece of a very large solution, but it helps energize the conversation on how to protect, grow, or extend our existing infrastructure. The model Mr. Sparber is currently building in PTC Integrity Modeler 8.2 he calls Risk Exploration Verification Integration Validation and Evaluation, REVIVE.

Effective and Efficient Preparation for the Unforeseeable

S. W. Hinsley, s.w.hinsley2@lboro.ac.uk; **M. J. Henshaw,** m.j.d.henshaw@lboro.ac.uk; and **C. E. Siemieniuch,** c.e.siemieniuch@lboro.ac.uk

ABSTRACT
This paper hypothesizes that a system-of-systems (SoS) that is not fit for purpose is so because it cannot implement the correct, timely, and complete transfers of material, energy, and information (MEI) between its constituents and with its external environment that are necessary to achieve a particular result. This research addresses the problem of maintaining a SoS fit for purpose after unpredictable changes in operation, composition, or external factors by creating a method, implemented as an engineering process, and supported by an analysis technique to enhance the affordance {"Features that provide the potential for interaction by affording the ability to do something" (Norman 1999)} of SoS constituents for MEI transfer and reveal potential undesirable transfers.

INTRODUCTION

This paper summarizes research addressing the problem of how to keep a system-of-systems (SoS) fit for purpose. The authors' observations motivated the choice of research topic from working in the defence industry with several companies over a career of 30 plus years. There appeared to be a tendency for delivered products and services that, although meeting their requirements, needed modification to maintain a desired capability from the composing SoS, and to be fit for purpose. Changes in the SoS situation, for example, the operational environment, requirement, or the SoS capability components often rendered the SoS unfit for purpose due to a combination of the two reasons below in various proportions:

- The SoS capability degraded and could no longer bring about the desired outcome for its intended design.
- The SoS needed to provide some different functionality from that of its intended design to achieve the desired outcome.

Causes of fit for purpose loss were dynamic and varied, often due to situational changes dictating that modifications to SoS constituent systems, to recover SoS fitness for purpose, had to be made close to, or at the point of utilization, frequently by the personnel working as part of the SoS using "workarounds." A workaround is an engineering solution that is sufficient but rarely optimal regarding efficiency or cost. Systems engineering shows that corrective action is most effectively and economically done early in the lifecycle, but we acknowledge that total avoidance of late-stage modifications is unfeasible, which poses the question "What could be done by suppliers to facilitate maintenance of fitness for purpose?" Let us note that one of the characteristics of a SoS noted by Maier (1998) is that it is evolutionary; this implies that development of a SoS always requires the adaptation of an existing (legacy) set of systems: there is no clear sheet.

RE-CONFIGURABILITY

The Need for Re-configurability

In his holistic approach to risk management Hopkin (2002) notes that risk is a "circumstance, action, situation or event (CASE) with the ability to impact key dependencies." Such impacts are equivalent to the SoS being unfit for purpose. More generally, the SoS being unfit for purpose (not able to do what the user requires) is often due to unforeseen circumstances, actions, situations, or events, so that personnel working as part of the SoS have to modify the constituent systems close to the point of employment so that they can converge towards their aims. If the necessary system modifications are not feasible, the users' subsequent failure to achieve their objectives may have consequences ranging from increased costs to loss of life or property.

CASE that might adversely affect a SoS fit for purpose can be in one of four categories: "known," "known-unknown," "unknown-known" and "unknown-unknown" (Rumsfeld 2002). Engineering actions can be taken to enhance fitness, and, the better known these CASE are (the more predictable the CASE), the more directly they can be affected by engineering actions (De Meyer et al 2002) such as design for robustness, resilience, and reconfigurability. Ring and Tenorio (2012) state "a system formalized by prescient design cannot respond to unforeseen situations." The realization of a systems capability is dependent on the simultaneous readiness of several components known as lines of development (LoD). The UK Ministry of Defence (MoD) eight defence

LoD, are a typical example, being:

- Training, equipment, personnel, information, doctrine, organization, infrastructure, and logistics.
- Let us note that this is not an exhaustive set of LoDs: it is likely that we need to consider others such as legal, commercial, and finance.

Robustness and resilience can be designed in a system, but any benefit they provide against unknown-unknown factors is serendipity. The major LoD contributing to fitness for purpose maintenance in this circumstance (the focus of this work) is often the personnel working as part of the SoS. In support of this, General Sir Rupert Smith states, "on every occasion that I have been sent to achieve some military objective to serve a political purpose, I, and those like me, have had to change our method and reorganize to succeed. Until we did this, we could not use our force effectively. From my lengthy experience, I have come to consider this as normal — a necessary part of every operation" (Smith 2005).

"Unknown-unknown" CASE poses a significant challenge, which we need to address. Building on Ashby's (1956) work on "requisite variety," Boardman and Sauser (2006) state "the uncertain and unknowable environment in which the SoS must operate presents a mystery of endless proportions, the only proper response to which is to have increasing variety, of a continually emerging nature, to deal with unforeseeable reality that eventually becomes clear and present danger."

The "How, Where, and Who" of Maintaining a SoS Fit for Purpose

SoS constituent and sub-system adaption has the potential to more effectively enhance MEI transfers at low levels than SoS reconfigurability, widely utilized because adaption actions are at a higher resolution and hence the adaption closely tailored to address a changed circumstance. Engineers employ reconfigurability and adaptation together in balance and proportion, tailored per case. The fit for purpose method, process, and technique facilitates SoS constituent system suppliers to equip their products and services affordably and conveniently with solution components, rather than solutions, to capitalize on the ingenuity and resourcefulness of utilizing personnel close to where they operate to efficiently and effectively address unforeseen changes when they occur. As Dalton (2013) commented, "but, ultimately, it is people who turn technology into capability; people who are experts in their profession with a comprehensive knowledge of the operational environment."

SOS CONSTITUENTS, TRANSFERS, AND AFFORDANCES

In 2008, the UK MOD defined defence LoDs as "the elements that must be brought together to deliver military capability to operational users" and states that "in addition to the defence LoDs, interoperability is included as an overarching theme that must be considered when any defence LoD is being addressed (UK MOD 2008)."

Systems engineering is "the management of the emergent properties (Burrowes and Squair 1999)." Emergent properties are not attributable to one component of the system, so similarly, systems and system-of-systems engineering has a strong focus on the interactions between constituents, and accordingly this research has a focus on the interactions between the constituents of a SoS. At the fundamental level, these interactions are transfers of matter, energy, and information (MEI). Thus, the designed operation of an instantiated SoS of interest depends upon the correct, timely, and complete transfers of MEI between the SoS constituents to achieve the purpose(s) of the SoS. This research identifies additional inherent and independent MEI sources, sinks, and bearers (SSBs) in a SoS constituent system not managed or captured by its defining documentation. These SSBs may cause undesirable emergent properties upon integration with other SoS constituents into a SoS or undergo exploitation to enhance the affordance for MEI transfer to address shortfalls.

SSB structures able to transfer MEI are affordances, defined by Sillitto at the INCOSE ASEC 2011 conference as "features that provide the potential for interaction by affording the ability to do something, as perceived by the user, to achieve some goal" (Sillitto 2011).

An illustration of the terms "intended," "inherent," and "independent" used to describe MEI transfers and SSBs may assist the reader here. For example, a maritime surveillance radar system is the system of interest (SoI). To electrically supply the radar control cabinet the designer specified an intended MEI (electrical energy) transfer from one of the ship's supplies to

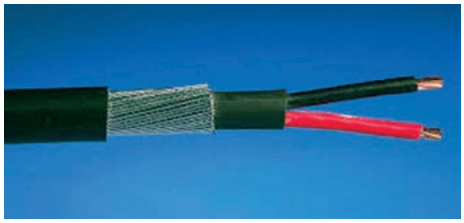

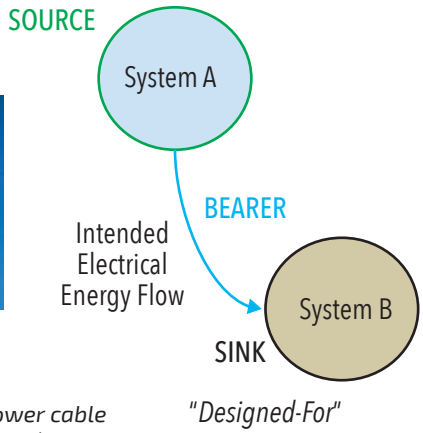

Figure 1. Steel wired armor direct current power cable and intended MEI transfer (DC electrical energy)

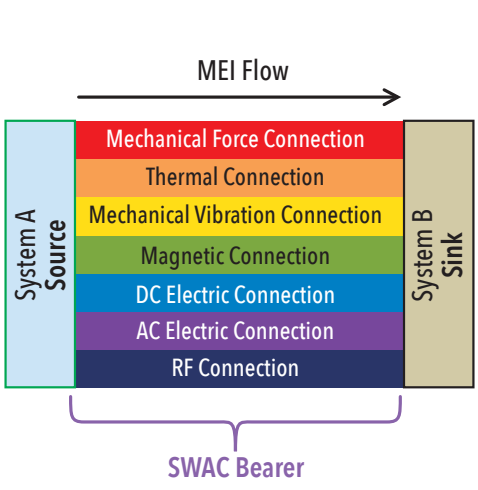

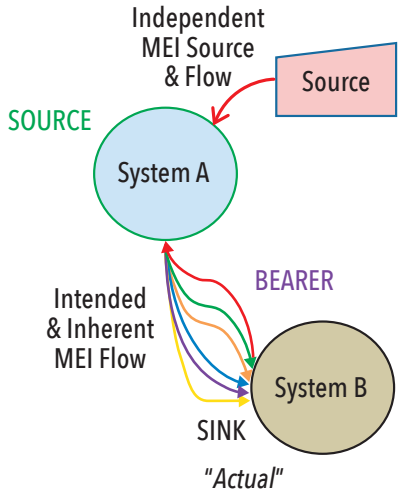

Figure 2: SWA DC power cable and inherent MEI transfers

the cabinet by a steel wire armored (SWA) cable.

This chosen bearer solution has inherent properties that enable it to conduct much more than DC electrical current, however, as illustrated on the left of Figure 2.

In this case, the SWA cable's inherent mechanical rigidity (mechanical energy bearer) of the power cable interfered with the correct operation of the cabinet's shock mounts. In addition, as shown on the right of Figure 2 (previous page), the structure of the vessel (mechanical energy bearer), the power cable, conducted vibrations from an independent MEI (mechanical energy) source, the vessels propulsion engines, to the radar antenna mast (mechanical energy sink) and degraded the radar's stabilization performance.

This section related capability elements to SoS and system constituents and introduced the concept of a systems MEI transfers as consisting of SSB in intended, inherent, and independent forms. The section following describes the fit for purpose method employing these concepts.

THE "VEE" MODEL AND THE FIT FOR PURPOSE METHOD

A common representation of the systems engineering process, used for illustrative purposes only here, is the "Vee" diagram shown in Figure 3 (INCOSE 2009).

The product lifecycle management (PLM) artefacts produced at each stage of the "Vee" by different companies implementations are functionally similar but tailored to their individual needs and constraints.

The fit for purpose method is a transform cascade, as shown in Figure 4. The cascade and analyses accommodate SoS constituents that are SoS. Although the cascade below suggests a waterfall process, in practice there is feedback, concurrent development, and iteration between the transformations. The data flows in Figure 4 correlate to the left-hand side of the systems engineering "Vee" model described in the next section.

The method facilitates changes in system capability desired to improve or expand the capabilities of SoS constituent systems to perform system level tasks, as well as those contributing to SoS capability level tasks.

- Transform 1 (top left) relates the SoS capabilities in the context of its operational concepts, to the MEI transfers across its boundary that result in the effects that the SoS is desired to have. Mission threads dictate the content and sequence of these MEI transfers producing the desired effects.
- Transform 2 in the cascade is a similar transformation to the first transform but at the SoS constituent system level.

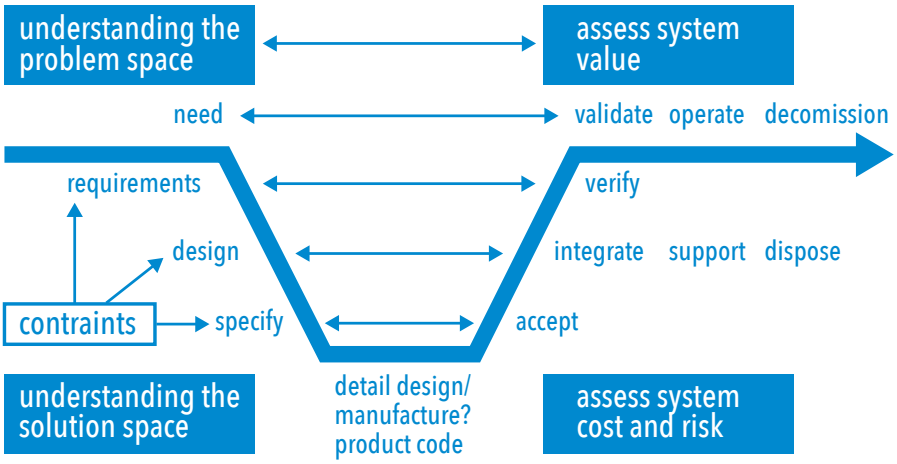

Figure 3. The "Vee" diagram, a common representation of the systems engineering process, which ranges, from conceptual models to assist comprehension of complex systems development to detailed product lifecycle and management models (©INCOSE UK Ltd, reprinted with permission)

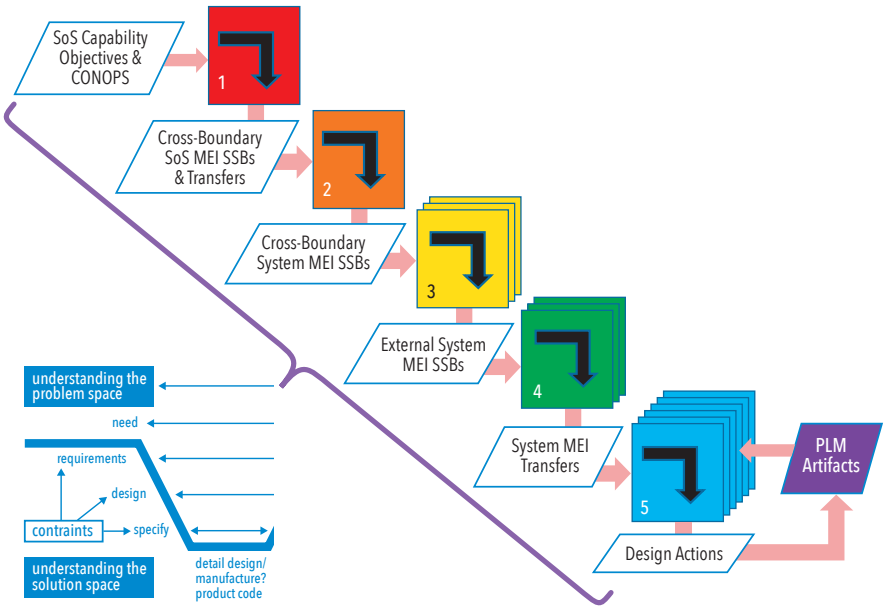

Figure 4. The FFP method – a cascade of transforms relates top-level SoS capabilities down to constituent systems via MEI transfers and SSBs to inform MEI enhancement

- Transform 3 groups all the prospective SoS constituent system MEI transfers into a set of system affordances for MEI transfer and identifies the major subsystems of interest.

This better enables examination and assessment of MEI transfer enhancement from a subsystem viewpoint.

- Transform 4 analyzes the affordances for MEI transfer and determines a subset as candidates for enhancement, by assessment at the system/subsystem level by the relevant specialist discipline engineers.
- Transform 5 associates system design actions with the system MEI transfer enhancement candidates, guided by the original system design actions and any others that are concurrent with subsystem MEI transfer enhancement.

The bottom-right "PLM artefacts" represents the project lifecycle management (PLM) documents, engineering drawings, computer aided design/computer aided manufacturing (CAD/CAM) models, and more into which we integrate the enhancement design actions with other planned actions.

Enhancement of system MEI transfer affordances are enablers for new system functionality and SoS capability. Figure 5

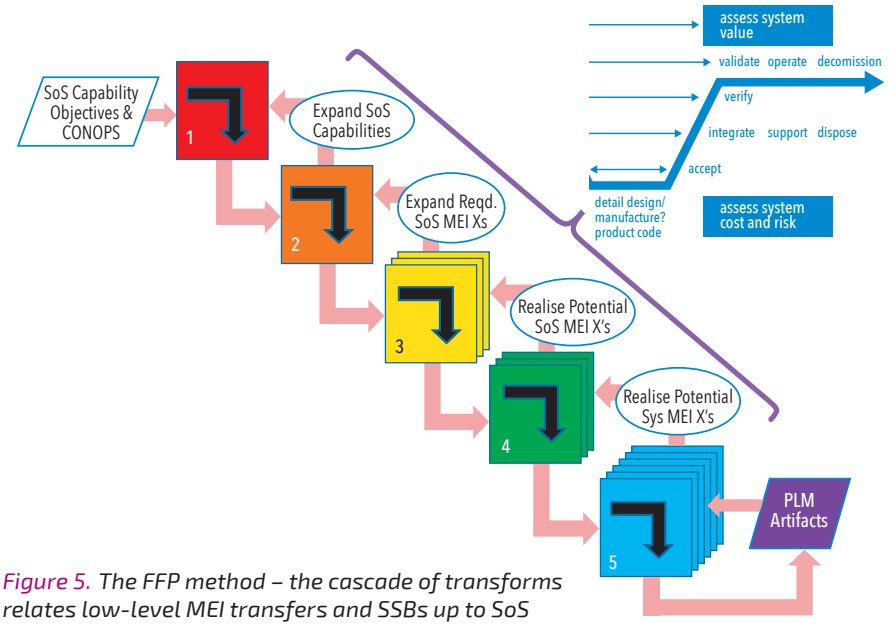

Figure 5. The FFP method – the cascade of transforms relates low-level MEI transfers and SSBs up to SoS capabilities via constituent systems

illustrates this and its correlation to the right-hand side of the systems engineering "Vee" model.

An MEI transfer enhancement activity is something one can think of regarding its own systems engineering "Vee" model superimposed on the "Vee" model of the system undergoing enhancement.

THE FIT FOR PURPOSE APPLICATION PROCESS AND TECHNIQUE

The fit for purpose process is an instantiation of the fit for purpose method tailored to the user organization's particular product lifecycle management (PLM) system. The fit for purpose process refers to PLM processes and utilizes PLM project artefacts (user and system requirements, system design, and test specifications) familiar to project staff to reduce the opportunity for error and maintain fidelity with the project/system-of-interest. In systems-oriented engineering companies, the product engineering process and the systems engineering process are often the same.

The notion of fitness for purpose "maintenance" is used to emphasize that this process applies at any stage in the lifecycle shown by the "Vee" diagram in Figure 3 as an opportunity arises. The FFP process is not a new approach to design. It offers a new perspective on projects and the engineering process. Analysis used to populate the matrices identifies opportunities to realize affordances at design opportunities such as scheduled major maintenance intervals, obsolescence resolutions, mid-life improvement (MLI) programs and improvement through spares programs.

To assist designers, we suggest a three-stage analysis technique. First, the designer is to identify the intended ('designed-for') MEI transfers, constituent MEI SSBs, and characterize them in the frequency domain to identify the bandwidth over which they are able to operate using a cognitive 'seismic to light' sweep. The designer then repeats this cognitive sweep for the inherent and independent SSBs that may affect the SoI. From this can be identified MEI SSBs that have the *ability* to transfer MEI due to their areas of common bandwidth, and hence form prospective MEI transfers. Next, an examination of the prospective MEI transfers in the time domain to promote those with the *opportunity* for MEI transfer by their SSB components being active at the same time to potential MEI transfers. Thirdly the susceptibility of the potential MEI transfer sinks to MEI conducted to them by the bearers is assessed to determine if any potential MEI transfer *can* be either problematic or present an exploitation opportunity to usefully enhance the SoIs affordance for MEI transfer.

The architecture and characteristics of the intended MEI transfers and the characteristics of their constituent MEI SSBs are in the product technical data pack (although probably distributed across several engineering disciplines and represented in several different formats) used to manufacture the product. The inherent and independent MEI SSBs are likely to be only sporadically captured by a few niche engineering specialties, for example, signature management, information architecting, spectrum management/scheduling, and process engineering. A suggestion for bringing these disparate parts together is an MEI meta-model that would provide a reference in the context of a product's technical data pack. Figures 6 and 7 on the following page illustrate how such a meta-model provides a complete view of a system of interest, in this case, a SoS with a single input and output.

The vehicle for data capture and visual analytics can be chosen by the fit for purpose adopter, and will be probably be dictated by what application integrates into their PLM system; this could perhaps be the Systems Modeling Language (SysML). At the time of writing, in use are a simple spreadsheet and open source bioinformatics software platform for visualizing molecular interaction networks.

FIT FOR PURPOSE SALIENT POINTS

Inherent & independent MEI transfers almost certainly will not appear in a product's design definition: they will not be controlled or managed, but they may be either problematic or being utilized by some stakeholders such as users and maintainers and thus will cause problems when they change or are subconsciously withdrawn from the product due to through-life development and modifications. For example, mechanical connections formed by electrical cables to a cabinet may transmit harmful shocks and vibrations to sensitive components within it. Specification compliant replacement components may not have the design margins of original components exploited by operators and maintainers. One way of encouraging considerations of a system's intended, inherent, and independent MEI transfers could be to make an fit for purpose analysis part of a project's systems engineering management plan.

As with systems engineering effort, MEI transfer enhancement may involve some additional cost, and the 'how much is enough' question requires an answer by the practitioners based on what they feel will provide the preferred cost-benefit to their particular case. However it also is an investment for the future which will reduce future implementation risk, reduce operational benefit latency and, by taking advantage of design opportunities, reduce the overall cost of maintaining fitness for purpose by engineering system capability currency to operational needs. MEI transfer enhancement should provide returns similar to other preparations for the future, such as product line architecture reduction, future spares provisioning, and "fitted for but not with" strategies. Commercial arrangements could share risk and benefit between customer and supplier, but stakeholders must assess the business case on a whole-life basis that includes the cost of upgrades.

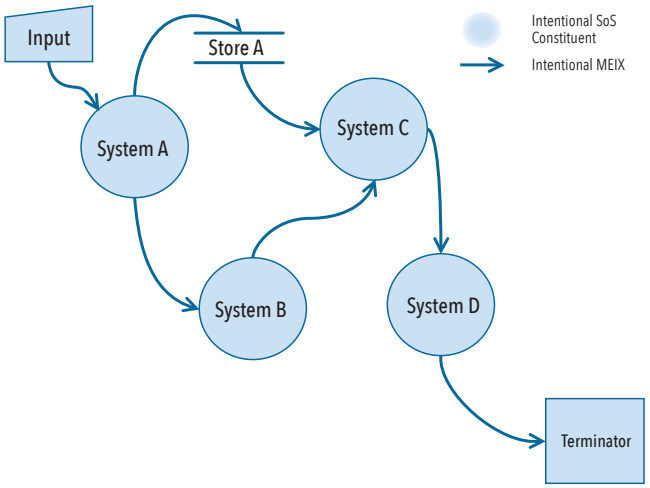

Figure 6. "What we think we have…": A MEI transfer diagram of a simple system-of-systems showing the intended ('designed-for') MEI transfers captured by the technical data pack(s)

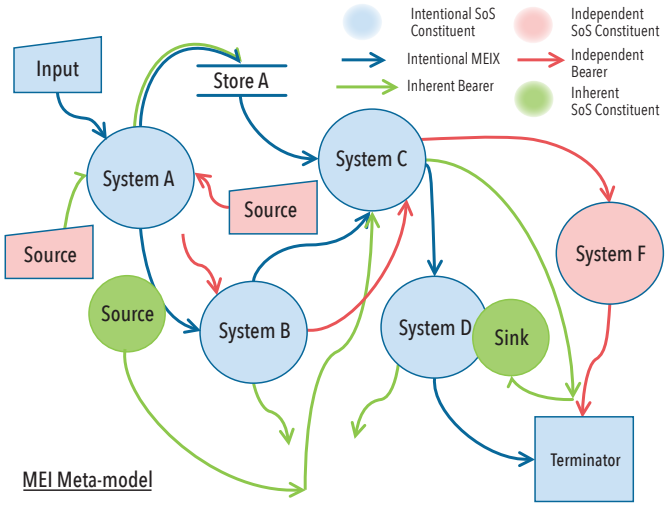

Figure 7. "…and what we actually have": A MEI transfer diagram of a simple system-of-systems showing the intended, inherent, and independent MEI transfers and SSBs seldom captured or managed.

Engineers can tailor the level of provision made for MEI transfer enhancement to the needs of the business: it may range from design only, functional models, virtual prototypes, board layout, and fitted components to live spares. These provisions could enhance MEI transfer enabling operational augmentation at the system level as well as at the SoS level.

Integration of the fit for purpose method into a supplier's PLM system will benefit the "creating" system by enhancing it to provide a complete understanding of "created" systems by capturing information that previously may have only been tacit. This process also identifies the major subsystems that will be affected by MEI transfer affordance enhancement and facilitates an examination by specialist engineers organized into work breakdown structure (WBS) subsystem teams that are familiar and experienced in specific areas. Fit for purpose embodied into a PLM system can examine a very large data set promptly for inherent and independent MEI transfers than could be problematic or provide opportunities that would be impractical otherwise.

An automated fit for purpose process will require the bandwidth, duty, and sink susceptibility characteristics of the MEI SSBs in the system definition to hold digitally in the hosting PLM system. These characteristics will exist as component libraries, engineering models, and schematics generated by specialist engineering disciplines and systems design artefacts generated by the project system engineers. However, the MEI SSB characteristics may not all be in a machine-readable form. The potential industrial exploiter will need to do a trade study between the desired degree of fit for purpose automation and the amount of machine-readable data currently in their PLM system with the work necessary to achieve the level of machine-readable data commensurate with the allocation of function (either manual, semi-automated, or fully automated) that they feel most cost-effectively delivers the desired benefits from incorporating fit for purpose into their engineering processes.

Designers incorporating enhancements into systems enablers for MEI transfers, which may come into play at some time in the future, need to be aware of the capabilities and facilities available to those enabling the enhancement. At first line, close to the point where the SoI is utilized, personnel will have fewer resources than are available at second line (deployable support and repair, field maintenance), and similarly less than those at third line (base workshop).

Any fit for purpose generated design actions need to harmonize with concurrent actions and existing processes and procedures. Engineers should incorporate MEI transfer enhancement and risk mitigation design actions on the selected major subsystems with the company PLM system and be able to integrate with other concurrent design actions, for example, those implementing MLIs, carrying out maintenance or repairs.

A fit for purpose goal is to stimulate thought that creates new design actions to realize a potential MEI transfer at the system and hence SoS level to enable a new SoS capability or mitigate risks that may only appear upon product deployment. We restrict the fit for purpose analysis as it appears in this paper to MEI spectra and duty: individuals may think of other parameters to add to the project definition to facilitate fit for purpose maintenance. This research is not intended to offer a universal and complete solution; it is a contribution that provides a complete view of a system which may well stimulate adopters to produce similar analyses tailored to best benefit their products and services.

Fit for purpose is neither a substitute for knowledge and wisdom nor a panacea for all ills. Fit for purpose adopters may well decide that fit for purpose would not provide an acceptable return-on-investment if applied in areas where nearly all the transfers are of one type, such as data handling and information management or where there is little latitude for modification.

FIT FOR PURPOSE ILLUSTRATION: OBSOLESCENCE RESOLUTION

This section illustrates how an obsolescence recovery exercise provided an opportunity to enhance a SoS constituent system's affordance for MEI transfer using the method and process described earlier. We identified the system's intended MEI transfers at the transform 4 on the cascade in Figure 4 and extrapolated both up the transform cascade to SoS capabilities and down to subsystem level respectively using requirements and design information in the PLM system. We identified intended, inherent, and independent MEI SSBs, characterized, and collated into an MEI meta-model, potential MEI transfers which we subsequently analyzed for risk and opportunity. We formulated design actions to reduce risks to an acceptable level and capitalize on opportunities to enhance system functionality and SoS capability as shown in Figure 5.

A Deck Approach Light Projector (DALP)

To assist aircraft landing on an aircraft carrier, an array of lights on the deck project beams towards the pilot to indicate the movement of the ship and their aircraft's deviation from the ideal approach angle and landing point. A "wave-off" (WO) lamp in the array illuminates if it is necessary for an approaching aircraft to abort the landing attempt. Figure 8 below shows a DALP equipment fitted to a carrier.

Obsolescence of some of the original equipment filament bulbs provided an opportunity to reduce downtime and maintenance cost by capitalizing on advances in light emitting diode (LED) technology. There were no requests for explicit enhancements to the DALP equipment at the time of the obsolescence resolution exercise, that is not in response to a function or performance upgrade requirement; however a fit for purpose analysis shows an opportunity to enhance the DALP's ability to transfer MEI in parallel with obsolescence resolution at little extra cost.

Figure 8. A deck approach light projector mounted on a naval vessel (Burrowes and Squair 1999) (photo courtesy of www.netmarine.net under creative commons https://commons.wikimedia.org/wiki/File:FS_CdG_Optics.jpg)

MEI Transfer Affordance Opportunity

The DALP WO MEI transfer affordance is a lamp that when flashed at 1Hz instructs the pilot of an approaching aircraft attempting to land to abort the landing. The fit for purpose method analysis determines the potential for MEI transfer enhancement of the DALP's MEI affordances and SSBs and identifies the potential of the WO affordance for enhanced information transfer.

The on/off and off/on response time of a LED is much faster than that of a conventional incandescent lamp which requires a finite time to heat up and cool down, so the faster switching characteristic of a LED lamp could provide an opportunity to enhance information transfer of the WO affordance, by modulating its output. The enhanced WO affordance would provide the carrier an available, non-broadcast, secure, low-latency, and jam-resistant communications link to an approaching aircraft.

The new LED lamps require new driver circuits, controlled by the DALP data bus. (Note: circuit is for illustration purposes only, and we do not show the data bus connections). The new circuitry required for the new LED lamps is on the left-hand side of Figure 9.

The new LED WO lamp driver circuit has a modulation input added to it. The circle shows the modification to the new WO lamp driver circuitry on the right-hand side of Figure 9.

The MEI transfer (in this case information) enhancement facilitates a new, secure, 'un-jamable' data transmission from the carrier to an approaching aircraft. The carrier houses extra functionality not occupying processor and memory space on the aircraft where it is at more of a premium.

DALP Enhancement Considerations

The DALP enhancement costs are small, as in this instance the equipment and associated documentation were to undergo modification already, and the number of DALPs in-service is small. The enhancement above does not significantly increase component count and type, and in this case should not significantly affect costs from bought-out materials, testing, equipment support publications, training, and more. The WO enhancement modification incorporates into the design actions forming part of the DALP project plans and documentation in accordance with the company lifecycle management process.

Fit for purpose analysis exploiters need to decide if there is merit at a design opportunity such as the DALP obsolescence recovery task in analyzing both the "before" and "after" situations. In the case of this obsolescence recovery task the move to LED lamps reduced the WO lamp energy transfer by a significant reduction in the infra-red (IR) region. This meant that although the new lamps retained the night vision goggle (NVG) and forward-looking infrared (FLIR) visibility of the old lamps, the inherent IR energy transfer of the LED lamps was insufficient to prevent ice build-up in artic conditions. Thus, we had to make a de-icing provision that the inherent IR output of the old filament lamps fulfilled.

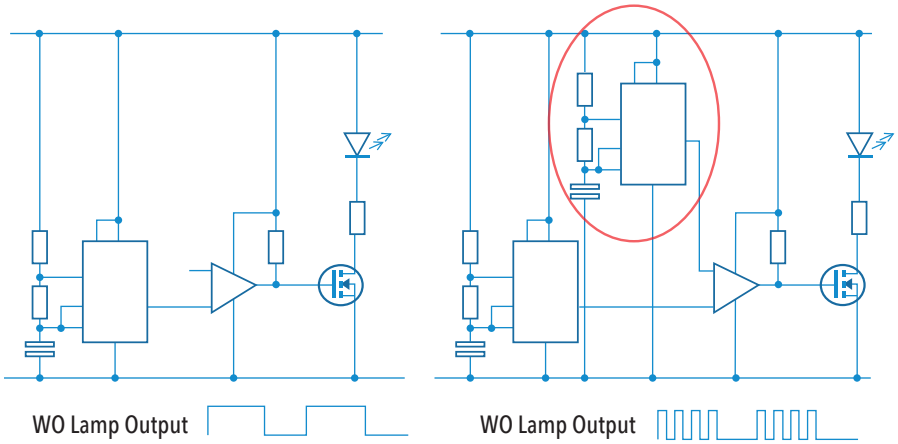

Figure 9. Illustrative LED lamp driver circuit and WO lamp driver circuit enhancement

DALP MEI transfer Affordance Enhancement Exploitation

The aircraft carrier is a central part of a SoS (a carrier group of vessels) that is conducting military operations in a littoral scenario, whose purpose is to provide intelligence/surveillance/target acquisition and reconnaissance (ISTAR) capability to friendly forces ashore. An unforeseen change in the political situation around the carrier groups operations meant that some military tasks achieved by manned aircraft became untenable. Because of the external change, the carrier group was no longer fit for the purpose of providing an ISTAR capability.

We examined candidate solutions to determine their feasibility, impact, and timeliness on both the problem and the capabilities available from the SoS resource, which included the necessary changes and enhancements to the MEI transfers. We chose a preferred solution of unmanned air vehicle (UAV) operations from a candidate set. Available UAVs have a core capability of operating from land, but do not have the ability to operate from a carrier. The UAVs are normally landed by a human pilot under remote control, but the latency in the control loop is too large to enable the remote pilot to compensate and make adjustments for the movements of the carrier at sea. The enhanced carrier to aircraft information transfer via the light projector provides a command link to an unmanned aircraft via its panoramic IR/TI camera, auto-tracker, and flight control system. This transfer provides a low-latency minor control loop to relieve the pilot of compensating for the movements of the vessel, enabling him to apply the flight commands to the UAV landing on the carrier much as he would do for a landing on the ground.

The provision made by the supplier during the DALP obsolescence recovery task came on-line whilst the carrier group was on-station, and enabled the available UAVs to operate from the carrier, thus maintaining fulfillment of the necessary ISTAR tasks without placing pilots, and expensive aircraft containing sensitive intellectual property in harm's way.

CONCLUSION

The interconnectivity and concomitant complexity of systems is rapidly increasing, meaning that engineers must now think in terms of the fitness of purpose of systems of systems, as opposed to single systems. The FFP method is offered as an holistic-thinking approach that will assist engineers identify concealed MEI sources, sinks, bearers, and transfers not otherwise included in SoS system definition which may lead to unexpected emergent phenomena either problematic and only revealed "late in the day", or being employed in-service unbeknown and not under the control of the design authority. Examples of such emergence in both naval and land-based domains emerge during the test and development of this work. The more complete insight from this approach enhances delivered products/services, and also improves the PLM engineering processes used to create them, and facilitates the task of identifying and implementing SoS adaptation (through reconfiguration at the SoS level or changes to constituent systems) in order to maintain that SoS as fit for purpose. This paper has provided an overview of the fit for purpose method and indicated how it may be implemented in organizational processes such that it may offer improved management of a complex SoS, enabling it to be maintained as fit for purpose to address new unforeseeable tasks, and/or changes, both internal and external, throughout the lifecycle at an affordable scale. ■

REFERENCES

- Ashby, W. R. 1956. *An Introduction To Cybernetics*, London, UK: Chapman & Hall Ltd. p. 202-208. Available at: http://pespmc1.vub.ac.be/books/IntroCyb.pdf.
- Boardman, J., and B. Sauser. 2006. "System-of-Systems – the meaning of "of."' Paper presented at the IEEE/SMC International Conference on System-of-Systems Engineering, Los Angeles, US-CA, 24-26 April, 118–123.
- Burrowes, D. and M. Squair. 1999. "Managing the Emergent Properties of a Design." Paper presented at the Ninth International Symposium of INCOSE, Brighton, UK, 6-11 June.
- Dalton, S. 2013. *The 21st Century Character of Air Power.* 1–11 Sir Sydney Camm Memorial Lecture. London, UK: Royal Aeronautical Society, 10 June.
- De Meyer, A., C. H. Lock, and M. T. Pich. 2002. "Managing Project Uncertainty: From Variation to Chaos." *MIT Sloan Management Review* 43 (2): 60-67.
- Hopkin, P. 2002. *Holistic Risk Management in Practice* 1st ed., London, UK: Witherby & Co. Ltd. 3.
- INCOSE, 2009. Z1 Guide, Revision 3, 2. Available at: http://incoseonline.org.uk/Program_Files/Publications/zGuides_1.aspx?-CatID=Publications.
- Maier, M. W. 1998. "Architecting Principles for Systems-of-Systems." *Systems Engineering*, 1 (4): 267-284.
- Norman, D. A. 1999. "Affordance, Conventions and Design." *Association of Computing Machinery: Interactions*. May-June, 38–42. Available at: http://www.jnd.org/dn.mss/affordance_conventi.html.
- Ring, J., and T. Tenorio. 2012. *INSIGHT* 12 (2): 9-12.
- Rumsfeld, D. 2002. US Department of Defense briefing to NATO, Brussels. NATO Press Conference, p. 8. Available at: http://www.nato.int/docu/speech/2002/s020606g.htm.
- Sillitto, H. 2011. "Integrating systems thinking, systems science and systems engineering – understanding the difference and exploiting the synergies." Presentation at INCOSE ASEC 2011 (p. 20), INCOSE.
- Smith, R. 2005. *The Utility of Force* 1st ed., London, UK: Penguin. 3.
- UK MoD. 2008. UK MoD AOF. Acquisition Operating Framework. Retrieved from https://www.gov.uk/acquisition-operating-framework.

ABOUT THE AUTHORS

[Editor: Author biographies were current when the paper was initially published in 2016.]

Steve Hinsley is a mature PhD research student under Professor Mike Henshaw and Professor Carys Siemieniuch in the Engineering Systems-of-Systems Group of the School of Electronic, Electrical and Systems Engineering at Loughborough University. The UK Engineering and Physical Sciences Research Council via Loughborough University Graduate School and Thales fund the research. Previously Steve was technical lead of the systems engineering capability for BAE Systems Advanced Technology Centre, and chief engineer at the Systems Engineering Innovation Centre. He is a chartered engineer, a member of the IET, and INCOSE. He holds a BSc (Hons) in electrical engineering from De Montfort University Leicester and a MSc in advanced systems engineering from Loughborough University.

> continued on page 38

Bringing Operational Perspectives into the Analysis of Engineered Resilient Systems

Valerie B. Sitterle, Valerie.Sitterle@gtri.gatech.edu; **Erika L. Brimhall,** Erika.Brimhall@gtri.gatech.edu; **Dane F. Freeman,** Dane.Freeman@gtri.gatech.edu; **Santiago Balestrini-Robinson,** Santiago.Balestrini@gtri.gatech.edu; **Tommer R. Ender,** Tommer.Ender@gtri.gatech.edu; and **Simon R. Goerger,** Simon.R.Goerger@erdc.dren.mil

Presented at the 26th Annual INCOSE International Symposium (2016) Edinburgh, UK-Scotland, 18-21.

ABSTRACT
Engineered Resilient Systems (ERS) is a Department of Defense (DoD) program focusing on the effective and efficient design and development of complex engineered systems across their lifecycle. An important area of focus is the evaluation of early-stage design alternatives in terms of their modeled operational performance and characteristics. The work in this paper ties together differentiated operational needs with requirements specification and maturation of previous analytical constructs toward a more operationally relevant viewpoint. We expand on the concept of Broad Utility as a high-level aggregated measure of robustness of fielded system capabilities with respect to operational requirements. The relation to requirements is more explicit, and systems are failing to achieve threshold requirements are penalized. The impact of this approach and how it offers a foundation from which to more fully explore sensitivity to Pre-Milestone A requirements are discussed.

ERS FOUNDATIONS AND RELATION TO TRADESPACE ANALYSIS

A large body of work currently exists concerning developing decision support methods and a tradespace toolset framework architecture in support of the Department of Defense's (DoD) Science & Technology priority for Engineered Resilient Systems (ERS). ERS is a U.S. DoD program focusing on effective, efficient design and development of complex engineered systems across their lifecycles and actively being implemented across a wide variety of engineering concepts, techniques, and design tools. Through ERS, the DoD seeks a transformation in defense acquisition processes via systems engineering throughout a system's lifecycle that will enable the DoD to better respond to an environment characterized by rapidly changing threats, tactics, missions, and technologies. ERS calls for adaptable designs with diverse systems models that can easily be modified and reused, the ability to iterate designs quickly, and a clear linkage to mission needs. An important area of focus is the evaluation of early-stage design alternatives regarding their modeled operational performance and characteristics. This includes research and development of methodologies to conduct Analysis of Alternatives (AoA) relevant to evaluating different dimensions of resiliency for these systems.

Towards this end, tradespace exploration and analysis enable decision makers to discover and understand relationships across capabilities, gaps, and potential compromises that facilitate the achievement of system objectives. These objectives undergo expression through requirements or other metrics. To be effective, decision makers must have deep knowledge of the component elements of a system, which includes how these elements interact internally to the system and externally with the operational environment (Spero et al. 2014). This requires development and maturation of executable and scalable analytical constructs and processes that must be able to be implemented within the context of a larger workflow to guide tradespace exploration and evaluate ERS resiliency concepts. Dr. Jeffery Holland defined the characteristics of a resilient system from an ERS perspective as i) trusted and effective in a wide range of contexts, ii) easily adapted to many others through reconfiguration and replacement, and iii) having a predictable degradation of function (Holland 2013). Goerger, Madni, & Eslinger (2014) matured this view to include the concept of "Broad Utility," a mission-focused perspective defined as the "*ability to perform effectively in a wide range of operations across multiple potential alternative futures, despite experiencing disruptions*." Conceptually, Broad Utility relates to concepts of robustness via

performance across a wide range of operations and possible mission contexts.

Using the guiding principles of ERS as a foundation, this work adheres to a specific context of evaluation. Namely, we do not seek to address the myriad of dimensions that constitute resilience, but rather seek to mature the concept of Broad Utility concerning its implementation as a computationally executable analytical construct to embody a more operationally relevant perspective. This effort builds on previous work (Sitterle, Curry, and Ender 2014, Sitterle et al. 2015) while still seeking to promote scalable, implementable methods that are transparent, intuitive, rational, and quantifiably traceable. We discuss how distinct operational scenarios may be defined to support a more operationally relevant context of analytical exploration in an executable environment, how these scenarios relate to requirements specified by the stakeholders, and how together these concepts can lead to an expression of Broad Utility that goes beyond the commonly used framework of additive multi-attribute valuation. Broad Utility as it is matured here is simply a starting point for systems engineering resiliency of engineered systems. It is intended to be composable with other analytical constructs and not as the single basis for evaluation. We aim to promote a better synergy between the design analysis and requirements generation processes, which are often not well integrated. And, in doing so, we offer more insight to requirements maturation for systems under development and how this relates to operational performance expectations.

THE LARGER SYSTEMS ENGINEERING PROCESS AND GENERATION OF A TRADESPACE

In Acquisitions, a Concept of Operations (CONOPS) is used to "examine current and new and proposed capabilities… and describes how a system will be used from the viewpoints of its stakeholders" (AcqNotes 2015). A CONOPS commonly expresses as a verbal or graphical statement of a commander's assumptions or intent concerning an operation or series thereof. It provides a bridge between vaguely expressed capabilities needed from a system and specific technical requirements needed to evaluate the system and enable it to be successful. These requirements "govern what, how well, and under what conditions a product will achieve a given purpose" (ANSI/EIA 2003), inherently describing capabilities necessary.

In specified mission contexts or for future operations. Capabilities – defined as the ability to execute specific courses of action.

Requirements may be classified according to type (Defense Acquisitions University 2011, Pflanz et al. 2012). Key Performance Parameters (KPPs) are key system capabilities that must be met for a system to meet its operational goals. Key System Attributes (KSAs) are capabilities considered crucial in support of achieving a balanced solution to a KPP or other key performance attribute deemed necessary by the stakeholder. Other Performance Parameters (OPPs, also called Tier III) are desirable but not critical toward providing required capabilities for mission success. KPPs, KSAs, and OPPs/Tier IIIs are considered "must," "should," and "could" haves respectively. It follows that KSAs are below KPPs in priority, while OPPs/Tier IIIs are below KSAs. Requirements such as KPPs and KSAs specifications are usually as quantitative metrics containing those attributes or characteristics of a system considered critical or essential to the development of an effective defense capability. They are expressed as having a threshold and objective levels, corresponding to the minimum acceptable and desired values (Defense Acquisitions University 2011). This structure is directly amenable to quantitative evaluation and offers a basis for external reference for valuation of individual requirements as explained in the "Evaluating Broad Utility" section.

We illustrate the larger systems engineering process that relates the expression of a CONOPS with the subsequent derivation of requirements and leads to the identification of early-stage designs and their evaluation by the idealized workflow shown in Figure 1. The figure highlights two distinct branches leading to a system description model: Measures of Effectiveness (MOEs)-to-Architecture Identification (vertically on the left), and Measures of Performance (MOPs)-to-System Design (horizontally at the bottom). In the former, evaluation measures derive from the operational requirements and stakeholder expectations. Next identification of KPPs and potential system design architectures occurs. The latter branch leads to the specification of system design variables that will aid analysis of the various evaluation measures. Both branches are iterative processes not always performed collaboratively, but the generation of a tradespace requires a synthesis of these concepts.

AoAs happen in large part through exploration of a tradespace. A tradespace is defined as the complete enumeration of the system alternative design variables together with the set of program and system parameters, attributes, and characteristics required to satisfy performance attributes associated with each system alternative. It is the complete solution space. Once a we create a tradespace we can then explore it. a significant amount of work goes into creating the problem, potential design architectures, modeling and simulation (M&S) components that will map the design variables to output measures on which tradeoffs will be assessed, and how the overall problem specification maps to stated stakeholder requirements and distinct operational scenarios.

Figure 2 illustrates this secondary process whereby a tradespace may be generated in a computational environment, allowing designs to be evaluated based on modeled fielded performance and essentially reversing the flow of Figure 1. Once system design variables and relevant M&S components are identified and/or developed, then we can create a tradespace for

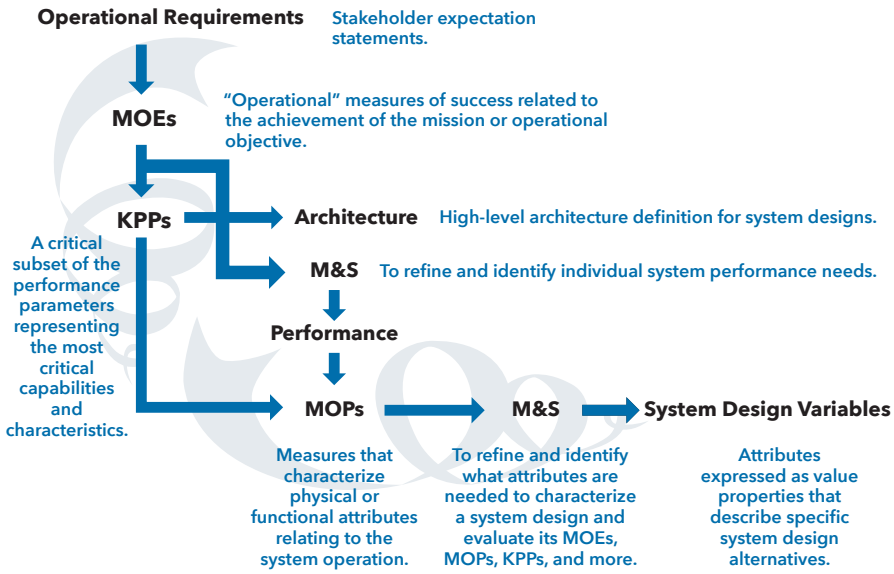

Figure 1. High-Level systems engineering process identifying evaluation measures and design variablwes

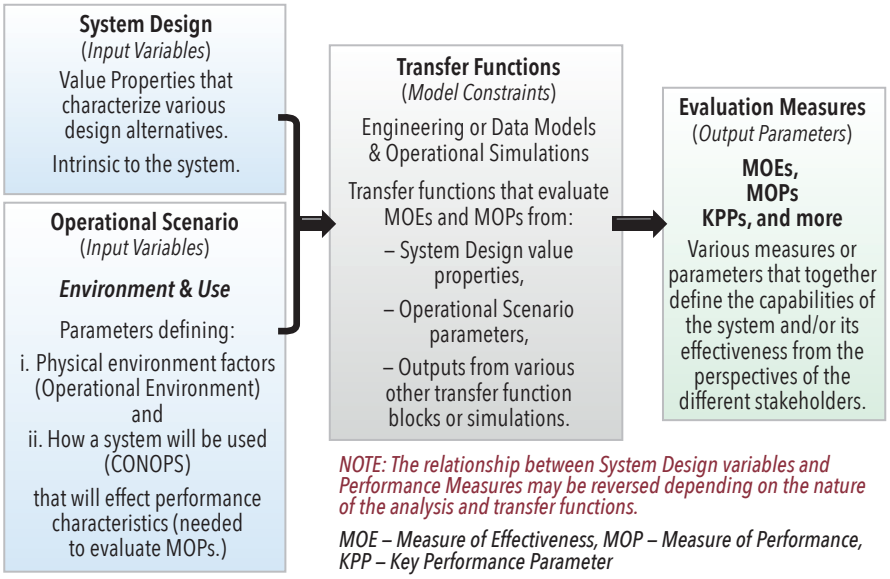

Figure 2. Development of a tradespace for analysis

evaluation of the MOPs, MOEs, and KPPs across the design alternatives. A tradespace analysis should be capable of capturing multiple viewpoints and analytical goals. Different stakeholders may value different sets of evaluation measures (MOEs, MOPs, KPPs, etc.) or may value the same measures differently. Analyses may take many different forms: i) define outcome measures and explore system designs that meet them, ii) converge on what these outcome measures should be through an iterative process, or iii) identify system design properties from an exploration of evaluation measures. While what constitutes a MOE, MOP, and KPP beyond the descriptions provided in Figure 1 is outside the scope of this paper, more detailed discussion is provided by the US Department of Defense (2013).

DEFINING UNIQUE OPERATIONAL SCENARIOS

When generating a tradespace, the process must enable the definition and subsequent analysis of the operational environment and CONOPS parameters necessary to evaluate the MOEs, MOPs, and technical requirements expressed as KPPs, KSAs, and others. To define an operational scenario as depicted in Figure 2, one must specify physical environment factors (ambient temperature, road conditions, humidity, winds) as well as parameters directly relating to system use (crew carrying capacity, maximum speed up a specified grade,) that will effect and therefore be required to evaluate system performance characteristics. The physical environment factors characterize the operational environment, and the use parameters correspond to CONOPS expectations. Delineating operational scenarios into these two parameter classes serves as intuitive scaffolding for modelers to specify attributes necessary to generate an operationally specific tradespace. Parts may be reused and expanded upon for creation of new operational scenario blocks. Like most aspects of model development, the attributes or measures the tradespace generation is intended to produce will drive what parameters to define in an operational scenario. For example, many defense systems are designed to operate in frigid, icy conditions as well as hot and humid conditions. The environment model blocks should capture precisely those parameters relevant to how the system under study will perform (vehicle acceleration or internal climate control, and air filtration). These model blocks may be dynamic simulations specific to certain scenarios.

Any number of operational scenarios may be defined in this way depending on the scope and needs of the analysis. Scenarios may be interested in the same output measures (cruise range) but impose different objective or threshold levels or require operation in vastly different environmental profiles that, in turn, alter the measured performance. Similarly, operational scenarios may be interested in entirely different output measures regardless of operational environment, and impose requirements not present in other scenarios. The level of detail required by the present stage of analysis, effort required to model and define them, and time and cost required to do so governs the complexity and extent of the definition of these parameters.

Implementation. ERS TRADESPACE is an integrated toolset and software architecture development undertaken collaboratively by Georgia Tech Research Institute (GTRI) and the Army Engineer Research and Development Center (ERDC). ERS TRADESPACE includes various components that serve as user interfaces to create and execute analyses or analytical blocks, a combination of hosting engines, and models and linking engines that coordinate, structure, and integrate analyses from different software components. ERS TRADESPACE thereby supports an end-to-end capability linking requirements and CONOPS specification to specification of design alternatives, generation of a trade space, and subsequent tradespace exploration and analysis (Balestrini et al. 2015a, Balestrini et al. 2015b). While the underpinnings of ERS TRADESPACE are beyond the scope of the present paper, it is the platform through which we have been investigating various degrees of modularity in implementation and how this may impact reusability, the backend data model, and the interface semantics necessary to facilitate subsequent analyses. The end goal is to use the description of the scenario and other building blocks to build an executable environment necessary to compute the required system metrics. Analyses that differentiate and integrate these scenario objects build from the backend data model specification. The methods described in the following sections, however, are toolset agnostic. They are described and written to be transparent and implementable in any executable environment.

EVALUATING BROAD UTILITY IN TERMS OF STAKEHOLDER VALUATION

Needs Contexts and Valuation from a Requirements Reference

A significant concern during early phases of acquisition, or during the Pre-Milestone An analysis of the DoD Acquisition process, is the resiliency of a system design across simultaneously competing or sequentially changing requirements on its performance attributes. A Needs Context, defined in previous work (Sitterle, Curry, & Ender 2014; Sitterle et al. 2015), is a scalable, applied methodology to capture certain resiliency dimensions related to how well a system performs its functions in the face of requirements perturbations. It builds from robustness as defined by Ryan, Jacques, & Colombi (2013) and the concept of Broad Utility advocated by Goerger et al. (2014), creating a requirements-based evaluation of the non-cost value of system design alternatives. Needs Contexts may be more completely described as characterizing "Robustness of Fielded System Capabilities and Capacity with respect to Operational Requirements." Contexts are defined based on flexible subsets of performance attributes relevant to the stakeholder(s) and ranking of those attributes within each. Succinctly, an individual Needs Context specifies a

subset of evaluation measures deemed critical to a stakeholder as the basis for analysis. The motivation is that choices must be made based on what is valued most by stakeholders, recognizing that some stakeholders may have a greater influence. Together, multiple Needs Contexts can be constructed to represent different viewpoints and can represent different or directly competing objectives for a system's performance:

- Different stakeholders, each with different or competing priorities in parallel
- Changes in requirements over time (future performance requirements differ in series)
- Different mission profiles with performance objectives, whether in parallel or series.

Requirements Basis for Value Functions. Value of a given system attribute is scaled against objective and threshold requirement levels using a KPP concept to promote comparability across analyses. The attribute value functions limit all possible valuations to the range of 0 to 1 by assigning any levels below the threshold or above the objective equal to 0 or 1 respectively. A tradespace may or may not cover the entire range. Value of a system design alternative is then assessed using an additive multi-attribute value (MAV) model, synergistic with the concept of evaluating Broad Utility via the robustness descriptor presented above. Since each Needs Context may be defined using different attributes, and different valuations and preference weightings, Needs Contexts can produce a different value for each system design alternative *k* (SD_k) within each *Needs Context m*,:

$$| \; U_k = w_i * v_i(Y_{ik}) + w_j * v_j(Y_{jk}) + \ldots w_n * v_n(Y_{nk}) = {}^{n}\Sigma_{i=1} \; w_i * v_i(Y_{ik}) \; |_{\text{Needs Context } m}$$

U_k denotes the overall value of system design alternative k (SD_k) for a given Needs Context, Y_{ik} represents a system attribute *i* for SD_k, each v_i is a value function expressing the relative value of the given system attribute level to a stakeholder, and wi are weights derived from preference rankings or other means. In keeping with traditional utility theory, overall system value is limited to the range of 0 to 1. Value functions are typically linear or exponential expressions but may be any monotonic function. Cost is a function of system design alternative characteristics, though it depends on other influences and variables as well. Utility and cost are therefore expressed as related dimensions, linked by an underlying SD_k.

Limitations of Additive Multi-Attribute Value Models. In the previously cited work, the Needs Context served as the basis from which an analyst could construct an overall valuation for each system design alternative from the perspective of the individual stakeholders. Though using a unique, requirements-based valuation construct, the overall valuation still relied on the commonly used additive multi-attribute value (additive MAV) model, also called the sum additive weight (SAW) method. This approach is scalable and intuitive, yet it does not adequately represent a more operationally focused perspective. For example, consider a set of five attributes each with equal weights (all $w_i = 0.2$). A system design that exhibits valuations of each attribute to a level of 0.8 (all $v_i = 0.8$) will produce the same measure of overall value, $U_k = 0.8$, as a design alternative where 4 of the 5 attributes meet the objective but one attribute fails to meet threshold ($v_i = [1, 1, 1, 1, 0]$). Similarly, as the number of attributes in the measure increases, the impact of attributes failing to meet threshold decreases. This same effect occurs with the traditional attribute value scaling to the given tradespace. In an operational environment, a defense system failing to achieve a key requirement threshold is not equally acceptable.

MATURING TO AN OPERATIONAL NEEDS CONTEXT

The challenge is to mature the overall valuation measure from a traditional additive MAV to a construct more representative of the operational viewpoint. The Needs Context is already well suited to represent disparate operational scenarios that may exist, and by definition, it captures those measures of performance deemed critical from a given operational perspective. However, failure to meet one or more critical requirement thresholds should be either readily apparent or carry some penalty that prevents the alternative from possessing a valuation on the same level as an alternative meeting all thresholds. Since there are analyses that may need data points representing designs that do well in many measures but fail in one or two to persist, we will not force the valuation for these designs to zero.

Penalty Function. Considering the operational perspective, we sought to modify the additive MAV model to include an "operational penalty" for alternatives with any one or more individual attribute value functions evaluating to zero. The additive MAV model value is the maximum valuation, an alternative could achieve (as it contains no penalty), while the overall valuation even with every attribute failing to meet threshold preserved the understood lower limit of zero. An exponential function of the value was used to generate a penalty effectively equal to the weight of any attribute with a value of zero and no penalty otherwise. A direct comparison of the traditional additive MAV model and the model with an operational penalty for a given system design alternative are as follows:

$$U_{+MAV,k} = {}^{n}\Sigma_{i=1} \; w_i * v_i(Y_{ik}) U_{OpPenalty,k} = U_{+MAV,k} * [\, 1 - {}^{n}\Sigma_{i=1} \; w_i * exp(-\theta * v_i(Y_{ik}))\,]$$

where n represents the number of attributes included in the value model, wi are the weights of each attribute, and vi are the values of the attributes for the given design alternative as obtained from the individual attribute value functions as before. The exponential penalty function produces $U_{OpPenalty} = U_{+MAV}$ when all requirements meet or exceed threshold levels, and a penalty effectively equal to the weight of the individual attribute if its level is below threshold such that $v_i = 0$. θ is chosen to be sufficiently large as to ensure this outcome for even the smallest feasible value. $\theta = 1000$, for example, reduces the exponential term to 4.54E–5 even if an individual valuation $v_i = 0.01$.

Surrogate Weighting. Methods used in ERS TRADESPACE and all methods described here are agnostic to how ranks, or even weights, are derived. Weights in an additive MAV model may originate from any number of methods including subject matter expert (SME) opinions, historical priorities, guided stakeholder discussions, and pairwise comparisons. Another approach that may be particularly useful for analysis of early-stage designs is to use surrogate weights based on the attribute rankings. If ranks are inconsequential or unknown across our subset of critical performance measures, equal weights are an appropriate starting point. If the ranks are known, and the preference order holds, different weighting methods may better reflect how the ranks are valued, such as linearly, exponentially, and more (Roszkowska 2013). Among these, rank order centroid (ROC) surrogate weights are one of the most robust options when there is some uncertainty in weights, but the rank preferences are clear. ROC weights are computed from the vertices of the simplex where $w_1 \geq w_2 \geq \ldots \geq w_n \geq 0$. Weights are t the coordinates of the centroid for the simplex, found by averaging the coordinates of the defining vertices. This approach assumes that the ranks specify the information set on the weights and that no point in the simplex is, therefore, more likely than another (weight density uniformly distributed over the simplex). Consequently, ROC weights are the

expected value weights for the respective probability density functions over the feasible weight space (Barron and Barrett, 1996) (This is readily demonstrated using a Monte Carlo simulation).

Despite their advantages, ROC weights alone do not solve issues of range sensitivity in decision analysis. Weights are usually adjusted from one tradespace analysis to the next because MAV functions traditionally normalize to the range of the local decision context, the current tradespace. Normalizing this way can produce very different decision outcomes when the tradespace range changes if the weights do not change. This is known as the "range dependence of weights" or "range sensitivity principle." The swing weight concept was developed specifically to preserve consistent decision outcomes in the face of changing tradespace ranges (Johnson et al. 2013). Swing weights and other weight-adjusting approaches focus only on altering the weights in the additive MAV model but work very well across different tradespace instantiations using range-dependent normalization.

However, intuitive perceptions of attribute importance are often independent of the range of the outcomes. We focused instead on how to adjust the value functions that effectively grade the individual attributes within the MAV model. Developing value functions that normalize to a basis external to the local decision context offers two primary advantages. Firstly, it preserves consistency in decision outcomes just as do the previously developed swing weight methods. Secondly, externally valuing attributes promotes direct comparability from one tradespace to the next while weight-adjusting methods with tradespace-dependent value functions do not. When using a value function basis external to the tradespace, weight-adjusting methods are not necessary. Our approach exploited the KPP/KSA requirements structure to form value functions not dependent on the current tradespace range also offers a clear analytical link to requirements. We can compare the impact of competing or changing requirements readily through construction of new Operational Needs Contexts. ROC weights, as expected values over the feasible weight space, are now a solid starting point for analyses when more rigorously obtained weight data are not available. When using an additive MAV model and the individual valuations are also expected values of the given attributes, ROC weights produce the expected MAV Broad Utility given the preference order established by the weights.

Requirements Differentiation. As discussed earlier, requirements expression may occur according to hierarchical type. While the Needs Context and Operational Needs Context valuation models presented in the previous section make no distinction between types of requirements, the models are easily amenable to do so. Weights in any MAV are scaling constants and, as such, may be "re-scaled" if necessary to differentiate between levels of priority or value. Returning to the concept of "must," "should," and "could" have requirements corresponding to KPP, KSA, and OPP/Tier III requirements types respectively, the equations used to assess Broad Utility may alter via a "requirement weight," β_i. For example, all "must have" requirements could be assigned $\beta_i = 1$, which reduces to the previous version of these value models. Measures of performance classified as "should have" and "could have" requirement types, might be assigned values of 0.8 and 0.6 for β_i respectively. The β_i values for these lower level requirement types are simply examples by but not directly based on DoD guidelines. The following valuation shows the function form when the requirement weight is only applied to the penalty term if a design fails to meet a requirement threshold:

$$U_{OpPenalty-\beta,k} = U_{+MAV,k} * [\,1 - {}^{n}\Sigma_{i=1}\, \beta_i * w_i * exp(-\theta * v_i(Y_{ik}))\,]$$

If we apply the β_i term in the traditional $U_{+MAV,k}$ model as well (yielding $U_{+MAV-\beta,k} = {}^{n}\Sigma_{i=1}\, \beta_i * w_i * v_i(Y_{ik})$ and subsequently a $U_{OpPenalty-\beta,k}$ model also using this form), analyses that choose to focus on OPP/Tier III measures will not yield valuations of Broad Utility on par with those based only on critical, KPP type measures even when meeting all requirements. That can be logical in the sense designs focused on meeting OPP/Tier III requirements should not be valued as highly as those meeting KPPs. The key to keeping such analyses meaningful, however, is consistency in application and documenting why that application is warranted.

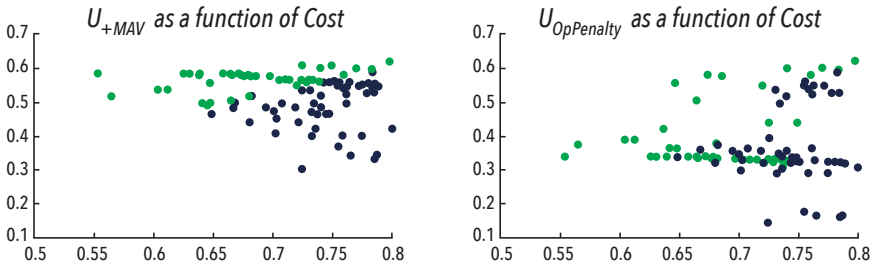

Figure 3. Example comparison of $U_{+MAV,k}$ and $U_{OpPenalty,k}$

Example. The Operational Needs Context matures the prior construct to include a penalty for failure of any performance attribute to meet its threshold value, producing a more operational view for alternative valuation. As an openly sharable example, a traditional additive MAV valuation model and the operational penalty valuation model were applied to the Iris dataset as shown in Figure 3. The Iris data set is a multivariate data set introduced by Fisher (1936) and is also included in the Seaborn Python visualization library Waskom (2015). We set threshold and objective values within the data range for each attribute, specified a rank order of {*petal_length, sepal_length, petal_width, sepal_width*}, and prioritized higher values of *petal_length and sepal_width*. We evaluated the cost as a model function of the attributes. Designs with any attribute levels below the threshold can be part of the broader Pareto set when using a $U_{+MAV,k}$ value model but not classified as such using a $U_{OpPenalty,k}$ value model. The "best set" taken from the traditional additive MAV approach in (a) underwent identification in a manner analogous to the "fuzzy Pareto set" described by Smaling & de Weck (2004). Instead of taking points off the Pareto frontier as a function of some K percent of the total range of the utility/cost data, however, we took successive Pareto layers. Specifically, we identified the Pareto frontier for the whole data set, then the Pareto frontier for the remaining data, and so on for a defined number of layers. This approach removed sensitivity to the range of data and helped preserve a transparent linkage between identifying a Pareto set and a "best" decision.

Relevance to Sensitivity and Uncertainty Analysis. Approaches investigating uncertainty on attribute weights in additive value models were well defined by Charnetski & Soland (1976, 1978) and expanded upon by Lahdelma, Hokkanen, & Salminen (1998), Lahdelma, Miettinen, & Salminen (2003), and Tervonen & Lahdelma (2007). They apply well to the $U_{OpPenalty}$ model. But, there are interesting ramifications when investigating uncertainty in the value functions comprising the model and the impact of this synthesis with the $U_{OpPenalty}$ construct. Firstly, if normalizing attribute values to the tradespace range, uncertainty must be characterized or propagated before normalization. Otherwise, valuations at range extremes can produce values below 0 or above 1. In contrast, distributions associated with uncertainty can be incor-

porated at any stage in the analysis when using an external value reference; attribute values will always be bound between 0 and 1. There are interesting dynamics for uncertain attributes with levels close to their objective and threshold. As shown in Figure 4, which investigated a uniform distribution of uncertainty on the iris attributes, data (shown as normalized histograms) skewing can occur near the objective (*sepal_width*) and discontinuous near the threshold (*petal_length*). Pulling value function distributions with these characteristics into a higher-level model such as $U_{OpPenalty}$ necessitates a sampling strategy since they are not readily mathematically convoluted with other distributions. Figure 5 extends these results to the evaluation of Broad Utility as characterized by the U_{+MAV} and $U_{OpPenalty}$ constructs. Figures 5 (a) and (b) show uncertainty only on the weights, evaluated through a Monte Carlo simulation as described by Lahdelma, Miettinen, & Salminen (2003). The distributions are understandably narrower when the rank order enforcement occurs as shown in (b). Figures 5 (c) and (d) then show the impact of uncertainty on the weights and iris "alternative" valuations when there is enforcement of preference ranks. Both distributions are broader than the comparable case for weights-only uncertainty in (b). The $U_{OpPenalty}$ case in (d) magnifies the effect from (c), resulting in a heavier distribution toward the lower values due to the uncertainty of an attribute near to its threshold (*petal_length*). This underscores the importance of rigorously investigating design alternatives with uncertain attribute levels near to thresholds and objectives, especially if classified as being in the "best set" of Pareto designs.

DISCUSSION

If we revisit the DoD acquisition process, the "system need is established, and high-level system requirements are defined" during the Materiel Solution Analysis (MSA) phase that culminates with the Milestone A decision (Baldwin et al. 2012). Notional system architectures are "often created to assist with the requirements analysis and definition for the preferred system concept," and reviews are conducted to ensure that the resulting requirements set "agrees with the customer's needs and expectations" (Baldwin et al. 2012). Our methodology using a requirements-based concept of value scaling for design alternative attributes supports these DoD goals, offering a direct analytical linkage between tradespace exploration and requirements maturation. In our implementation via the ERS above TRADESPACE toolset, we do not force any given method on systems engineers. Rather, we enable the flexibility to tailor analyses to needs. Engineers are therefore able to use simply an additive MAV model or apply the operational penalty as described here when constructing Needs Context perspectives and analyses. With or without an operational penalty or weighting according to requirement prioritization, the method offers direct means for an analyst to investigate the impact of various requirements on Broad Utility to stakeholders.

The Operational Needs Context is a modular construct designed to preserve scalability. Using an exponential penalty function, for example, enables ease in execution compared to attribute-by-attribute "if-then" statements for each alternative. The method is based on selecting key evaluation measures, focused on what matters most to stakeholders. Its

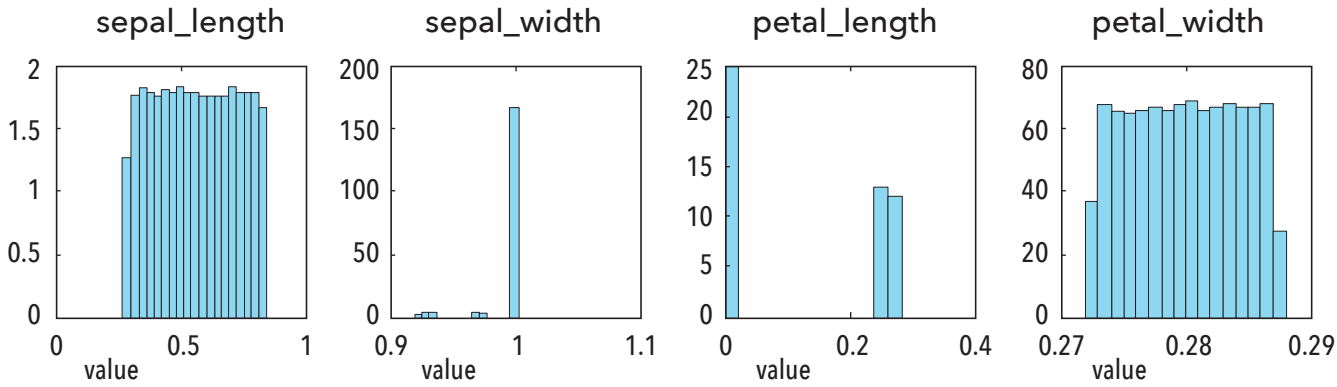

Figure 4. Illustration of impact of uncertainty on value function results

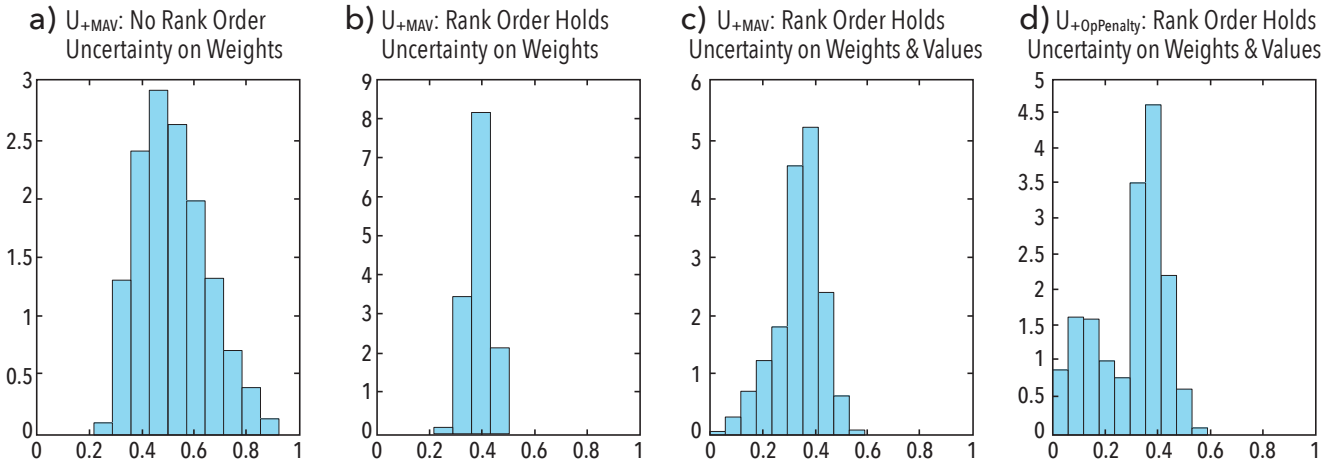

Figure 5. Illustration of relation between broad utility and uncertainty

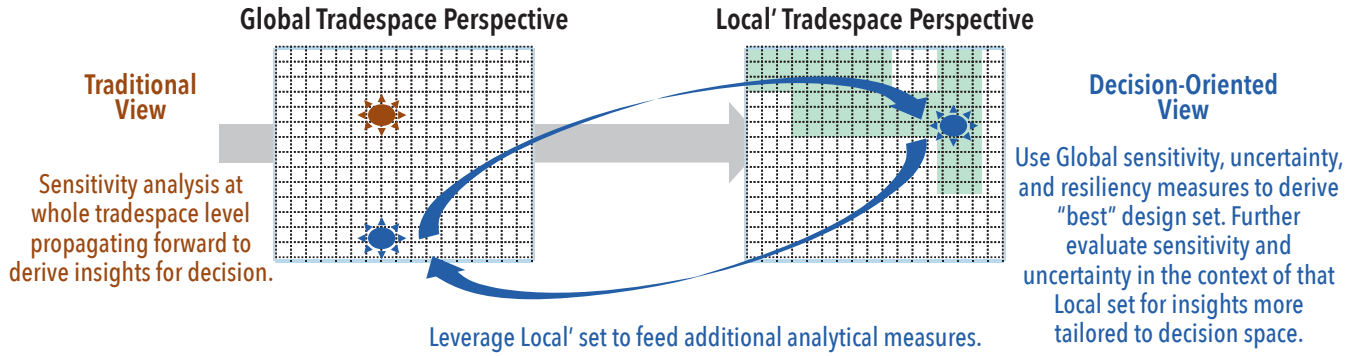

Figure 6. Comparison of a traditional data view with a more decision-oriented view

design supports the analysis of operational scenarios that differ according to the operational environment, the system use (to do what – a CONOPS), and even competing stakeholder priorities across the desired capabilities. Prioritizing agility and acceleration for a vehicle design, for example, may not produce alternatives well suited to driving through mud or with strong underbody blast protection. Some designs may be robust across multiple stakeholder operational needs; some may not. The Operational Needs Contexts can help highlight where compromise may or may not be possible while providing a traceable, quantitative basis for reducing a set of options for further analysis.

This work has also shown the importance of uncertainty, its treatment, and interpretation. Uncertainty analysis concerning identifying a "best set" of alternatives should focus on what aspects of uncertainty change our decision about which design alternatives to include in that set. Figures 4 and 5 illustrate how valuation uncertainty can result in highly skewed or discontinuous distributions, and uncertainty on weights and values can produce a bimodal aggregate distribution when an operational penalty is applied. Design alternatives in a "best set" with attribute levels near the objective and especially threshold values must be carefully evaluated. Simply taking a mean, standard deviation, or quartile representation may not represent the uncertainty well. As with any mathematical method, selection and effectiveness of synthesis with other methods are problem- and process-dependent. And, addressing composability highlights the need to understand global versus local tradespace perspectives.

COMPOSABILITY AND GLOBAL-TO-LOCAL ANALYSIS PERSPECTIVES

Within an ERS perspective, evaluating resiliency of systems consists of making decisions (trades) at a high level across broad measures of performance that themselves derived through a process of making trades across lower level variables. High-level resiliency dimensions may include Broad Utility (Goerger et al. 2014), reliability, manufacturability, flexibility to engineering change (Sitterle et al. 2015), system development cost, overall lifecycle cost, and so on. There will likely not be any one answer or view across the numerous resiliency dimensions that exist and yet to undergo development. For early-stage design, Pre-Milestone A in the DoD materiel acquisitions process, we typically do not have the *a priori* insights to specify the exact nature of the relationships across various dimensions of resiliency. More often, we evaluate multiple designs and architectures, bubbling up to the design space regions that are feasible and desirable. Through continued analysis of these desirable regions, we begin to analyze and gain insights concerning how the resiliency dimensions for hypothetically realized design concepts interrelate.

In traditional decision analysis of tradespace data, the context of evaluation basis is on the whole of a specific instantiation of a tradespace. An ERS perspective requires the maturation of a more holistic approach toward integrating the whole tradespace analytical view with analyses targeted toward the "best" set of designs. For small sets of system design alternatives, it may be quite feasible to evaluate various measures of performance feeding resiliency measures (in turn a higher-level MOP or MOE) for all design alternatives. As the number of design alternatives increases, say to 10,000 or 1M or more, the global versus local treatment becomes more complicated. Some sensitivity analyses may need to be global, performed across the entire tradespace to reduce a set of parameters or attributes under evaluation. In other analytical treatments, a global approach can compromise insights that may be specific to the "better" design alternatives with bias from the "poor" alternatives as well as inhibit the ability to scale both computationally and visually. (This is especially true for component-based tradespaces.) Focusing some aspects of the analysis on a more tailored, decision-oriented space, a "*Local*" data set, will enable deeper insights regarding sensitivity, uncertainty, and correlations across the "best" set of designs. Toward this end, the Operational Needs Context can alter which alternatives are in that set via the penalty function. The Operational Needs Context may be used after a different analytical treatment narrows the alternatives or attributes thereof, or it may be used as a precursor, reducing the "input tradespace" to the next analytical treatment to a more rational set of designs. This may be done in any number of ways: taking a top set of designs including and off of the Pareto front as described earlier, taking a specified top percentage of alternatives with the highest values (irrespective of cost), and more. Which way to take a "best set" to propagate to the next analytical processes depends entirely on the needs of the next treatment and the overall process. These concepts are illustrated in Figure 6.

By treating the Operational Needs Context as an analytical building block that helps identify promising designs for additional types of resiliency analyses, we must understand that workflow matters when seeking to compare results across different analytical efforts. Previous work began to investigate how different constructs might be synthesized when one is serving to reduce a set of design options (Sitterle et al. 2015). To help promote a more local tradespace context, one analysis used a weighted local covariance, a covariance calculation taken from machine learning applications that prioritize the contributions of nearest neighbors with appropriate selection of the weighting kernel. Even so, results from a measure of flexibility change when the analysis is applied to the global (entire) tradespace versus the local, "best set" tradespace filtered by Broad Utility. This is intuitive but serves to underscore the importance of understanding where and when to take local perspectives. In some tradespaces and types of analyses,

including all points may hinder effective evaluation when some points characterize completely different systems.

SUPPORTING DECISION ANALYSIS FOR ERS

The focus for early-stage design is not identifying an optimal solution, but rather narrowing down the potential solution space for more rigorous data collection, generation, and evaluation. We are interested in how our broader Pareto set changes across Operational Needs Contexts and how that set changes if there is uncertainty associated with weights assigned to the attributes prioritized in each context and valuation of those attributes. With explicitly quantifiable and traceable links to requirements, the analytical constructs described here offer a means for direct exploration of the sensitivity to threshold and objective values. In keeping with the goals of ERS, this capability can aid early-stage design requirements maturation and offer an approach that can be synthesized with other analytical approaches to build a complete characterization of resiliency.

Even so, the challenge for ERS extends much further than analytical methods. Active development continues on ERS TRADESPACE, the aforementioned integrated toolset by GTRI in collaboration with ERDC, especially on methods and tools to support complex analyses across both local and non-local model and simulation components. A limitation of model-based tradespace generation and analysis is the level of time, cost, expertise, and effort required to develop and implement relevant M&S components. There is a limit to how many threats and how many simulations may be defined and evaluated for each system. Our guiding philosophy is therefore to support an integrated and yet highly flexible design and analysis process focused on identifying the set of alternatives most likely to meet requirements based on the information (data, design architectures, models) available.

Throughout this effort, we adhere to the philosophy that the most important goal is insight, not numerical treatment or inference. In decision analyses, qualitative concepts and judgments are often required to be translated into quantitative measures to enable scalable, consistent, and traceable analyses. Even so, workflow matters. Consistency across treatments and processes by which they apply are critical to meaningful (actionable) insights. Methods, processes, and tools (MPTs) should be engineered together to promote transparent, intuitive, rational, and quantifiably traceable foundations for resiliency analyses. Analytical processes for characterizing resiliency still, however, depend on the specific stage of design or place in the acquisition lifecycle. Mature frameworks for resiliency evaluation appropriate to various types and stages of the acquisition lifecycle will emerge from the more extensive application and lessons learned. We hope that by sharing the perspectives described here, we will help promote a collaborative, communal maturation of resiliency analyses for ERS across researchers and DoD customers alike. ■

ACKNOWLEDGEMENTS

Portions of this material come from work supported, in whole or in part, by the United States DoD through the Systems Engineering Research Center (SERC) under Contract HQ0034-13-D-0004. SERC is a federally funded University Affiliated Research Center managed by Stevens Institute of Technology. The views and conclusions are those of the individual authors and participants, and should not be interpreted as necessarily representing official policies, either expressed or implied, of the DoD, any specific US Government agency, or the US Government in general.

REFERENCES

- AcqNotes. 2015."JCIDS Process: Concept of Operations (CONOPS)." http://acqnotes.com/acqnote/acquisitions/concept-of-operations-conops.
- ANSI/EIA. 2003. *Processes for Engineering a System. ANSI/EIA 632-2003*. Philadelphia, US- PA: American National Standards Institute /Electronic Industries Association.
- Baldwin, K., et al. 2012. "The United States Department of Defense Revitalization of System Security Engineering Through Program Protection." Systems Conference (SysCon), IEEE International. Vancouver, British Columbia (Canada): IEEE.
- Balestrini-Robinson, S., D. F. Freeman, and D. C. Browne. 2015. "An Object-oriented and Executable SysML Framework for Rapid Model Development." *Procedia Computer Science* 44: 423-432.
- Balestrini-Robinson, S., D. F. Freeman, J. Arruda, and T. R. Ender. 2015. "ERS TRADESPACE: A Collaborative Systems Engineering Framework." *Proceedings. AHS Systems Engineering Technical Specialists' Meeting*. Huntsville, US-AL: AHS.
- Barron, F. H., and B. E. Barrett. 1996. "Decision quality using ranked attribute weights."
- *Management Science* 42(11): 1515-1523.
- Charnetski, J. R., and R. M. Soland. 1976. "Technical Note – Statistical Measures for Linear Functions on Polytopes. *Operations Research*, 24(1): 201-204.
- Charnetski, J. R., and R. M. Soland. 1978. "Multiple-Attribute Decision Making with Partial Information: The Comparative Hypervolume Criterion." *Naval Research Logistics Quarterly* 25(2): 279-288.
- Defense Acquisition University. 2011. IPS Element Guidebook, 2- Design Interface: 2.3 Key Performance Parameters (KPPs) and Key System Attributes (KSAs). *U.S. Office of the Assistant Secretary of Defense for Materiel Readiness.*
- Fisher, R. A. 1936. "The Use of Multiple Measurements in Taxonomic Problems." *Annals of Eugenics* 7 (2): 179–188.
- Goerger, S. R., A. M. Madni, and O. J. Eslinger. 2014. "Engineered Resilient Systems: A DoD Perspective," *Procedia Computer Science* 28: 865-872.
- Holland J. P. 2013. "Engineered Resilient Systems (ERS) Overview" Presentation at the U.S. Army Engineer Research and Development Center (ERDC) December.
- Johnson, E. R., G. S. Parnell, S. N. Tani, and T. A. Bresnick. 2013. "Perform Deterministic Analysis and Develop Insights. "In *Handbook of Decision Analysis*, 166-226.
- Lahdelma, R., J. Hokkanen, and P. Salminen. 1998. "SMAA-Stochastic multiobjective acceptability analysis." *European Journal of Operational Research* 106(1): 137-143.
- Lahdelma, R., K. Miettinen, and P. Salminen. 2003. "Ordinal Criteria in Stochastic Multicriteria Acceptability Analysis (SMAA)." *European Journal of Operational Research* 147(1): 117-127.
- Pflanz, M., C. Yunker, F. N. Wehrli, and D. Edwards. 2012. "Applying Early Systems Engineering: Injecting Knowledge into the Capability Development Process." *Defense Acquisition Research Journal* 19(4): 422-433.
- Ryan, E. T., D. R. Jacques., and J. M. Colombi. 2013. "An Ontological Framework for Clarifying Flexibility-Related Terminology via Literature Survey." *Systems Engineering* 16(1): 99-110.

- Roszkowska, E. 2013. "Rank Ordering Criteria Weighting Methods – A Comparative Overview." *Optimum Studia Ekonomiczne* 5 (65): 14-33.
- Sitterle, V. B., M. D. Curry, D. F. Freeman, and T. R. Ender. 2014. "Integrated Toolset and Workflow for Tradespace Analytics in Systems Engineering." In *INCOSE International Symposium* 24(1): 347-361. Las Vegas, US-NV: INCOSE.
- Sitterle, V. B., D. F. Freeman, S. R. Goerger, and T. R. Ender. 2015. "Systems Engineering Resiliency: Guiding Tradespace Exploration within an Engineered Resilient Systems Context." *Procedia Computer Science 44*: 649-658.
- Smaling, R. M., and O. L. de Weck. 2004. "Fuzzy Pareto Frontiers in Multidisciplinary System Architecture Analysis." *AIAA Paper 4553*: 1-18.
- Spero, E., M. Avera, P.Valdez, and S. Goerger. 2014. "Tradespace Exploration for the Engineering of Resilient Systems." *Procedia Computer Science* 28: 591-600.
- Tervonen, T., and R. Lahdelma. 2007. "Implementing Stochastic Multicriteria Acceptability Analysis." *European Journal of Operational Research* 178(2): 500-513.
- US Department of Defense. 2013. Interim Department of Defense Instruction 5000.02. *Operation of the Defense Acquisition System*. Washington, DC (US): Office of the Under Secretary of Defense for Acquisition, Technology, and Logistics.
- Waskom, M. 2015. http://stanford.edu/~mwaskom/software/seaborn/index.html.

ABOUT THE AUTHORS

[Editor: Author biographies were current when the paper was initially published in 2016.]

Dr. Valerie B. Sitterle is a Senior Research Engineer at the Georgia Tech Research Institute. Her primary areas of research include developing frameworks and associated analytical methods for the design and characterization of complex defense systems and systems-of-systems. Dr. Sitterle earned a Ph.D. in Mechanical Engineering at Georgia Tech, a BME and MS in Mechanical Engineering from Auburn University, and an MS in Aerospace Engineering and Engineering Science from the University of Florida.

Erika L. Brimhall is a Research Engineer II at the Georgia Tech Research Institute. She has nine years of experience in aerospace and defense, including working on the Space Shuttle main propulsion system at Kennedy Space Center, radar signal processing at Raytheon Missile Systems, and systems modeling at GTRI. She earned her B.S. in Aerospace Engineering from the Florida Institute of Technology and her M.S. in Systems Engineering from the Johns Hopkins University. She is currently leading a system architecture project.

Dr. Dane F. Freeman is a Research Engineer II at the Georgia Tech Research Institute in the Systems Engineering Application Branch. His primary area of research includes the development of systems engineering tools, probabilistic design space exploration, and product family design. He earned a B.S. in Mechanical Engineering from the University of New Orleans, and an M.S. and Ph.D. in Aerospace Engineering from Georgia Tech.

Dr. Santiago Balestrini-Robinson is a Senior Research Engineer at the Georgia Tech Research Institute (GTRI). His primary area of research is the development of opinion-based and simulation-based decision support systems for large-scale system architectures. He earned a B.S., an M.S., and Ph.D. in Aerospace Engineering from the Georgia Institute of Technology.

Dr. Tommer R. Ender is a Principal Research Engineer at the Georgia Tech Research Institute and serves as Chief of the Systems Engineering Research Division. His primary area of research includes the development of systems engineering tools and methods as applied to complex systems-of-systems, concerned with supporting decision-making through a holistic treatment of various problems. Dr. Ender is an instructor in Georgia Tech's Professional Masters in Applied Systems Engineering. He earned a B.S., M.S., and Ph.D. in Aerospace Engineering from Georgia Tech.

Dr. Simon R. Goerger currently serves as the Director of the Institute for Systems Engineering Research, U.S. Army Engineer Research and Development Center (ERDC). He was previously served as a Colonel in the U.S. Army, where his appointments included Director of the DoD Readiness Reporting System Implementation Office in the Office of the Undersecretary of Defense, Joint Multinational Networks Division Chief for the U.S. Army Central Command in Kuwait, and Director of the Operations Research Center of Excellence in the Department of Systems Engineering at the United States Military Academy (USMA). He received his B.Sc. from the USMA, and his M.Sc in Computer Science and his Ph.D. in Modeling and Simulation both from the Naval Postgraduate School.

S. W.Hensley et al. continued from page 29

Michael Henshaw is professor of systems engineering and leads the Engineering Systems of Systems (ESoS) Research Group. His research focuses on integration and management of complex socio-technical systems, with a particular emphasis on the challenges of through-life management of systems and capabilities. Professor Henshaw graduated in applied physics, and his early research focused on laser-plasma interactions, using computational fluid dynamics to investigate various phenomena in applications such as X-ray lasers. He joined British Aerospace (later BAE Systems) as an aerodynamicist and worked for seventeen years in aeronautical engineering tackling problems associated with unsteady aerodynamics (computational and experimental) and, later, multi-disciplinary integration. He received an appointment to a chair in systems engineering at Loughborough in 2006 to direct the large multi-university, multi-disciplinary program in network-enabled capability.

C. E. Siemieniuch is a professor of enterprise systems engineering and a member of the Engineering Systems of Systems Research Group in the Department. She is a chartered ergonomist and human factors specialist with expertise across the full range of systems-related human factors topics. Her key skills are in knowledge lifecycle management systems, organizational, and cultural aspects of enterprise modeling techniques, organizational systems architectures, dynamic allocation of function, and the design of complex systems. She is active in both the military and civilian domains.

Extending Formal Modeling for Resilient Systems Design

Azad M. Madni, Azad.Madni@usc.edu; **Michael Sievers,** Michael.Sievers@usc.edu; **Ayesha Madni,** amadni@usc.edu; **Edwin Ordoukhanian,** ordoukha@usc.edu; **Parisa Pouya,** pouya@usc.edu

■ **ABSTRACT**
Resilience is a much-needed characteristic in systems that are expected to operate in uncertain environments for extended periods with a high likelihood of disruptive events. Resilience approaches today employ ad hoc methods and piece-meal solutions that are difficult to verify and test, and do not scale. Furthermore, it is difficult to assess the long-term impact of such ad hoc "resilience solutions." This paper presents a flexible contract-based approach that employs a combination of formal methods for verification and testing and flexible assertions and probabilistic modelling to handle uncertainty during mission execution. A flexible contract (FC) is a hybrid modelling construct that facilitates system verification and testing while offering the requisite flexibility to cope with non-determinism. This paper illustrates the use of FCs for multi-UAV swarm control in, partially observable, dynamic environments. However, the approach is sufficiently general for use in other domains such as self-driving vehicle and adaptive power/energy grids.

INTRODUCTION

Resilience, a non-functional characteristic, allows a system or system-of-systems (SoS) to continue to provide useful service in the face of disruptions (Neches and Madni 2011). Disruptions can be external, systemic, or human-triggered (Madni and Jackson 2009). Examples of disruptions in the operational context of multi-UAV swarms include hacked or compromised swarm member, loss of communication within the swarm or between specific swarm members, and loss of visibility due to extreme weather or sensor malfunction. Resilient responses to such disruptions can take a variety of forms depending on environment observability and available intelligence. These include: circumvent disruptions if we can anticipate; withstand disruption if within the designed performance envelope; and recover rapidly from the negative effects of disruptions outside the performance envelope. Practically speaking, this means dynamically extending system capacity to cope with disruptions restructuring or reconfiguring system under disruptions; and continuing to operate at a somewhat diminished but acceptable level. The system's design envelope includes system models and adaptation logic incorporated within the system model to produce the necessary resilient responses when such disruptions occur. Doyle (2016) defines resilience as "the ability to recognize unanticipated perturbations that fall outside the purview of the system model designed to help the system adapt to disruptions that lie outside the system's design envelope. This definition implies that resilience is concerned with monitoring the boundary conditions of the system's model for competence (how well resilience strategies match disruption demands), and then adjusting or expanding that model to accommodate better changing demands (Neches and Madni 2011). The critical issue here is assessing an organization's adaptive capacity (resource buffers that allow resources of a particular type to increase on demand to a maximum limit) relative to the challenge posed by the disrupting event to that adaptive capacity. Boundaries in the multi-UAV swarm context define the system's competent performance envelope relative to specific classes of disruptions and uncertainties. Therefore, resilience engineering in a certain sense is concerned with introducing transparency into an organisation's safety model with the purpose of determining when the model needs revision. In other words, resilience engineering is about monitoring a system's decision making to assess the system's risks and risk envelope relative to an unsafe operating boundary.

Risk monitoring implies proactive and automatic/semi-automatic monitoring of buffers, margins, and tolerances. Buffer capacity concerns the magnitude and type of disruptions a system can absorb or adapt to, without substantial degradation in system performance, or breakdown in the integrity of system structure. Flexibility is the ability of a system to restructure or reorganise itself in response to external changes or pressures (Madni 2009). The margin is the proximity of a system's operation regime relative to its designed operational performance envelope or boundary. Tolerance is the ability of a system to degrade gracefully (as opposed to collapsing) as stress/pressure increases, or when disruption magnitude or severity exceeds its adaptive capacity.

This paper presents a model-based approach that combines formal and probabilistic modelling to engineer a resilient system and verify their designs.

METHODS: FORMAL AND PROBABILISTIC MODELING OF SYSTEMS AND SOS

Formal modelling introduces rigour in system verification, testing, and reasoning. However, formal modelling has limita-

tions. The rigour in formal modelling comes at the expense of flexibility. Ideally, one wants sufficient formality to support model verification and testing, and sufficient flexibility to scale and cope with uncertainty. This recognition provided the impetus for this research.

Our modelling approach extends the concept of a "contract" in contract-based design (CBD) to address uncertainty and partial observability that contribute to non-deterministic system behaviour (Madni 2015 and Sievers 2014). CBD is a formal method for explicitly defining, verifying and validating system requirements, constraints and interfaces. An implementation satisfies a design contract if it fulfils guarantees when assumptions are true. This is the "assert-guarantee" construct used in CBD. The rationale for choosing CBD is that statements in contracts are mathematically provable. The limitation of a traditional contract or CBD is that the assertions are invariant. The key innovation in our approach is the relaxation of invariant assertions requirement to introduce flexibility in the contract. The resulting resilience contract (RC) is a hybrid modelling construct that combines traditional contract, and flexible assertions, with partially observable Markov decision processes (POMDP). A POMDP is a unique form of a Markov decision process that includes unobservable states and state transitions trained during system use. POMDPs introduce flexibility into a traditional contract by allowing incomplete specification of legal inputs and flexible definition of post-condition corrections (Madni 2015 and Sievers 2014). A resilience contract extends a traditional deterministic contract for stochastic systems.

Figure 1 shows a hierarchical resiliency model using SysML block definition notation. The system comprises two subsystems as shown. Each subsystem and the system have individual RCs that comprise parameters and operations associated with its POMDP. As described below, RCs are software agents that update a belief state (Figure 5) and determine the next action (Equation 2) based on observing element outputs, the current belief state, the transition probabilities associated with the current state, and the reward (or penalty) for taking a given action. A belief state represents an entity's most probable state.

The assertions associated with a resilience contract (RC) are flexible, and the techniques employed include in-use learning, uncertainty handling, and pattern recognition. An RC is developed at design time, and trained during system use ("learning"). It allows trading of model verification for model flexibility, and vice versa.

A POMDP model consists of a set of states, *S*, set of actions *A*, a set of observation *O*, a transition model, an observation model, and a reward model. The Markov assumption implies that we only need knowledge of what state we are in and not the trajectory taken to get to that state. Each state is associated with an action policy that determines whether to remain in that state or assert controls that guide the system along with a trajectory to another state. Implicit in this concept is determining the most likely system state which may be hidden and must be inferred by evaluating system outputs. Since this is inherently a stochastic process, you need flexibility for deciding what actions are needed.

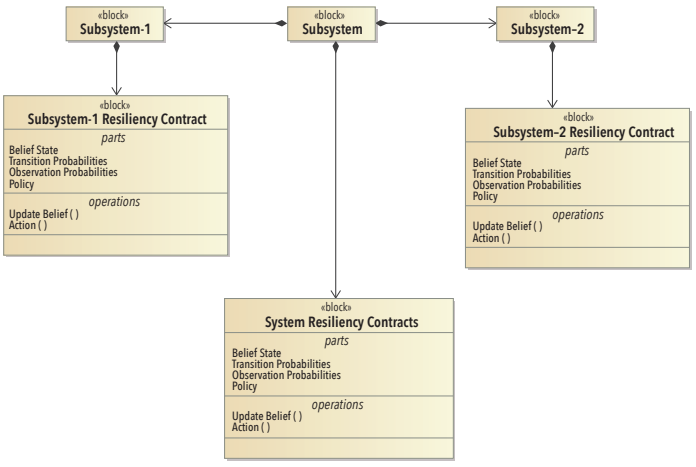

Figure 1. Resiliency model

Contract flexibility is introduced in several including relaxing the time invariance restrictions on the state space and action space, adding evaluation metrics for determining best action, or updating the emission and transition probabilities of hidden states. A critical insight in introducing flexibility in a standard contract is replacing the "assert-guarantee" construct with a "belief-reward" construct. We point out that this change provides the basis for incorporating flexibility into contracts without compromising model verification and testing benefits of traditional contracts to any appreciable degree.

RESULTS: UAV SWARM MODELING, BEHAVIOR PATTERNS, AND USE CASES

We chose UAV swarm control as our application domain. A UAV swarm is a system-of-systems (SoS) in which the elements can be either homogeneous or heterogeneous. The elements in the SoS cooperate to perform their assigned mission, or mutually agreed to tasks, and coordinate as needed. Each UAV in the swarm has sensors and communication facilities. UAV swarms participate in a variety of missions in the military and civilian sector. Exemplar missions include search and rescue, reconnaissance and surveillance, humanitarian assistance, and disaster relief.

The UAV swarm in our illustrative example is a swarm of quadcopters. To model and evaluate the system and SoS resilience, the questions that we need to answer are about model fidelity, model verifiability, and model flexibility. Fidelity pertains to the depth of modelling and the perspectives needed to answer the questions posed. Verifiability pertains to model correctness analysis. Flexibility pertains to the ease of extending or augmenting the model with reasoning mechanisms that introduce various forms of resilience. Ideally, we want just enough fidelity, and adequate flexibility to respond to disruptions. At the single UAV level, just enough flexibility means rudimentary dynamics of the UAV (quadcopter), basic sensor model, and a basic collision avoidance algorithm. The model could be run offline to generate parametric curves that could then be used to accept commands from the probabilistic model and generate new locations that can be used by the graphic visualisations. The model needs to support waypoint navigation and trajectory following. Moreover, the model should be easily replicable to realise SoS behaviour.

At the SoS level, the model needs to support different missions, communication protocols, and SoS configurations. The model should be capable of reflecting the behaviour of hacked or compromised UAV in the SoS, loss of communication, loss of a UAV, loss of sensing, and malfunctioning SoS member. At both the individual UAV level and the swarm level, it should be possible to evaluate different resilience concepts.

Behaviour patterns and use cases UAV swarm behaviours conveniently group into four behaviour patterns: deployment, en route, action on objective, and redeployment. Each behaviour pattern, associated with a mission phase, is discussed next.

Deployment (or takeoff) pattern: the act of putting SoS into operation. UAVs initiate operations and take flight. Variations in pattern come in the form of *Takeoff Method: Vertical* (VTOL), *Horizontal or Conventional* (CTOL), *Assisted* (*Mechanical or Human Catapult*, piggybacked from aircraft, propulsion assistance for short takeoff), and more; *Takeoff Order: Sequential* vs. *Parallel; Swarm Size, Hierarchy*, and *Homogeneity; Mission*: new, clean

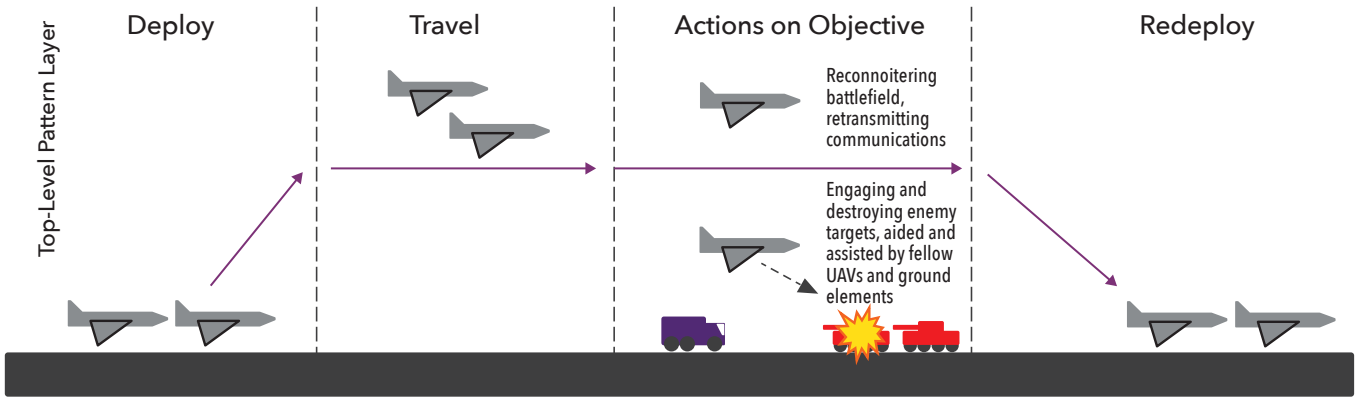

Figure 2. UAV SoS CONOPS

sheet deployment, or are UAVs reinforcing another UAV swarm; *Platform*: airfield, airport, grass field, naval ship, and improvised (such as a road or building top). The key factors affecting operation are mission-enemy-troops-terrain-and-weather-time available-civilian (METT-TC). An example of METT-TC factor is "an enemy has robust air-defence in area necessitating unique flight manoeuvres on takeoff."

Enroute (or cruise) pattern: the act of deployed swarm flying from one location to another in pursuit of the overall mission. UAV SoS objectives navigate as appropriate in support of global mission, pathfind at a local level, maneuver through terrain, weather, other UAVs in SoS, and neighboring systems not a part of SoS (coalition aircraft, enemy aircraft, and noncombatant aircraft), as well as making trade-offs in pathfinding and navigating in light of METT-TC. Variations in *Pattern* comes in the form of *Tactical Flight Considerations*; high altitude vs mid-altitude vs. nap of the earth vs. a combination; formation and disposition during the cruise; swarm size, composition, and capabilities (swarm heterogeneity factors); enemy air defence capabilities and presence; and weather.

Actions on Objective pattern: an essential part of overall CONOPS. Swarm achieves commander's intent and mission purpose. For example, *Reconnaissance, Observation, Sensing, Collecting, Aerial Communications Retransmit, Kinetic: Destroy* enemy assets; neutralise enemy unit.

We can tactically address UAV SoS *Objectives* at a local level both as individual systems and as a swarm to successfully execute actions on objectives and deploy UAV Systems as a SoS to achieve desired tactical and operational objectives in the battlespace. Variations in pattern – highly METT-TC dependent; examples: coordinated payload delivery to destroy a bridge and conduct recon; battlefield sensing and communications retransmission to support a focused, ground-based operation; routine mapping and imagery collection; search and rescue operation to locate downed aircraft in suspected geographical "crash window."

Redeployment pattern: the act of safely taking SoS out of operation. UAVs must RTB (return to base) and land while preserving themselves and collected data (if held onboard). Variations of pattern: *Landing Method: Vertical* (VTOL), *Horizontal* or *Conventional* (CTOL), *Assisted* (tail hook and cable, parachute landing or drag chute once landed), *Landing Order: Sequential* vs. *Parallel, Swarm Size*, *Hierarchy*, and *Homogeneity*, *Mission*: new, clean sheet deployment or are UAVs reinforcing another UAV swarm, platform: airfield, airport, grass field, naval ship, improvised (a road or building top), other METT-TC factors: such as enemy has robust air-defence in area necessitating unique flight manoeuvres on landing. Hasty landing: such as a damaged UAV improvises and lands in a clear area and sends out a distress signal.

Each basic pattern can adapt and decompose into multiple more nuanced, specific scenarios using METT-TC considerations that apply to the SoS mission. Fundamental concepts for top layer patterns are adapted and developed for highly specific use cases (fundamentals of an attack apply, but tactics behind attacking an enemy tank column vary – in the open versus enemy ground troops in wooded mountains). The right level of decomposition and detail for each top-level pattern help answer questions about where to introduce resilience and how best to incorporate resilience logic/reasoning within the SoS.

Figure 3 shows the state transition diagram for a quadcopter. In this figure, some transitions are labelled with belief values, for example b (failed) ≥ 0.95 is the threshold of transition from *normal motors* to *failed motor*, transition happens if belief ≥ 0.95 that a motor has failed. Some transitions have fixed assertions, such as *failed motor and Operational, Transition* from *Evaluate Environment* to *Auto Plan Enabled* has three beliefs with different probabilities in our example. Auto planner determines the course

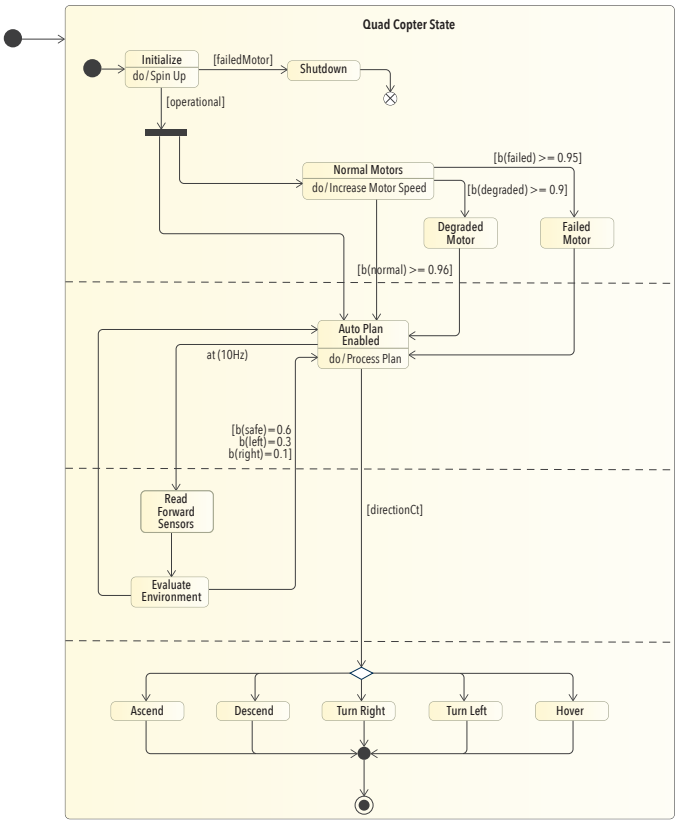

Figure 3. State transition diagram for multi-UAV SoS

Table 1. UAV swarm is a SoS

▪ Operational independence of UAVs
• UAVs operate independently to satisfy mission requirements
▪ Managerial independence of UAVs
• UAVs can be governed independently while being part of the swarm
▪ Evolutionary development of SoS
• development and existence is evolutionary with functions, and purposes added, removed, and modified with experience and need
▪ Emergent SoS behavior
• UAV-SoS performs functions and carries out purposes that do not reside in any single UAV
• UAV-SoS behaviours are emergent – cannot be realised by a single AV
▪ Geographic distribution
• UAVs are displaced in space and time and primarily exchange information

of action to take based on environment beliefs, motor condition beliefs, and the goals (action taken is the one that maximises reward or minimises penalty).

A UAV swarm can be viewed as a system-of-systems (SoS) because multiple UAVs need to cooperate to accomplish an end-to-end mission. A UAV swarm, especially a different swarm, exhibits the characteristics of a SoS (Table 1).

UAV SWARM CONTROL ARCHITECTURE AND CONOPS

Figure 4 presents swarm control architecture based on creating an optimal policy based on belief estimates provided by the state estimator. The state estimator relies on observations from the UAV swarm, environment sensors, and MDP belief model to generate updated belief estimates. Policy actions act on the UAV swarm and are used by the state estimator to update state information.

A simple example is presented to convey the fundamental ideas of UAV swarm control. In this example, the UAV swarm needs to turn either left or right to avoid an obstacle. There is uncertainty regarding the location of the threat, in that the threat could be to

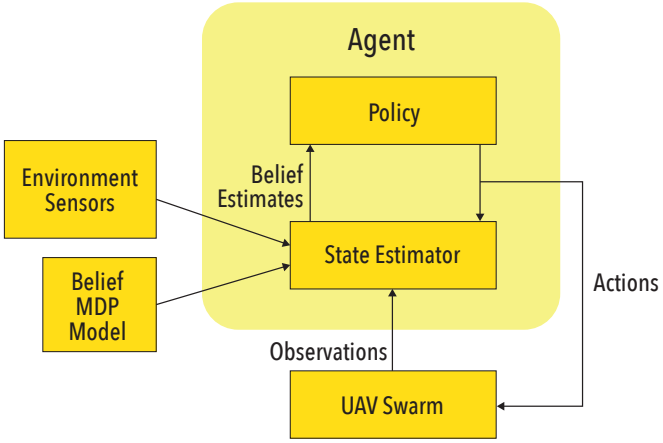

Figure 4. Example swarm control architecture

the left or the right of the swarm. A decision needs to be made to veer left or veer right. If the swarm veers right and the threat is to the right, serious consequences could ensue. The same is true if the swarm veers left and the threat is heading left. There are three possible actions that the swarm can take: veer left; veer right; continue flying straight ahead and collect more observations on the threat. POMDP policy for this simple concept of operations (CONOPS) has to deal with considerations such as UAVs not crashing into each other; all UAVs getting safely to their destination; UAVs avoiding potentially disruptive events; if one or more UAVs is shot down, the remaining UAVs reorganize and reallocate functionality to ensure accomplishment of mission objective to the extent feasible. The key ideas behind an optimal POMDP policy are two-fold: a POMDP policy maps current belief into action, and an optimal POMDP policy is a continuous solution of a belief MDP. Figure 5 shows the equation for summation of outcomes based on the path the UAVs take. The equation normalises the rewards and penalties. As shown in Figure 5, the system starts with a 50-50 belief that the threat could be to the left or the right. The system makes an observation. The system notices a potential threat to the left. So, the system moves its belief to the left as shown in the figure. That is, there is a greater belief that the threat could be to the left. Also, the system does not observe anything to the right. Thus, belief undergoes an update in accord with Bayesian analysis using observation and current state.

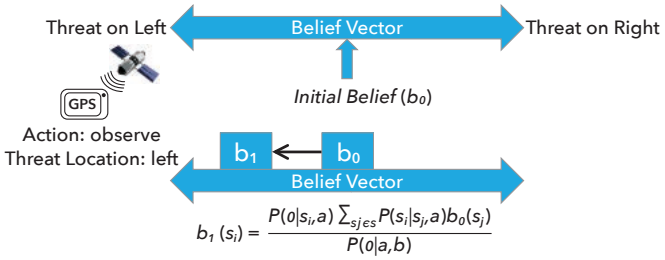

Figure 5. Iterative update of beliefs

A fundamental problem with state space models is that they are subject to combinatorial explosion. Several methods can be applied to contain this explosion, including pruning (Bellman 1957); branch and bound (Morrison et al. 2016), heuristic search (Szer 2012), Monte Carlo search (Browne 2012), and policy tree (Golovin 2010). Additionally, we relax the strict Markov assumptions by including heuristic analyses that use state trajectories when necessary for reducing ambiguities that increase the cost of computing the most likely belief state.

ILLUSTRATIVE EXAMPLE

The illustrative example is associated with quadcopters tasked to accomplish a mission while avoiding obstacles in an environment that is only partially observable. The architecture for this small SoS layers with each layer assigned to a particular model type. There are different types of models associated with this small SoS: vehicle physics model; behavioural model; and Markov decision process model (Figure 6). We discuss each next.

Vehicle Physics Modeling. For vehicle physics models, we need "just enough fidelity" to accept action commands from a probabilistic model and drive various visualisation on the dashboard for situation awareness. The number of UAVs can grow, we need a sparse representation for vehicles. To this end, we chose quadcopters for our research because of their relative simplicity. Quadcopters, in general, are under-actuated systems in that 6 degrees of freedom (*X, Y, Z*, roll, pitch, and yaw) are controlled by only 4 rotors. These vehicles are nonlinear systems that require two non-linear controllers at the physics level, one for controlling

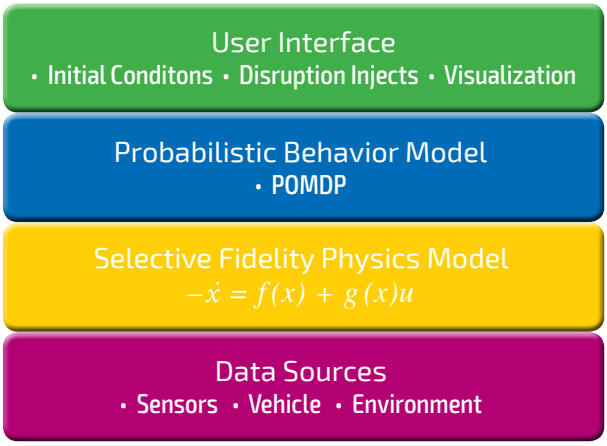

Figure 6. Layered system architecture

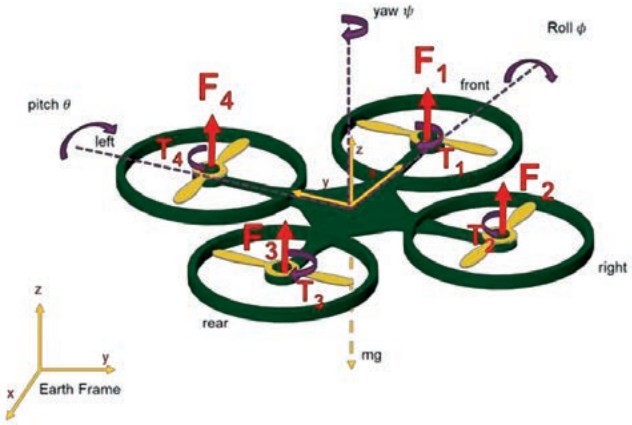

Figure 7. Quadcopter system

attitude, and the other for controlling position. Figure 7 shows a quadcopter for a particular orientation.

The physics model makes the following assumptions: a) the quadrotor is a rigid body with symmetric mass distribution; b) propellers are rigid; c) centre of gravity and body fixed frame origin are co-located; d) Earth's gravitational field (g), quadrotor's mass (m) and body inertia matrix (J) are constants; e) thrust factor and torque factor of motors are constants; f) inertia of motors and rotors is negligible; g) aerodynamic drag force is proportional to translational velocity; and h) rotation of the Earth relative to distant stars is negligible. With these assumptions, the model becomes simpler but requires further simplification to reduce computation.

Waypoint-waypoint path generation is done using an analytical function called Wymore's standard scoring function (Wymore 1993). Since the vehicle in the illustrative example is in a near hover mode, and errors in *X, Y*, and *Z* directions are negligible, we can predict the vehicle's position by calculating path coordinates without having to run a full dynamics model. Wymore's function used for path generation takes the form:

$$Position\_X = \frac{1}{1+\left(\frac{B-L}{t-L}\right)^{2*S*(B+t-2*L)}} \quad (1)$$

Where *t* is simulation time; *B* is midpoint time between two waypoints; *S* is minimum speed (S=*1*/(*B*–*L*)); and *L* is required time from a waypoint to keep *X, Y, Z* bounded.

Probabilistic Behavior Modeling. A basic but essential aspect of "sense-plan-act" model for UAV control pertains to decision making in uncertain operational environments. The decision-making associate members with UAV navigation and UAV health status monitoring make decisions that need to take into account: pre-defined UAV actions, overall UAV status, and environmental inputs. The vehicle uses sensed environmental observations for decision making. Against this backdrop, we employ Markov decision process models for representing vehicle behaviour. We illustrate the highlights of this approach in the following paragraphs.

We assume that the UAVs are operating in a "perfect world." This scenario means that the environment is wholly known (the location and orientation of obstacles and moving objects are known). The decision making reduces to a navigation and health monitoring problem that can be solved by Markov decision process models. To this end, our simple scenario calls for three UAVs that assigned a mission in which they are required to fly from a starting location to a pre-defined destination. The UAVs are capable of taking three types of actions (hover, land, and move north/east/south/west based on the health and status of fuel, batteries, and so on). For simplicity, we define an aggregated, state-of-health value in the range [*0, 3*] that represents observations of UAV fault, guidance sensor, and battery status. A value of *3* implies full capability, no faults, and battery power with margin to complete the mission and return to base. A value of *2* implies degraded operation, but the UAV is still able to complete its mission and return safely to base. The UAV has suffered a severe condition when its health is *1* and must return to base immediately. Finally, a value of *0* implies a survival condition in which the UAV must land immediately. In reality, UAV actions must account for the individual observations health and environmental monitors.

The navigation problem is solved by approximating the environment within a grid and then using Bellman's policy update equation (Bellman 1957) as shown in Equation 2.

$$\pi^*(S) = \text{argmax}_a \Sigma_{s'} \; T(S' \mid S, a)U(S')$$
$$U(S) = R(S) + \gamma \max_a \Sigma_{s'} T(S' \mid S, a)U(S') \quad (2)$$

In the above equations, π^* and *U* are the optimal policy and utility vector at the current state *S*, *a* is the current action, and *T* is the probability of transitioning from *S* to *S′*, given *a*.

The trajectories of the three UAVs, based on UAV health-status monitoring and MDP output, are shown in both 3D and plan-view in Figure 8. The motion (navigation) of the UAVs and their locations within the environment are shown in both the plan view and 3D diagrams. The stars are vehicles and the circles are their locations. Hovering in each location (represented by small cubes/squares) is represented as red circles in the plan-view diagrams, and altitude change in the 3D diagrams show the landing action. In this scenario, the Z axis remains constant so that the MDP applies to a 2D problem. The UAVs fly from the starting point on the ground to three parallel, pre-defined planes in the 3D environment and move towards the rendezvous point (goal). For safety purposes, multi- UAV coordination is done beforehand to maintain a safe horizontal distance among vehicles to minimises the probability of vehicle-to-vehicle collision during the conduct of the mission. The collision can happen when a UAV in a higher plane malfunctions and has to land immediately. When the UAV in a higher plane attempts to land it could collide with the UAV in the lower plane. By scheduling them one after each other, we avoid this problem. Thus, each UAV waits at the start point (defined on its plane) until the UAV at the next higher altitude flies some distance away.

In Figure 8a, the first UAV flies from the starting point, (*0, 0, 0*), to (*0, 0, 6*) and successfully travels towards the rendezvous point,

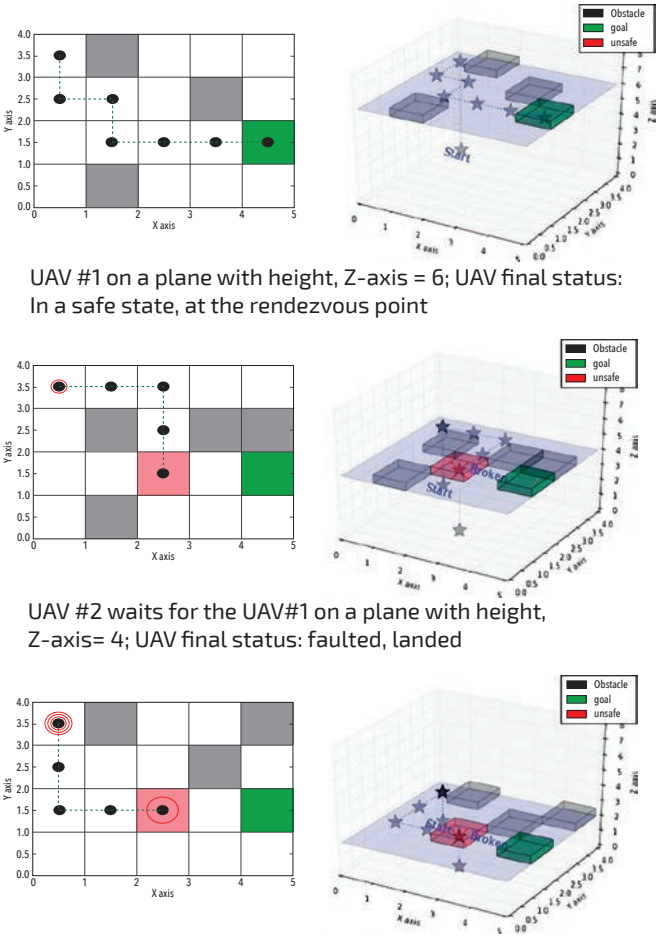

UAV #1 on a plane with height, Z-axis = 6; UAV final status: In a safe state, at the rendezvous point

UAV #2 waits for the UAV#1 on a plane with height, Z-axis= 4; UAV final status: faulted, landed

UAV #3 waits for the UAVs #2 & 1 on a plane with Z=2; UAV final status: hovering, faulted, and landed

Figure 8. Three UAVs flying to a rendezvous point; plots in the left-hand-side: plan views; plots in the right-hand-side: 3D environment

(*4, 1, 6*), since the state-of-health value of the first UAV remains above a pre-defined safety margin (above *2* on a scale of *0 – 3*). The second UAV, Figure 8b, flies from the start point to (*0, 0, 4*) in the second horizontal plane, waits for the first UAV to fly some distance away while it hovers at its current location, and then begins to move towards its destination. However, for this UAV, the state-of-health value drops to a meagre number (below *1* on a scale of *0 – 3*) indicating that the UAV has entered its survival state and must immediately land.

Finally, in Figure 8c, the third UAV is at its starting point in plane *Z=2*, awaits the departure of the first and second UAV. When the first and second UAV fly some distance away, the third UAV starts to head towards its destination. In this scenario, during its flight, its health value drops by *1* on a scale of *0 – 3*. When this happens, it hovers at its current location. Eventually, its health value falls below the safety margin (less than 1 on a scale of *0 – 3*) and it is forced to land.

Partially Observable Markov Decision Process Modeling: POMDP modelling can apply to both the localisation and mapping problems. In the localisation problem, we assume that the local map of the area of operation is known. However, the UAV does not know its position and orientation in its environment. POMDP modelling is applied to determine the location and orientation of the UAV.

In the second problem, the geography of the environment is unknown, but the UAV knows its position and orientation in the environment.

In the scenarios shown in Figure 8, we employed an MDP model to solve the navigation problem under the assumption that the UAV knows the location of obstacles. However, in the real world, the UAV may not wholly know its environment. Thus, the problem becomes a POMDP mapping problem in which the vehicle employs sensors to observe the environment, updates information about the environment, and then take actions. Since we can make no guarantees about the surrounding environment based on partial observability, the combination of observations and actions helps the vehicle continuously update its understanding of obstacles by updating a "belief state" using Equation 3. The UAV uses this belief state to make decisions about the direction of movement in the environment.

$$b(S') = P(S'|O,a,S) = \frac{\Omega(0,a,S')\sum_{S \in States} T(S',a,S)b(S)}{P(o|a,b)} \qquad (3)$$

Equation 3 illustrates how the belief state gets updated for state *S* as the UAV makes observations in that state ($\Omega(0, a, S)$). For instance, consider a scenario in which a UAV located within a 3 x 3 grid at location 6 and uses its sensors to identify obstacles in adjacent locations. Initially (step 0), the UAV has made no observations. In subsequent steps, the UAV points to adjacent locations and makes observations. A belief state is associated with the probability that there is an obstacle at each location. Table 2 shows the probability that an observation correctly identifies an obstacle. The entries with action *~Observe* imply that there is a sensor fault preventing an observation. The entries with action *Observe* allow for false positive observations. Table 2 presents the observation probability associated whether there is a wall or not depending on the observation made and action taken. In Figure 2, observation refers to whether there is a wall in front or not; action refers to whether you are performing the correct action, with or without observation; and probability refers to whether there is a wall or not depending on the observation and action. Thus, the first row of Table 2 implies that the probability of observing a wall without performing the "observe" action is *0.5*. The second row specifies the probability of "not observing a wall" without performing the "observe" action which is *0.5* again. The third row provides the probability of observing a wall while performing the correct action "observe" and it is *0.9*. The probability associated with the final row is *1– 0.9 = 0.*1.

Table 2. Observation probability

Observation	Action	Probability
Wall	*~Observe*	0.5
~Wall	*~Observe*	0.5
Wall	*Observe*	0.9
~Wall	*Observe*	0.1

There can be nine belief states corresponding to the belief there is an obstacle in a given location (N, NE, NW, S, SE, SW, E, W). If the UAV observes "no obstacle" in a specific direction, then the belief that there is an obstacle in the adjacent cell decreases; conversely, if the UAV observes an obstacle then the belief for the next cell increases.

Figure 9 shows the UAV in location 6 before any observations. At this point, the belief state vector is *1/8* for all locations and *0* for location *6* since the UAV is in location 6.

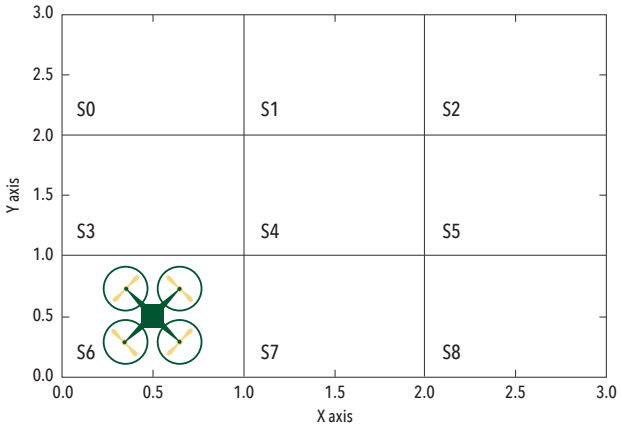

Figure 9. Step 0: action= none; observation = none

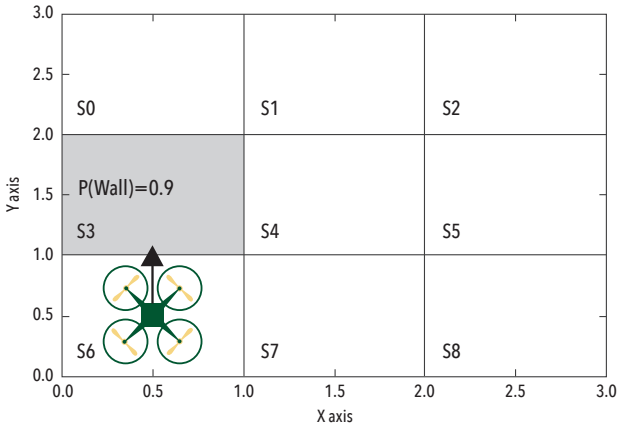

Figure 10. Step 2: action = observe; observation = 0.9 (Wall)

In Figure 10, the UAV makes an observation at the location associated with belief state, *S3*. Here we assume that the sensor has temporarily malfunctioned. This corresponds to the *~Observe* entries in Table 2, and the probability of an obstacle is *0.5*. The lack of an observation slightly reduces the probability of no obstacle at the location corresponding to belief state *S3*. In Figure 10, on the second attempt, the UAV observes an obstacle in the location corresponding to *S3*. From Table 2, the probability of a real positive is *0.9*. Now the value for belief state *S3* is significantly increased.

The UAV now makes another observation as shown in Figure 11. No obstacle is observed in the location corresponding to belief state *S7* so the probability of an obstacle from Table 2 is *0.1* and the belief there is an obstacle in *S7* decreases. Since the UAV knows that *S3* is not passible, the UAV could move to the location corresponding to state *S7* and again make observations.

Interestingly, although it was at *S6* and there were no obstacles there, it will again look to *S6* to determine whether a potential threat arrived. Similarly, the UAV will again observe the location corresponding to *S3* because the obstacle might have moved away or the obstacle might have been a false positive that confused the sensor. Ultimately, the actions taken will require more than simply computing probabilities, that is, there will be heuristics needed that reduce the state space and reduce the impact of Byzantine fault conditions (Sievers 2017).

SUMMARY AND CONCLUSIONS

This paper has presented a model-based approach for designing resilient systems. The approach combines formal modelling and probabilistic modelling to ensure requisite system verifiability and flexibility. The approach combines traditional contract from contract-based design (CBD) with flexible assertions and partially observable Markov decision process (POMDP) to create a hybrid modelling construct called a resilience contract (RC). An RC is well-suited to modelling complex systems such that we can address both system model verification and system flexibility. The approach enables both system and SoS model verification and offers the requisite flexibility to respond to disruptions. The approach shown is in the context of multi-UAV swarm control with several simplifications that do not limit the feasibility of the approach. For example, system state is a multi-faceted term that includes system's location, health, fuel status, sensor status, and more. For simplicity, we have used system location as system state in the figures presented. However, the approach is sufficiently general to be applied to a variety of SoS including autonomous vehicle networks and smart grids.

Future work will focus on solving the UAV navigation problem using POMDP. To this end, the first step will be to employ belief states along with sensed information from the environment to update the pre-defined reward and penalty values for the grid. In other words, the probabilities associated with the belief state (array) will be combined. Correctly, the pre-defined grid value and the new values will be used as new rewards/penalties to determine the best policy at that time. This will result in the UAV's action (movement) within the operational environment. The process of updating the belief state and grid values (rewards/penalties) will continue until the UAV reaches its destination. ■

ACKNOWLEDGEMENTS

This work was supported in part by Department of Defense Systems Engineering Research Center (SERC), RT-166 contract No. HQ0034-13-D-0004.

REFERENCES

- Bahill, T. and A. M. Madni. 2016. *Trade-off Decisions in System Design*. Cham, CH: Springer
- Bellman, R. A. 1957. "Markovian decision process." *Journal of Mathematics and Mechanics*, 1:679-84.
- Bellman, R. A. 1957. *Dynamic Programming*. Princeton, US-NJ: Princeton University Press, ISBN- 0-486-42809-5
- Browne, C., E. Powley, S. Lucas, P. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, and S. Colton. 2012. "A Survey of Monte Carlo Tree Search Methods." *IEEE Transactions on Computational Intelligence and AI in Games*, 4(1):1-49.
- Carlson, J. M. and J. Doyle. 2000. "Highly optimized tolerance: Robustness and design in complex systems." *Physical Review Letters*, 84 (11):2529.
- Csete, M. E. and J. C. Doyle. 2002. "Reverse engineering of biological complexity." *Science*, 295(5560):1664-9.
- Goerger, S. R., A. M. Madni, and O. J. Eslinger. 2014. "Engineered Resilient Systems: A DoD Perspective." *Procedia Computer Science*, 28:865-72.
- Golovin, D. and A. Krause. 2011."Adaptive Submodularity: Theory and Applications in Active Learning and Stochastic Optimization." *Journal of Artificial Intelligence Research*, 42:427-486
- Madni, A. M. and S. Jackson. 2009. "Towards a Conceptual Framework for Resilience Engineering." *IEEE Systems Journal*, 3(2):181-91.
- Madni, A. M. and M. Sievers. 2015. "A Flexible Contract-Based Design Framework for Evaluating System Resilience Approaches and Mechanisms." Presented at the IIE Annual Conference and Expo, Nashville, US-TN, 30 May – 2 June.

- Morrison, D., S. Jacobson, J. Sauppe, and E. Sewell. 2016. "Branch-and-bound algorithms: A Survey of recent advances in searching, branching, and pruning." *Discrete Optimization*, 19:79-102
- Neches, R. and A. M. Madni. 2012. "Towards affordably adaptable and effective systems." *Systems Engineering*,16(2):224-34.
- Sievers, M. and A. M. Madni. 2014."A flexible contracts approach to system resiliency." *Systems, Man and Cybernetics (SMC), IEEE International Conference*. https://ieeexplore.ieee.org/document/6974044
- Sievers, M. and A. M. Madni. 2017. "Contract-Based Byzantine Resilience in Spacecraft Swarms," *AIAA SciTech*. https://arc.aiaa.org/doi/abs/10.2514/6.2017-0644
- Szer, D., F. Charpillet, and S. Zilberstein. 2012. "MAA*: A Heuristic Search Algorithm for Solving Decentralized POMDPs." https://arxiv.org/pdf/1207.1359.pdf
- Woods, D. D. "Essential characteristics of resilience." 2006. *Resilience engineering: Concepts and Precepts*, 21-34.
- Wymore, A. W. 1993. *Model-based Systems Engineering*. Boca Raton, US-FL: CRC Press.

ABOUT THE AUTHORS

[Editor: Author biographies were current when the paper was initially published in 2018.]

Dr. Azad M. Madni is a professor of astronautics and executive director of the Systems Architecting and Engineering Program at the University of Southern California. He is the founder and CEO of Intelligent Systems Technology, Inc., a high-tech R&D company specialising in complex systems engineering. He is a life fellow of IEEE, IETE, and SDPS, and fellow of INCOSE, AAAS, and AIAA. He is the recipient of numerous awards and honours including the 2011 INCOSE Pioneer Award. He received his PhD, MS, and BS degrees in engineering from the University of California, Los Angeles. His research interests include adaptive cyber-physical-human systems and machine learning, formal methods in engineered resilient systems, and interactive storytelling in virtual worlds.

Dr. Michael Sievers is a senior systems engineer at Caltech's Jet Propulsion Laboratory and a lecturer in the System Architecting and Engineering Program at the University of Southern California. He conducts research in model-based systems engineering as well as contributing to many JPL's flight missions. He is a principal investigator on many harsh environment, fault-tolerant, high-performance computing research projects. He holds a PhD in computer science from the University of California, Los Angeles. He is an INCOSE fellow, a life senior member of IEEE, and life associate fellow of AIAA.

Dr. Ayesha Madni is a research scientist at the University of Southern California's Department of Astronautical Engineering. She specialises in developing high-performance teams, and in fostering creativity and innovation within research projects. She has worked as a research scientist over the past 10 years with diverse organisations ranging from non-profits to industry and government. Her research sponsors include the Gates Foundation, DARPA, ONR, NSF, CDE, and GM among others. This year she was awarded the Distinguished Achievement Award for STEM education by the Engineering Council of Southern California for her STEM-focused research in simulation and game-based learning.

Edwin Ordoukhanian is a PhD candidate in USC's Astronautics Department specialising in systems architecting and engineering. He is pursuing his doctoral degree under Professor Azad Madni, his dissertation adviser. Edwin is a teaching assistant and a research assistant on research projects led by Professor Madni in model-based approaches for driverless cars networks and UAV swarm control. Edwin received his MS in aerospace engineering from USC and his BEng in automation and control from the National Polytechnic University of Armenia.

Parisa Pouya is a PhD student in the Astronautics Department specialising in systems architecting and engineering. She is pursuing her doctoral degree under Professor Azad Madni. She is currently a teaching assistant in the SAE program and research assistant on research projects led by Prof. Madni. She received her MS in machine learning from the University of London, and MS in systems engineering from Loyola Marymount University.

Systems Theory Principles and Complex Systems Engineering Concepts for Protection and Resilience in Critical Infrastructure: Lessons from the Nuclear Sector

Adam D. Williams, adwilli@sandia.gov

SAND2019-PEER REVIEW. *Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the US Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525*

■ **ABSTRACT**
Part of the Presidential Policy Directive 21 (PPD-21) (PPD 2013) mandate includes evaluating safety, security, and safeguards (or nonproliferation) mechanisms traditionally implemented within the nuclear reactors, materials, and waste sector of critical infrastructure—including a complex, dynamic set of risks and threats within an all-hazards approach. In response, research out of Sandia National Laboratories (Sandia) explores the ability of systems theory principles (hierarchy and emergence) and complex systems engineering concepts (multidomain interdependence) to better understand and address these risks and threats. This Sandia research explores the safety, safeguards, and security risks of three different nuclear sector-related activities—spent nuclear fuel transportation, small modular reactors, and portable nuclear power reactors—to investigate the complex and dynamic risk related to the PPD-21-mandated all-hazards approach. This research showed that a systems-theoretic approach can better identify interdependencies, conflicts, gaps, and leverage points across traditional safety, security, and safeguards hazard mitigation strategies in the nuclear reactors, materials, and waste sector. As a result, mitigation strategies from applying systems theoretic principles and complex systems engineering concepts can be (1) designed to better capture interdependencies, (2) implemented to better align with real-world operational uncertainties, and (3) evaluated as a systems-level whole to better identify, characterize, and manage PPD-21's all hazards strategies.

INTRODUCTION

Meeting the Presidential Policy Directive 21 (PPD-21) mandate that "Critical infrastructure must be secure and able to withstand and rapidly recover from all hazards (PPD 2013)" includes evaluating safety, security, and safeguards (or nonproliferation) mechanisms traditionally implemented within the nuclear reactors, materials, and waste sector of critical infrastructure. Critical nuclear infrastructure harnesses the energy released during nuclear fission, where atomic and subatomic particles collide in a sustainable chain reaction. Related benefits include baseload quantities of electricity or significant volumes of desalinated seawater (arguably in a manner that reduces carbon emissions), as well as generating radionuclides for medical uses (cancer treatments) and advanced technological development (oil well logging). However, some nuclear fission by-products become radioactive because

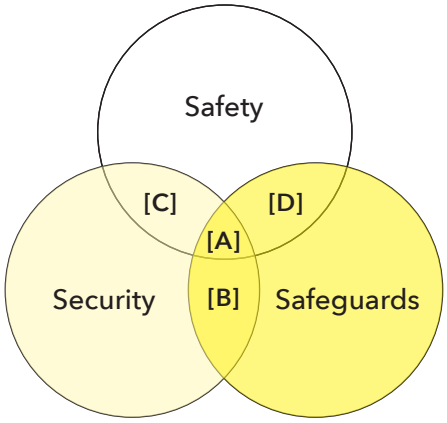

3S Interaction	Representative Example [Location on Venn Diagram]
Interdependency	Coordination of 3S responsibilities during emergency operations [A]
Conflict	Intrusive access control could impede evidence of peaceful uses (*increase safeguards risk*) [B]
Gap	Passive safety systems could be new targets for malicious acts (*increase security risk*) [C]
Leverage Point	Safeguards inspections could reveal a reactor vessel integrity issues (*reduce safety risk*) [D]

Figure 1. Types of interactions between safety, security, and safeguards in the critical nuclear infrastructure sector, with representative examples

of unstable nuclei which dissipate excess energy by spontaneously emitting alpha, beta, and gamma rays. Uncontrolled radiation can result in particular and psychologically fear-inducing impacts on human (poisoning and latent cancers) and environmental (land contamination and agricultural spoilage) health effects. To maintain these benefits—and minimize these health effects—the nuclear sector applies technologies, training, policies, and protocols to meet safety (preventing unintentional radiological releases), safeguards (preventing military use of nuclear technologies), and security (protecting against intentional radiological release or theft) objectives.

Protection and recovery efforts within the nuclear domain must include addressing not only traditional concepts of security, but also the long-standing emphasis on safety and the unique need for international safeguards. From this perspective, protection and resilience for nuclear facilities each consist of a complex and dynamic set of risks that are consistent with the PPD-21 call for investigating mechanisms to "strengthen all-hazards security and resilience" for critical infrastructure (PPD 2013). In the nuclear realm, this perspective is reflected internationally in calls by the World Institute of Nuclear Security (an international non-governmental organization) for an all-hazards approach to securing nuclear materials and the facilities (2019) and domestically by a National Academy of Sciences Committee conclusion that "The NNSA should adopt…a 'total systems approach' to characterize the interactions and dependencies of security (Committee on Risk-Based Approaches for Securing the DOE Nuclear Weapons Complex 2011, 1)." More specifically, in the words of former Deputy Director-General for Safeguards at the International Atomic Energy Agency, Olli Heinonen:

> Safeguards, security, and safety are commonly seen as separate areas in nuclear governance. While there are technical and legal reasons to justify this, they also co-exist and are mutually reinforcing. Each has a synergetic effect on the other, and authorities should carve out avenues for collaboration to contribute to the effectiveness of the nuclear order. For instance, near real-time nuclear material accountancy and monitoring systems provide valuable information about the location and status of nuclear material. This in turn is useful for nuclear security measures. Similarly, such information enhances nuclear safety by contributing as input to critical controls and locations of nuclear materials (2017).

Thus, to meet the primary PPD-21 objective of being able to "withstand and rapidly recover from all hazards" for the nuclear sector, strategies for the protection and resilience of nuclear materials and facilities must adequately address safety, safeguards, and security (3S) challenges—and the interactions between them (figure 1). In response, Sandia has explored the ability of systems theory principles and complex systems engineering concepts to better understand the complexities of the interactions between traditional safety, safeguards, and security mitigations in the nuclear sector. By investigating the complexity and dynamism in international spent nuclear fuel transportation, small modular reactors, and portable nuclear power reactors, this Sandia research identified key commonalities and unique outliers necessary to support a PPD-21-mandated all-hazards approach to protection and resilience for critical infrastructure.

SYSTEMS THEORY AND COMPLEX SYSTEMS ENGINEERING FOR PROTECTION AND RESILIENCE

Sandia's studies began by asserting that systems theory principles and complex systems engineering concepts provided a useful framing for characterizing the complexity in—and interactions between—nuclear safety, safeguards, and security in real-world operations. One such systems theory principle is hierarchy; wherein we articulate functional descriptions in terms of levels of complexity within a system. Systems theory argues that hierarchy is a useful framework for understanding, defining, and evaluating the characteristics that generate, separate, and connect these levels of complexity. By extension, this logic of hierarchy also asserts that higher ranking components/influences constrain the range of possible behaviors of components at lower levels. For example, the research indicated that the size (power output) of a given nuclear reactor constrains the types of safety, safeguards, and security mitigations implemented—and, thus, influences levels of protection and resilience for the nuclear activity.

The principle of hierarchy is directly related to the observed phenomena by which behaviors at a given level of complexity are irreducible to (and thus, inexplicable by) the behavior or design of its component parts. Called emergence, this concept describes how interactions among components within a system (or with environmental influences) drive system-level behaviors. Going beyond the ability for individually selected technologies, policies, and behaviors to achieve component-level goals, the logic of emergence captures the importance of the interactions between such components on achieving system-level objectives. Recent Sandia research concluded that considering nuclear activities as complex

systems afforded the benefit of evaluating safety, safeguards, and security as emergent properties—which matches the complexity observed when implemented in international or transboundary environments.

Given the importance of emergence, there is a need to better understand how interactions between components and with environmental influences impact the ability of systems to achieve their desired objectives. This is the principle of interdependence and describes how actions (or outcomes) in one component impact actions (or outcomes) in another. The principle of interdependence also addresses the concept of feedback—where output from component A's interaction(s) with other components (or environmental influences) influences the next set of inputs back into component A actions. In this research, the team evaluated safety, safeguards, and security for nuclear sector activities in terms of how each impacted—and was impacted by—both technical and non-technical (or, socio-political) components.

Current efforts in systems engineering aim to better combine these systems theory principles to design and operate ever increasingly complex systems. As systems increase in complexity, according to Keating, et al. (2003, 38), "it is naïve to think that problem definitions and requirements will be isolated from shifts and pressures stemming from highly dynamic and turbulent development and operational environments." If this is true, then engineering for complex infrastructure should also be cognizant of—if not explicitly incorporate—risk mitigation processes that form part of its operational environment. This Sandia research aimed to better address the multidomain interdependencies between long-established nuclear safety practices, internationally-mandated nuclear safeguards processes, and socio-technical nuclear security systems. Complex systems engineering offers the mechanism by which to design nuclear facilities in such a way to account for these safety-safeguards-security interdependencies by expanding design options to include non-traditional influences on system performance. Thus, it seems that invoking these systems theory principles and complex systems engineering concepts provide a strong foundation on which to build all-hazards strategies and mitigations for critical nuclear infrastructure protection and resilience.

SANDIA'S SYSTEM-THEORETIC APPROACH TO NUCLEAR SAFETY, SAFEGUARDS, AND SECURITY

In several studies—summarized in the next section—Sandia researchers demonstrated that 3S risk stems from interactions between technical, human, and organizational influences within critical nuclear infrastructure as complex systems. These studies also offer several useful conclusions for evaluating 3S risk complexity for critical nuclear infrastructure. First, integrated 3S approaches can help identify interactions—such as interdependencies, conflicts, gaps, and leverage points—across nuclear traditional safety, security, and safeguards approaches. Second, including the interactions between safety, safeguards, and security better aligns with real-world operational uncertainties and better describes the risk complexity associated with multi-modal, multi-jurisdictional systems in which critical nuclear infrastructure must operate. Third, we can design risk mitigation strategies resulting from integrated 3S risk assessments to better account for interdependencies not included in independent S assessments.

Table 1. Summary of systems engineering design goals for each type of interaction evaluated in Sandia's systems-theoretic approach to nuclear safety, security, and safeguards

3S Interaction	Systems Engineering Design Goal
Interdependency	Identify & (possibly) decouple
Conflict	Identify, eliminate, and/or reconcile
Gap	Identify, eliminate, and/or reconcile
Leverage Point	Identify & exploit

Other efforts in the nuclear sector have taken a range of approaches to explore 3S integration. One endeavor identified overlaps in regulations, procedures, and instrumentation to offer "3S-by-design" as a potential resource savings for nuclear utilities (using shared video surveillance data between safety, safeguards, and security) (Stein and Morichi 2012). Another used traditional risk management approaches to highlight analytical consistencies between these domains—namely by pairing the traditional security-related issue of sabotage with safety and traditional security-related issue of theft with safeguards (Cipollaro and Lomonaco 2016). In contrast, the Sandia grounded their studies in systems theory and complex systems engineering to illustrate interactions (Table 1) between risks and mitigations (interdependencies), characterize oppositional forces in operational risks (conflicts), identify missed operational risks (gaps), and capture natural redundancies or compensatory effects to mitigate risks (leverage points).

For this research, interdependencies refer to aspects of expected individual S operations whose operations are directly impacted by the behavior from operations in another S. Such relationships could include, but are not limited to, technical components that are collocated and/or use the same infrastructure; temporal processes that must be completed sequentially; or organizational policies that are predicated on specific technological capabilities. Sandia's 3S analysis sought to identify any interactions within the evaluated nuclear infrastructure sector that impacted—either positively or negatively—expected safety, safeguards, or security behaviors. For example, one interdependence for critical nuclear infrastructure relates to desired responses to a fire alarm. For safety, the primary goal is to evacuate facility personnel as quickly as possible. Yet, for security, the emphasis is on ensuring that the alarm is not a diversion for a malicious act (an adversary using the chaos as an escape mechanism). From this perspective, the interdependent need for security to verify the location of all personnel from sensitive areas of the nuclear facility while also meeting the safety need for timely evacuation presents a complex systems engineering design problem.

Often, integration-based analyses focus on identifying—and mitigating—conflicts. For this research, conflicts refer to aspects or objectives of expected individual S operations that negatively overlap with expected behaviors from a different S. Systems engineers commonly capture conflicts using various forms of trade space analysis within systems engineering by tracing their origins to either implementation, design, or requirements decisions. This research sought to expand on this tradition to identify negative interdependencies between safety, safeguards, and security—particularly where an improvement in the operations of one S resulted in a deleterious effect on behaviors in another S. For example, one conflict in critical nuclear infrastructure are

Table 2. Representative set of enhanced mitigation design goals identified from interdependencies, conflicts, gaps, and leverage points in safety, security, & safeguards activities

Case	Safety	Security	Safeguards	[3S Interaction Type] Systems Engineering Design Goal
SNF Transport	Better SNF access can help prevent unplanned radiological release	Focus on preventing unauthorized access during SNF transport	Fewer people with SNF access can enhance continuity of knowledge during transport	**[Leverage Point** Identify and exploit multiple benefits of focusing on preventing unauthorized access
Small Modular Reactors	Strict access controls challenge emergency operations	Strict access control procedures to offset fewer onsite security personnel	Strict access controls can provide assurance to safeguards inspectors	**[Conflict]** Identify, eliminate, or reconcile impact of access controls on emergency operations
Reactors	Scuttling as a last-ditch response to an accident	Scuttling raises questions on protection responsibilities	Scuttling raises questions on reporting and accountancy responsibilities	**[Gap]** Identify, eliminate, or reconcile benefits of scuttling on security and safeguards responsibilities

common practices related to transporting hazardous materials. For security of nuclear transportation, one point of emphasis is "need to know," or limiting who is informed about the transportation details (route and timelines). Yet, national safety regulations often require clear (and distinct) markings indicating that a given vehicle is carrying nuclear materials. So, an improvement in hazardous material marking for first responders directly impedes implementing "need to know" to meet security obligations. From this perspective, we could address this conflict by invoking systems theory principles into technical or procedural redesign.

In addition to conflicts, integrating across safety, safeguards, and security behaviors can identify operations or behaviors that we have not yet identified. For this research, gaps refer to aspects or objectives of expected individual S operations that we have not captured, mitigated, or otherwise addressed. Yet, this perspective also demonstrates that gaps can be positive and represent missed opportunities to improve system behaviors. For example, one gap common across critical nuclear infrastructure is coordination during emergencies involving nuclear materials. Much like emergencies with other hazardous materials, the safety (protect the public from undue harm) and security (protect the materials from malicious use) are well known. One unique (and often missing) aspect of critical nuclear infrastructure emergencies are the safeguards—which, from this perspective, is a gap that represents an opportunity for enhanced emergency operations. More specifically, coordinating completion of safeguards actions (maintaining continuity of knowledge of the location and amounts of nuclear materials) can improve safety and security operations by streamlining hazardous clean-up efforts and clarifying who has had access to the nuclear materials, respectively.

Lastly, in this research, leverage points refer to aspects or objectives of expected individual S operations that positively overlap with expected behaviors from a different S. In contrast to conflicts, leverage points are force multipliers between safety, safeguards, and security when an improvement in one S results in a simultaneous improvement in expected behaviors in another S. This research purposefully sought such relationships to demonstrate the concept that there are situations in which interdependence is desired. For example, consider the multiple responsibilities involved when nuclear material is in transit and must cross a national (or international) border. Because of the importance of adhering to all safety, safeguards, and security responsibilities along the entire transportation route, border crossings represent a transition in risk mitigation responsibility that can stretch traditional (isolated) inspection approaches. From this perspective, we could assign aspects of safeguards inspections to safety inspectors to take advantage of the larger number of qualified safety inspectors worldwide. Thus, already existing safety operations augment the need to meet continuity of knowledge of nuclear material responsibilities by designing jurisdictional transition inspections to leverage data commonly collected for safety purposes to meet safeguards obligations.

LESSONS FROM ACROSS THE NUCLEAR INFRASTRUCTURE SECTOR

Evaluating the risk complexity for different pieces of nuclear infrastructure demonstrated the applicability of this research to meeting the PDD-21 mandate for critical infrastructure. This section summarizes the technical evaluations of an integrated 3S approach to risks for three different nuclear infrastructure sector-related activities—spent nuclear fuel (SNF) transportation, small modular reactors, and portable nuclear power reactors. A representative set of how identifying interdependencies, conflicts, gaps, and leverage points can enhance 3S risk mitigation strategies is summarized in Table 2. In addition, these studies illustrate how using systems theory principles and complex systems engineering concepts can meet the PPD-21 call for an all-hazards approach.

Case 1: International Transportation of Spent Nuclear Fuel

Recent interest in new nuclear programs (United Arab Emirates and Vietnam) and the increasingly popular fuel take back agreements from existing nuclear power programs (Russia) indicate an expected increase in the amount of spent nuclear fuel (SNF) transported across the globe, including transfers of SNF casks between transportation modes (road to rail to water) and across geopolitical or maritime borders. SNF is nuclear material that has undergone fission within a reactor vessel and is now significantly radioactive. Risk mitigation for the international SNF transportation is challenging because of the likelihood that related mitigation resources and regulations along approved routes will be inconsistent. To investigate the resulting complexity in achieving 3S objectives, this study (Williams et al. 2017b) used a hypothetical SNF transportation across fictitious borders and between multiple conveyances. (For details

on the hypothetical case description, see Williams et al. 2017a)

Results from this study demonstrated that different analysis techniques, albeit in different ways, incorporated systems theory principles and complex systems engineering concepts to identify interdependencies, conflicts, gaps, and leverage points for risk mitigation. One interdependency identified how the negative health effects of the radiological release from exposure to SNF—an important factor in designing adequate security responses to SNF transportation accidents–would directly impact security responders' effectiveness. Consider advanced notification of SNF transportation details to local first responders as an example of a conflict. While the timeline for advanced notice can both shorten response and public evacuation times, it can also increase the possibility for and adversary to obtain route information. Two new states of increased risk—uncoordinated implementation of both standard operating procedures and operational emergency plans—emerged from several gaps identified in expected SNF transportation behaviors which evaluating safety, safeguards, and security individually missed. Other results identified leverage points for better mitigating the risks of SNF transportation, including how improved prevention of unauthorized access to the cask (for the security goal of preventing theft) also results in better mitigation of unplanned radiological releases (from a safety accident) and enhanced continuity of knowledge of material location (a safeguards issue).

Though representative of the larger study, these results highlight how hierarchy (constraining end-to-end SNF transportation risk), emergence (ensuring that inspections meet objectives), and interdependence (accounting for the impact of security protocols on security performance), as systems theory principles, better capture the real-world risk facing international SNF transportation. Similarly, identifying gaps (the potential for there to be no shipment oversight entity), interdependencies (the need to coordinate between security and emergency personnel after a notional train derailment), conflicts (inspectors may have contradictory safety and safeguards responsibilities), and leverage points (using security procedures to maintain continuity of knowledge for safeguards) provides the opportunity to use complex systems engineering to design better risk mitigation strategies. Using these insights resulted in a systems-based all-hazards approach for managing risk complexity in multimodal and multijurisdictional international SNF transportation.

Case 2: Small Modular Reactors

By design, small modular reactors (SMRs) will have a smaller operational footprint and generate substantially less energy than the current nuclear power plants (NPPs), thereby offering a significant relative cost reduction to current-generation nuclear reactors—increasing their appeal around the globe. In addition, SMRs offer a variety of passive (no additional energy is necessary for initiation) safety features intended to provide adequate core cooling to delay (or prevent) core damage in the event of a short-term station blackout. When combined with the small core size and lower power density design characteristics, the passive safety systems may provide an inherent degree of resilience to beyond design basis events not typically seen in traditional NPPs. Yet, this shift in focus from engineered active safety systems to passive safety measures has potential implications for not only safety, but also for safeguards and security of SMRs. This study conducted a technical evaluation on a hypothetical SMR (for more technical details, please see Lewis et al. 2012) based on light water reactor-based concepts and designs across a range of safety, safeguards, and security scenarios. (For more study details, please see Williams et al. 2018). Given the novelty of SMR technologies, this study identified the need to achieve the same levels of 3S risk reduction with reduced resources and applicability of current 3S technical analysis and best practice rules of thumb to SMRs as challenges to meeting the PDD-21 all-hazards approach.

Overall, the focus on this study on interactions between technologies, processes, and procedures related to safety, safeguards, and security identified several instances where traditional assumptions of independence did not fully capture likely SMR operational realities. In one example of an interdependency, SMR passive safety systems can reduce the chances of a safety incident, but simultaneously offer new potential targets that increase the security risk. For an example of a conflict, consider the popular argument that SMRs will have very few personnel and strict access controls. While such restriction of access can increase security against both external and internal adversaries (and increase the assurance of appropriate safeguards-related access), they can also challenge the ability for emergency personnel to adequately respond to accidents at the facility. This study also identified the gap in understanding how the tradition of physically separating reactor trains to reduce common cause safety failures also increases the complexity of an NPPs layout and potentially makes it easier for an aspiring proliferant to guide inspectors around sensitive facility areas. Despite some incongruity between SMRs and best practices, this study identified the possibility for increased safeguards inspections frequency (due to the technical reactor characteristics and assumed attractiveness of the nuclear materials) would also reduce chances for an insider adversary to perpetrate a malicious act against the facility.

Though seemingly obvious, these interactions are not often accounted for individual technical analyses available in the public domain. These representative examples also illustrate how key systems theory principles like hierarchy (the role of a smaller facility footprint on traditional safety, safeguards, and security mitigations), emergence (statements regarding by-design approaches for both security and safeguards in SMRs), and interdependence (the need to adequately secure passive safety systems) can improve risk mitigation for critical nuclear infrastructure. None of the interdependencies, conflicts, or gaps, identified in the study presented significant challenges to SMRs meeting safety, safeguards, and security objectives. Yet, they did identify leverage points where we could implement complex systems engineering concepts—designing for safety-safeguards-security interdependencies as part of the operational environment, for example—to gain efficiency and effectiveness in an all-hazards approach for protection and resilience for SMRs.

Case 3: Portable Nuclear Reactors

A recent solution to siting and construction challenges of traditional NPPs are portable nuclear reactors (PNRs), or power-generating reactors that we can move between locations with sub-gigawatt electricity generation capability. Several nations are in the beginning stages of deploying and operating PNRs—including the Offshore Floating Nuclear Plant by the Massachusetts Institute of Technology, the US Army's proposed mobile very small modular reactor (vSMR), China's floating small modular reactor, and Russia's floating PNR, the Akademik Lomonosov (which, according to media reports, docked in December of 2019 and has supplied 10GWh of electricity through January 2020 [Nuclear Engineering International 2020]). While such flexible redeployment comes with many operational benefits,

there remain many unanswered questions about PNRs and how their risks may differ in form from traditional land-based reactors. One of the most unique aspects is the fact that each PNR can be transported as a complete NPP, resulting in changing risk profiles as the PNR moves between territorial and international borders or as water-borne travel challenges the assumption that a PNR (and its safety systems) will remain upright for the duration of any accident. In response, this study conducted a technical evaluation on a hypothetical PNR based on the scant technical information available in the public domain. (For more study details, please see Williams et al., forthcoming)

The results of this study on PNRs highlighted the value of systems-level analysis of safety, safeguards, and security interactions in developing all-hazards strategies for critical infrastructure that differs significantly from the status quo. Take, for example, the interdependency between the need to scuttle (or, purposely sink) a floating PNR to prevent an adversary act from succeeding and the safeguards reporting and inspections obligations for the sunken nuclear material. This also represents a conflict—while scuttling a floating PNR might serve as an ultimate security risk mitigation for preventing theft and sabotage, doing so also directly impedes the safety objectives of protecting maritime environments and associated commercial interests from undue exposure to radionuclides. Other 3S conflicts for PNRs are directly related to the potential for inconsistent and different interpretations of international maritime laws. One interesting gap identified in the study relates to the implications of the potential loss of control of the entire floating PNR vessel—as this scenario may allow a non-nuclear state access to a fully functioning nuclear reactor, even if it is only for a short period of time. In contrast, one similarly interesting leverage point identified in the study relates to how we could use the anticipated increase in safety-related inspections of PNRs between use locations as opportunities for additional safeguards-related inspections and reporting.

The preliminary results from this study are a first step in identifying, mitigating, and preventing such risks from negating the tremendous opportunities—like more flexible, cost-efficient electricity generation for remote civilian areas—presented by PNRs. Overall, this technical evaluation concluded that the researchers expect no significant public health impacts, current international safeguards approaches will be challenged, and, we will need to overcome jurisdictional ambiguity (and current technological shortcomings) for adequate security. This study also illustrated how hierarchy (defining constraints by level of PNR mobility), emergence (ensuring 3S risk mitigations are adequate across all possible PNR states), and interdependence (accounting for more dynamism between 3S mitigations during PNR motion) as systems theory principles helped address the anticipated increase in complexity for PNR operations. Combining these principles with complex systems engineering concepts provides an integrated approach better capable of including operational environments into PNR designs. In so doing, it may be possible to develop general PNR performance requirements designed to ensure that systemwide, safety, safeguards, and security risk remains acceptable—a conclusion of this study that supports the PDD-21 all-hazards approach for critical infrastructure protection and resilience.

CONCLUSIONS AND IMPLICATIONS

In calling for an all-hazards approach for protecting critical infrastructure, PDD-21issued a new challenge to designing and implementing resilient systems and structures to meet societal needs among increasingly complex operational environments. Moreover, PDD-21 is a charge that implicitly points to insufficiencies in traditional approaches that seek to optimize individual domains in isolation—as exemplified by how the nuclear infrastructure sector traditionally treats safety. While seeking to optimize nuclear safety (or nuclear safeguards or nuclear security) may yield apparent improvements in risk reduction, doing so disregards key aspects of risk complexity that can significantly impact overall performance. In response, three recent Sandia studies evaluated the impacts and implications of exploring the interactions between safety, safeguards, and security risk mitigation in the nuclear infrastructure sector. Across these studies of international spent nuclear fuel transportation, small modular reactors, and portable nuclear reactors, incorporating systems theory principles (hierarchy, emergence, and interdependence) and complex systems engineering concepts (designing to include the operational context) produced higher fidelity results. These results included descriptions of risks missed by more traditional approaches and requirements for improving mitigation designs toward improved protection and resilience. Ultimately, these three studies demonstrated the utility of using systems engineering to incorporate interdependencies between safety, safeguards, and security controls for enhancing the overall performance of critical nuclear infrastructure.

Several important implications result from the conclusions of these three studies. First, risks for critical infrastructure are not necessarily independent—implying that protection and resilience efforts should address the potential for interdependency. Second, systems theory principles provide a useful mental model for describing interdependencies and complex systems engineering concepts help characterize potential solutions. More specifically, these principles and concepts help identify risks that traditional approaches miss, while simultaneously offering a wider set of potential mitigations to improve overall performance. Third, evaluating interdependencies, conflicts, gaps, and leverage points helps incorporate elements of the operational environment into system design—which has traditionally been a source of notable uncertainty in critical infrastructure risk. For example, explicitly evaluating desired safety, safeguards, and security behaviors as emergent properties in terms of these interactions directly results in opportunities to overcome traditional obstacles in risk reduction. Lastly, we enhance designing for protection and resilience in terms of all-hazards strategies when accounting for interdependence—whether between elements of risk itself or between isolated mitigations against elements of risk.

Though representative, these Sandia study results highlight opportunities to leverage interactions between critical infrastructure operations (and with operational environments) to guide desired behaviors to meet PDD-21's three strategic imperatives. Refining and clarifying functional relationships, employing systems theory principles and complex systems engineering concepts to design for leverage points and gaps/conflicts can strengthen critical infrastructure protection and resilience. Consider using SNF security inspections at a border crossing to support safeguards and clarifying security/safeguards responsibilities for floating PNRs in territorial waters, respectively. These principles and concepts can also enhance information exchange by providing a common mental model (focus on emergent behaviors in an operational environment for PNRs) and by coordinating multi-domain risk mitigations toward the same protection and resilience goals (3S coordination for SMR operations). Lastly, explicitly evaluating integration in terms of interdependencies, conflicts, gaps, and leverage points offers a

wider analysis function to develop solutions in support of critical infrastructure decisions more creatively. These results from Sandia's critical nuclear infrastructure studies describe the unique position that systems engineering has in meeting PPD-21's call to "address…in an integrated, holistic manner this infrastructure's interconnectedness (PPD 2013)"—and speaks to the role systems engineers can play in developing appropriate all-hazards strategies to enhance protection and resilience of critical infrastructure. ■

REFERENCES

- AP News. 2019. "Russia's Floating Nuclear Plant Sails to Its Destination." 23 August. https://www.apnews.com/208ba688c-44c4ff49de735ecb9a3963f.
- Cipollaro, A., and G. Lomonaco. 2016. "Contributing to the Nuclear 3S's Via a Methodology Aiming at Enhancing the Synergies Between Nuclear Security and Safety." *Progess on Nuclear Energy* (86): 31-39.
- Committee on Risk-Based Approaches for Securing the DOE Nuclear Weapons Complex. 2011. "Understanding and Managing Risk in Security Systems for the DOE Nuclear Weapons Complex (Abbreviated Version)." Report, National Research Council of the National Academies. https://www.nap.edu/catalog/13108/understanding-and-managing-risk-in-security-systems-for-the-doe-nuclear-weapons-complex.
- Divison of Nuclear Security. 2015. *Nuclear Security Series Glossary, Version 1.3*. Vienna, AT: International Atomic Energy Agency.
- Heinonen, O. 2017. "Nuclear Terrorism: Renewed Thinking for a Changing Landscape." Briefing at the Open Debate of the United Nations Security Council, New York, US-NY, 13 February. https://www.belfercenter.org/sites/default/files/files/publication/21317_Olli_UN_Briefing.pdf.
- Keating, C., R. Rogers, R. Unal, D. Dryer, A. Sousa-Perez, R. Safford, W. Peterson, and G. Rabadi. 2003. "System of Systems Engineering." *Engineering Management Journal* 15 (3): 36-45.
- Lewis, T., B. Cipiti, S. Jordan, and G. Baum. 2012. "Generic Small Modular Reactor Plant Design (SAND2012-10406)." Report, Sandia National Laboratories.
- Nuclear Engineering International. 2020. "Russian Floating Nuclear Plant Supplies 10GWh of Electricity to Chukotka." *Nuclear Engineering International Magazine*, 27 January.
- PPD (Presidential Policy Directive). 2013. PPD-21. *Directive of Critical Infrastructure Security and Resilience*. Washington, US-DC: PPD, Administration of Barack Obama.
- Stein, M., and M. Morichi. 2012. "Safety, Security and Safeguards by Design: An Industrial Approach." *Nuclear Technologies* (179): 150-155.
- Williams, A., D. Osborn, J. Bland, B. Cohn, C. Faucett, L. Gilbert, S. Horowitz, and J. Rutkowski. Fourthcoming. "System Studies for Global Nuclear Assurance and Security (GNAS): 3S Risk Analysis for Portable Nuclear Reactors (Volume I)—Technical Evaluation of Safety, Safeguards, & Security (SAND2019-TBD)." Report, Sandia National Laboratories.
- Williams, A., D. Osborn, J. Bland, J. Cardoni, B. Cohn, C. Faucett, L. J. Gilbert, R. Haddal, S. Horowitz, M. Majedi, and M. K. Snell. 2018. "System Studies for Global Nuclear Assurance & Security: 3S Risk Analysis for Small Modular Reactors (Volume I)–Technical Evaluation of Safety, Safeguards, & Security (SAND2018-12447)." Report, Sandia National Laboratories.
- Williams, A., D. Osborn, K. A. Jones, E. A. Kalinina, B. Cohn, M. Thomas, M. J.Parks, E. Parks, and Amir H. Mohagheghi. 2017a. "Hypothetical Case and Scenario Description for International Transportation of Spent Nuclear Fuel (SAND2017-13661)." Report, Sandia National Laboratories.
- ——. 2017b. "System Theoretic Frameworks for Mitigating Risk Complexity in the Nuclear Fuel Cycle: Final Report (SAND2017-10243)." Report, Sandia National Laboratories.
- World Institute for Nuclear Security. 2019. "WINS 2019 Annual Report: The Golden Threat of Nuclear Security." Report, World Institute for Nuclear Security. https://wins.org/wp-content/uploads/2019/03/WIN-102-0219_WINS_Annual_Report_2019_SCR.pdf.

ABOUT THE AUTHOR

[Editor: Author biography was current when the paper was initially published in 2020.]

Adam D. Williams is a principal R&D systems engineer in the Nonproliferation and Cooperative Threat Reduction Center, Sandia National Laboratories. He has a BS in mechanical engineering, MA in international affairs/national security, and PhD in human and systems engineering.

Engineering a Cyber Resilient Product Line

Patrice Williams, patrice.dillon.williams@raytheon.com; **Paula Moss**; **Susan Bataller**; and **Suzanne Hassell**

OVERVIEW OF A PRODUCT LINE STRUCTURE SUPPORTING MULTIPLE ARCHITECTURES

A product line consists of a managed core set of composable systems with scalable features and customizable variations. Critical mission threads may differ across the product line, but key product line architecture components support the implementing capabilities supporting a specific customer mission.

The choice to adopt a product line engineering strategy allows an organization to manage its assets for efficient use across business opportunities. This article uses an illustrative product line containing two separate but related architectural solutions, which include some similar and some unique hardware assets. Developing shared hardware assets conforming to both architectural constraints facilitates asset usage across the entire product line. Using the product line engineering factory configurator adapts the shared software asset supersets to these hardware assets. This approach comes from Meyer and Lehnerd 1997.

Figure 1 illustrates an example dual architecture product line for cryptographic solutions, supporting a core and an adjacent market. Each architecture

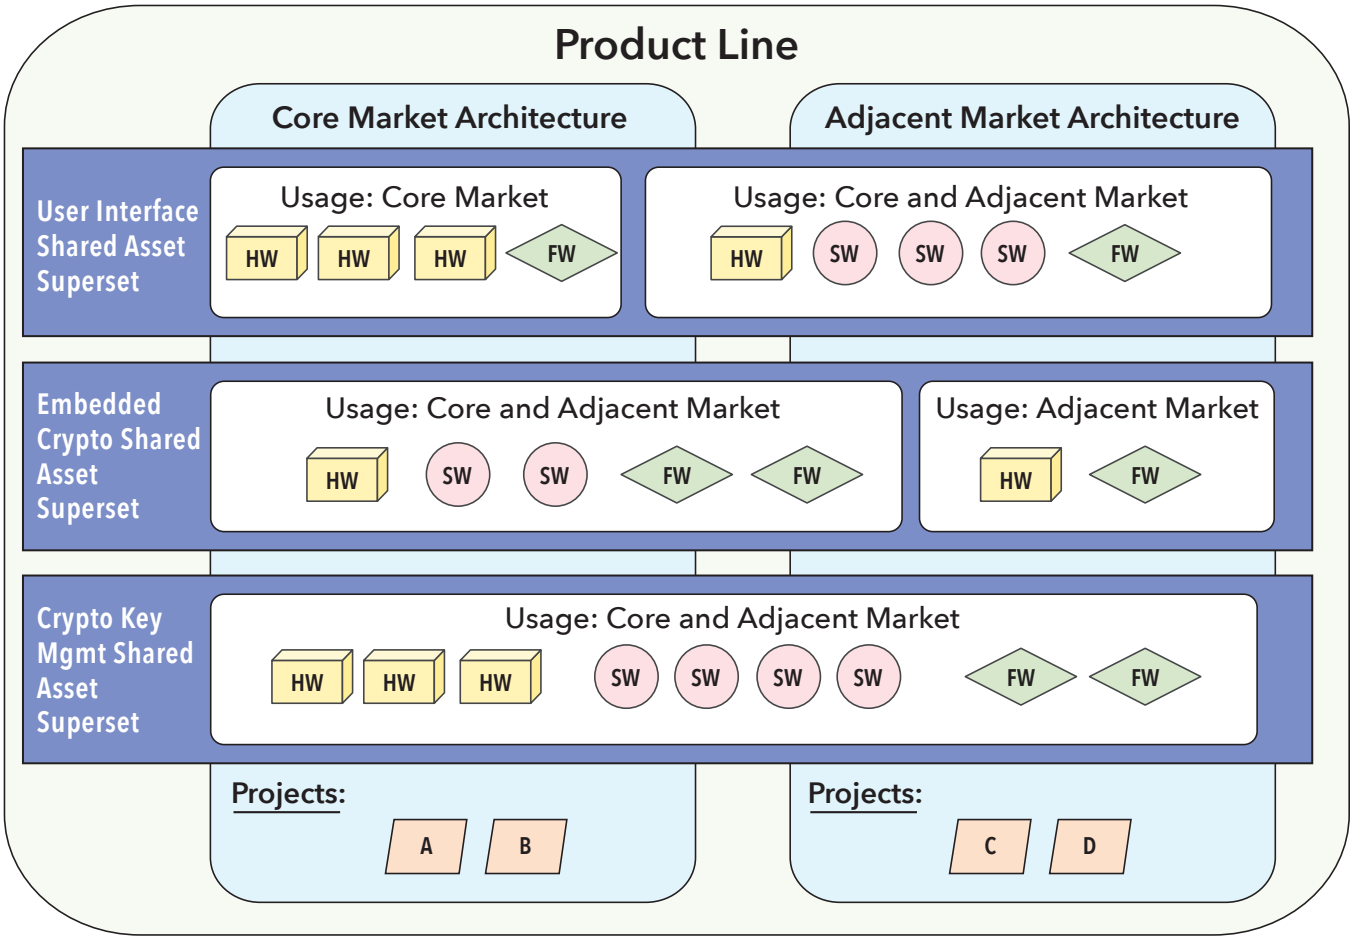

Figure 1. Example product line structure

solution supports the requirements and environmental considerations of the business opportunities within that market. The product line has organized its shared asset supersets into separate functional areas to create embedded cryptographic modules, crypto key management systems, and user interfaces supporting both architectures. Within each shared asset superset is a set of hardware, software, and firmware assets supporting the core's architecture and/or adjacent business opportunities. Finally, each architecture supports several projects, with each being a unique architecture instantiation.

A product line can provide a structured approach to effectively managing commonality and diversity in its product offerings. An organization may decide to stand up a product line based on its business forecasts in its core markets. Alternatively, once established in a core market, an organization may wish to leverage its existing products to establish a presence in identified adjacent markets. The potential cost savings associated with effectively leveraging existing projects in the new market is a significant motivator for adopting a product line approach.

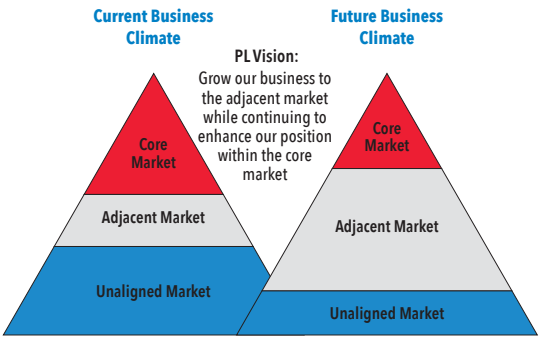

Figure 2. Product line business climate

Figure 2 shows the business climate for our example product line. The organization has an established business presence in its core market. The future business climate shows increasing opportunities in an adjacent market, with a corresponding decrease in core business opportunities. The organization anticipates reusing technologies and products developed in the core market will provide competitive advantages in the newly expanded adjacent market. Described below are the key components needed to establish a product line.

Vision: Create the vision for the product line, which identifies and binds its scope. This includes identifying the business opportunities within the current and adjacent markets the product line will support, and the business opportunities not in product line's scope.

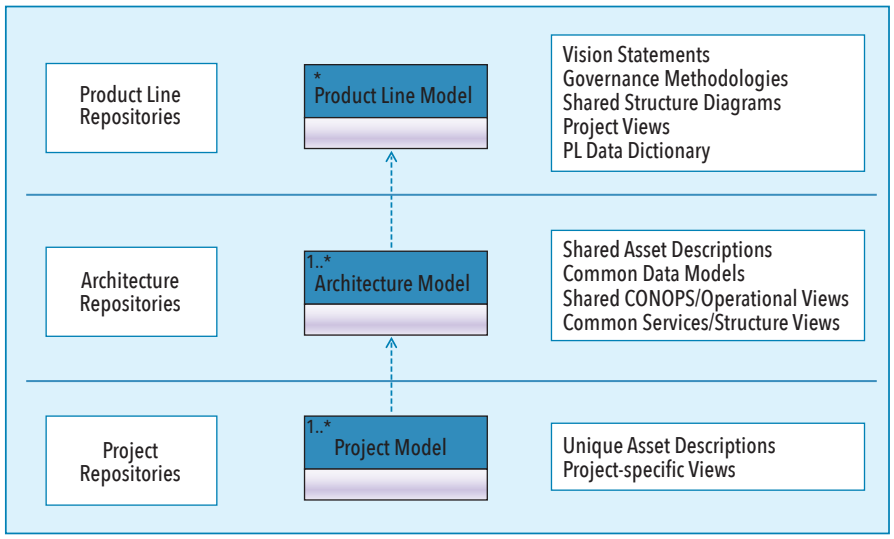

Figure 3. Example product line artifact organization

Product Line Guiding Practices: Identify overarching practices used across the product line. These include identifying development and change control strategies, funding models, and strategies for shared asset development with defined variation mechanisms.

Digital Environment: Establish a digital environment supporting the product line, which organizes digital models and repositories to promote artifact reuse. Defining, characterizing, organizing, and managing the product line artifacts is essential for efficient governance. Creating customizable templates for common product line documents can reduce the effort for each supported project.

Figure 3 provides a layered artifact organization for the example product line. The product line repository stores artifacts defining the product line structure, established governance methodologies, project portfolio views, and product line evolutionary plans. The architecture repositories, one for each identified architecture, store information common to that architecture, such as architecture views, data models, operational views, and shared asset inventory. Finally, each project creates a project-specific repository to contain its unique views and inventory of project-unique assets.

Table 1. Building block approach

QAs Supporting Building Block Approach	
Modularity	Composed of discrete components; minimized impact of change propagation
Interoperability	Ease by which one asset can exchange data with another; different operating environments
Adaptability	Ability to adjust to new conditions; the ability to adapt to new use or purpose

Table 2. Secure architectures

QAs Supporting Secure Architectures	
Confidentiality	Not disclosing information to unauthorized individuals or processes
Integrity	Not modifying or deleting protected data in an unauthorized and undetected manner
Availability	Ability to adjust to new conditions; modification ability for new use or purpose
Resiliency	The ability to adapt to changing conditions and to withstand and recover rapidly from disruptions.

The following recommended activities define each architecture within the product line.

Characterize the Architecture: Characterize the architecture by identifying the architectural principles and quality attributes which will define the architecture. Table 1 and Table 2 below show quality attribute examples (QA).

Create Common Architectural Views: Common operational, data, and system views define the common architecture supporting the member projects.

Identify Shared Architectural Constraints: Shared architectural constraints facilitate building coherent systems from the shared asset superset. Examples include:

- Using a Modular Open System Approach (MOSA) and associated support for a specific Open System Architecture (OSA)
- Hardware architecture constraints, such as size, weight, and power restric tions
- Software architecture constraints, such as layered architecture, service-oriented architecture, specific middleware, and programming languages

Identify Shared Security Attributes and Capabilities: Security capabilities provided by the architecture may exceed the protection's scope required by any one project.

Creating a project as part of a defined architecture has advantages for both the project and the product line. The following depicts the typical steps undertaken when adding a new project to an architecture solution within the product line. The project benefits by inheriting a defined architecture, and a set of shared assets available to provide the required capability. The project identifies available shared assets, and the variation points used to create the architecturally compatible asset instance. Developing new assets within a shared asset superset can provide value to other projects in the product line.

Figure 4 shows an example project's product line use. Project 'A' belongs to the Core Market business, and therefore inherits its overall architecture. The project develops two functions, for stand-alone cryptography, and for key generation and distribution. For the stand-alone cryptography function, assets identified within the crypto and the user interface shared asset supersets are good fits for the new function. The project also identifies the need for new assets, developed within the shared asset superset. Similarly, the key generation and distribution function will use existing shared assets from the crypto key management shared asset superset. Since project 'A' has some unique requirements for this function, unique assets will provide this functionality.

A key concern in a crypto architecture is ensuring the resulting system can achieve mission success in a cyber-compromised environment. This is also known as 'Cyber Resiliency,' associated with the "Resiliency" product line quality attribute. Section II describes the approach for applying cyber resiliency analysis within the product line.

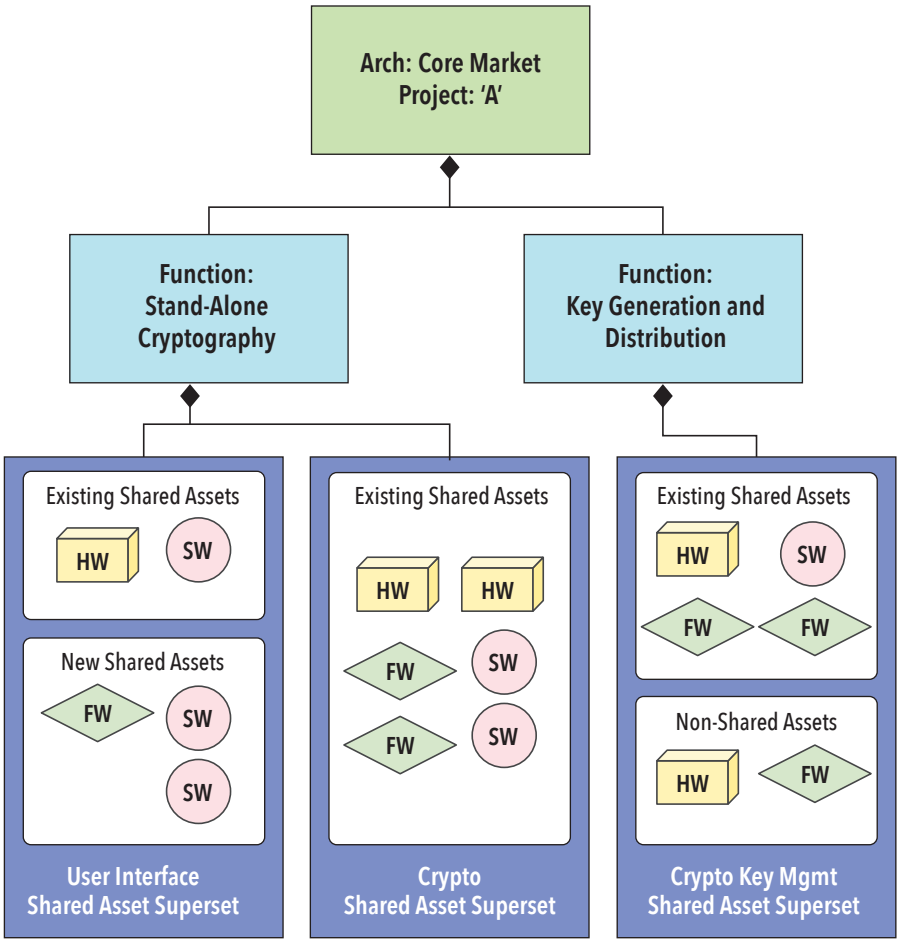

Figure 4. Example project use of shared and unique assets

CYBER RESILIENCY AND PRODUCT LINES

Cyber Resiliency is achieving mission success in a cyber-compromised environment. It anticipates a compromised system. Cyber hardening of systems is insufficient to ensure systems continue operating in a cyber-compromised environment. Brittle systems may result in unreliable system performance and failed missions in an environment with ever changing threats. Reusing Commercial-Off-the-Shelf (COTS) and Government-Off-the-Shelf (GOTS) hardware, software, and firmware has created a vast attack surface including undiscovered or unpatched vulnerabilities. Vulnerabilities can also invade the system at any point in the system supply chain. "Resilient computer network defense must anticipate the emergence of new vulnerabilities, take action to exploit these vulnerabilities, and disrupt the actions of successful intruders to increase their work factor and minimize their impact. The focus of resilience is the assumption that attackers are inside the network, we cannot detect them, and yet engineers must ensure mission survival (Hassell 2015)."

"Cyber secure and cyber resilient approaches focus on both protection from and reaction to a cyber threat. Cyber secure approaches focus on keeping the adversary out of the system. Cyber resilient approaches focus on mission success if an adversary can get into the system. The cyber resiliency wheel applies these techniques to interim system architecture products demonstrating the architecture decisions made to improve cyber resiliency (Hassell, Wilson, and Williams 2020)."

Cyber Resiliency analysis addresses key concerns assumed to happen during system operation. You may not know the specific cause but anticipate the resulting effect on the Systems of Systems. Cyber Resiliency has a focus on key

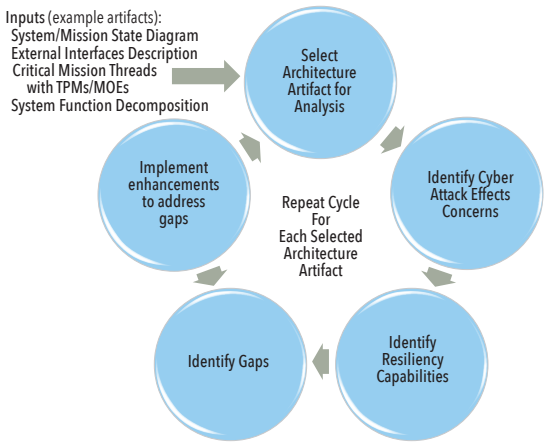

Figure 5. Resiliency wheel

Mission Threads and their associated Key Performance Parameters, Technical Performance Measures and Measures of Effectiveness. The Mission Thread analysis is an architecture-based tabletop analysis of customer concerns based on known or anticipated attacker effects and capabilities engineers have applied or will develop and field offsetting those concerns. This tabletop analysis includes a cross functional team including Operational subject matter experts, Software Engineers, Architects, and System Security Engineers. The Mission Thread binds the analysis timeframe to the time period during specific Mission Thread execution by the Systems of Systems. When identifying gaps, resources strategically allocated to implement enhancements close the gaps and any not closed track as a program risk. Figure 5 shows the process for performing the resiliency analysis.

The cyber resiliency analysis is made tractable by focusing only on key mission threads. Improving key mission thread component resiliency increases other mission threads' resiliency if they exercise the functionality in the improved components. If the improved components are a part of a product family, the resiliency lift applies across the product family.

Resiliency Concerns describe cyber attack effects on the System of Systems resulting from a cyber exploit. Architecturally, Use Cases describe normal system behavior. Misuse Cases describe resiliency concerns. A Misuse Case example is what happens to the Systems of Systems when it is under a Denial of Service attack. Table 3 provides a Resiliency Concerns list. It does not include all concerns applying from Systems of Systems inception to retirement.

Resiliency Capabilities can be capabilities built into the Systems of Systems or training and processes established for the Systems of Systems offsetting Resiliency Concerns. Resiliency Capabilities are proactive. They adhere to sound architecture principles such as "separation of concerns" and understanding and maximizing Quality Attributes such as "Trust." Table 4 derives, with some modifications, capabilities identified by Harriet Goldman (2010) of MITRE.

APPLYING THE RESILIENCY WHEEL: CRYPTOGRAPHY EXAMPLE

The following example uses the 5-step Resiliency Wheel to analyze the cryptography architecture.

Step 1: Using the Critical Mission Threads, Key Performance Parameters, and Measures of Effectiveness for the domain; determine the applicable architecture artifacts for the analysis. These include the Concept of Operations diagram, System Block Diagram, Activity Diagrams, Sequence Diagrams, and State Diagrams.

For our cryptography example, the critical mission thread is the Crypto Key Management Product Platform. This includes creating embedded cryptographic modules, crypto key management systems, and user interfaces.

The Customer Key Performance Parameters and Measures of Effectiveness for the Cryptographic System are:

Table 3. Resiliency concerns

Resiliency Concerns (Effects of Exploit)		
Data Exfiltration	**False Representation**	**Physical Effects**
Disrupt Connection	**Force Code Execution**	**Social Engineering**
False Information	**Force Supervisor Protected State**	**Software Exfiltration**

Table 4. Resiliency capabilities

Resiliency Capabilities		
Adaptive	**Diversity**	**Non Persistence**
Containment	**Forensics**	**Pre-emption**
Cyber Modeling	**Integrity**	**Prioritization**
Deception	**Least Privilege**	**Pro-active**
Detection	**Monitoring**	**Randomness/Unpredictability**
Distributedness	**Cyber Maneuver**	**Reconstitution**
		Redundancy

Table 5. Cryptography system resiliency concerns

Resiliency Concerns (Effects of Exploit)	
Disrupt Connection	Inappropriate storage of keys—Keys easily recovered by an attacker
Data Exfiltration	Key Re Use—Allows the attacker to crack the key
False Representation	Insider threat—Employees have access to keys

Table 6. Cryptography system resiliency capabilities

Resiliency Capabilities	
Adaptive	Audit log of key management
Containment	Policy to prevent reuse of keys—lifecycle management
Least Privilege	Role-based access to keys
Detection	Plan to detect key misuse within software
Diversity	Key rotation

- User Interface
- Key Distribution
- Key Management

The architecture information supports architecture tabletop discussions with the stakeholders. The stakeholder group includes the customer, architect, safety, and security engineers, and key system developers.

Step 2: Identify Resiliency Concerns. Table 5 describes the resiliency concerns resulting from the tabletop discussion.

Step 3: Identify Resiliency Capabilities mitigating the Concerns. Table 5 describes the resiliency concerns resulting from the tabletop discussion.

Step 4: Identify Gaps:

False information: Attacker cracks and manipulates keys (inaccurate information, malicious content attached).

Disrupt Communications: Keys stored improperly.

False Representation: Observe Operations for future malicious intent.

Step 5: Implement enhancements to the system(s) to mitigate the Resiliency Concerns.

The Resiliency Wheel should repeat when there are significant design changes to the system or changes to the operating environment raise new threat vectors. Cyber resiliency awareness should be an integral part of program system engineering activities.

Increasing cyber resiliency has emerged as a significant concern for both commercial and defense systems. When related systems belong to a product family, the effectiveness of adding resiliency to product modules accrues across the product family, reducing cost, schedule, and, most importantly ensuring mission success.

SUMMARY

Product line engineering provides a tremendous opportunity for organizations. Utilizing proven practices and technology allows an organization to focus on enhancements and features benefiting their customers. While many benefits to engineering a product line exist, adding Cyber Resilient practices need attention. These Resiliency measures help ensure mission success in a cyber-compromised environment. Applying the cyber resiliency wheel techniques and focusing on critical mission threads throughout the system, helps engineers evaluate the organization's most vital needs. Although, it is impossible to build a product line hardened against every cyber-attack, it is possible to build a product line with confidence using Cyber Resiliency techniques. ■

REFERENCES

- Goldman, H. 2010. "Building Secure, Resilient Architectures for Cyber Mission Assurance." Paper presented at the 201 Secure and Resilient Cyber Architectures Conference, McLean, US-VA, 29 October.
- Hassell, S. 2015. "Using DoDAF and Metrics for Evaluation of the Resilience of Systems, Systems of Systems, and Networks Against Cyber Threats." *INSIGHT* 18 (1): 26-28.
- Hassell, S., B. Wilson, and P. Williams. 2020. "Cyber Secure and Resilient Techniques for Architecture." Paper presented at the INCOSE International Symposium, online virtual event, 20-22 July.
- Meyer, M. and A.P. Lehnerd. 1996. *The Power of Product Platforms: Building Value and Cost Leadership*. New York, US-NY: The Free Press.

ABOUT THE AUTHORS

[Editor: Author biographies were current when the paper was initially published in 2020.]

Patrice Williams is a system security engineer with a focus on cyber security, MBSE, DevSecOps and secure architecture. Patrice joined Raytheon as a software engineer in 2009. She is a Raytheon cyber security resiliency architecture framework subject matter expert and has contributed to several cyber red team activities. She is currently participating in the 2020 CODE Center Cyber Rotational Engineering Program. Patrice has a BA in computer science and a Masters in cyber security with a focus in digital forensics.

Paula Moss is an engineering fellow at Raytheon Technologies and is a Raytheon certified enterprise architect. She has extensive experience with software, systems engineering, and architecture for command and control systems and is a cyber resiliency subject matter expert.

Susan Bataller is a senior engineering fellow at Raytheon Technologies and is a Raytheon certified enterprise architect. She has work extensively on military satellite communications systems in the systems, architecture, and software domains. Susan has been principally involved in conceptualizing, developing, and maintaining a product line architecture for Raytheon's military SATCOM products.

Suzanne Hassell is a principal engineering fellow at Raytheon Technologies and is a Raytheon certified enterprise architect. She has been at Raytheon since 2005, and is the Raytheon cyber resiliency subject matter expert and Raytheon Technologies Raytheon Intelligence and Space Chief Cyber Architect. She was the principal investigator for the US Army CERDEC Morphing Network Assets to Restrict Adversarial Reconnaissance (MORPHINATOR) program and led Raytheon resiliency research projects. Prior to coming to Raytheon, Suzanne did systems engineering, architecture, and software research and development in the communications industry for 23 years. She has 12 US patents.

Harmonizing the Domains of Loss-Driven Systems Engineering

Keith D. Willett, Keith.Willett@incose.org

■ **ABSTRACT**
System characteristics include *what it is* (structure, **state**), *what it does* (function, **behavior**), *where it resides* (**environment**, containing whole), *what it uses* (**resources**, energy source, raw material), *what it contains* (**content**), and *why it exists* (**value delivery**). An adversity produces a disturbance that can induce stress in a system so it may suffer some loss within one or more of these characteristics. Loss-driven systems engineering (LDSE) is an approach to address systemic loss in all forms helping ensure value delivery. LDSE domains include *reliability, sustainability, survivability, risk management, resistance, resilience, agility, safety*, and *security* which all work in harmony to avoid, withstand, and recover from loss. Traditional systems engineering treats these as separate domains with varying degrees of detail, rigor, and results. LDSE proposes consolidating these domains for a comprehensive, cohesive, and consistent approach to address system loss. This paper establishes interrelationships among the LDSE domains to harmonize role, fit, function, and impact among the domains focusing on sustaining value-delivery.

■ **KEYWORDS:** Loss-driven systems engineering, risk management, safety, security, agility, resistance, resilience, reliability, sustainability, survivability.

INTRODUCTION

To achieve *value delivery*, a system performs functions to produce *desired results*. To *sustain* value delivery while undergoing adversity, the loss-driven systems engineering (LDSE) domains contribute to system *viability* and *relevance*. LDSE describes an approach to address all forms of loss. Initially, LDSE domains include reliability (consistency), sustainability (renewable, waste management), survivability (continued existence), risk management (loss probability), resistance (retain desired status), resilience (regain desired status), agility (dynamic adaptation), safety (accidental loss), and security (malicious loss).

Viable and Relevant

Viable is capable of working successfully; being effective, efficient, and elegant. For example, we want our clean water supply to remain consistent, our food supply to remain plentiful, and the plane within which we fly to remain airborne until we reach our destination. *Relevant* is appropriate to current interest, or *current order* conformance. Absent of any adversity, the current order may change thus defining new desires. To remain *relevant*, a system may need to adjust the value it delivers and delivery method. LDSE helps ensure viability and relevance.

Context

Expressing a system's meaning and value, and expressing what constitutes loss and the loss degrees may vary according to *context*. For example, a commercial airplane is aluminum (structure), its function is flying (behavior), and its purpose is transport people and cargo (value-delivery). LDSE provides the lexicon and method to consistently and cohesively express loss, the loss degree, and how to address loss in various contexts such as structure, behavior, content, and value-delivery.

ELABORATING ON SYSTEM-OF-INTEREST LOSS

An action sequence has a chronology of results: *impact, effect*, and *consequence*. *Impact* is one object forcibly contacting another. *Effect* is a first-order result of contact. *Consequence* is the importance or relevance. For example, the pool stick strikes the cue ball (**impact**) which moves from its current location and knocks the eight ball into the side pocket (**effect**) which wins the game (**consequence**). Impact may be a **literal** contact or a **virtual** contact. The former is some hard contact (physical) where the latter is soft (psychological or cyberspace). A cyberspace attack includes bit flows (electrons) causing a virtual impact. An impact result may also either be virtual or literal. The former includes data exfiltration (confidentiality loss), data modification (integrity loss), or data destruction (availability loss). In cyber-physical systems, malicious electron manipulation may cause a physical explosion resulting in loss of property or life.

The system of interest (SoI) may suffer a *direct* or *indirect* loss from a recent encounter (impact), a recent change resulting from an encounter (effect), or an implication from its inability to produce desired results

Table 1. LDSE Domain Descriptions

Domain	Description / Comments
Reliability	Consistency for system characteristics; dependency.
Sustainability	Resource management, environment management, waste management, using renewable resources versus depletable resources.
Survivability	Continue to exist; remain compatible with the current order.
Risk management	Predicts the loss probability. Related to all LDSE aspects.
Resistance	Retain some desired status for system characteristics.
Resilience	Regain some desired status for system characteristics.
Agility	Dynamic adaptation; adaptable processes (development), adaptable solutions (systems), and adaptable workflows (operations).
Safety	Addresses accidental loss (not exclusively).
Security	Addresses malicious loss (not exclusively).

(consequence). Distinguishing loss nuances is important when considering *system assurance* (SoI focus) and *mission assurance* (focus on the SoI's containing whole or that which motivates the need for the SoI)) such as tactical versus strategic impact, effect, and consequence.

Impact types include:

- **Disclose**: losing intellectual property or other sensitive information negatively affecting competitive posture
- **Modify**: change to one or more system characteristics
- **Loss of X**: X∈ (overall system functionality, system access, system); system does not work at all, losing virtual system access, losing physical system access, or system destruction
- **Theft/loss**: possession loss either via malicious or accidental act
- **Misled**: suffering from deceit; conclude a thought or perform an action based on falsehood
- **Loss of effectiveness**: cannot perform intended purpose; system still active but cannot produce one or more intended results
- **Compliance driver violation**: system is or acts in some manner incompatible with legal authority, regulation, policy, or some other authoritative requirement
- **Deplete**: misdirect resources or consume resources unnecessarily or without authority; use up a resource, produce excessive waste, incur unnecessary cost
- **Deniable**: lack of accountability
- **Defile**: spoil the environment

Impact degrees include:

- **Destroy**: end the SoI ability to produce desired results
- **Disrupt**: temporarily incapacitate the SoI ability to produce desired results
- **Degrade**: deteriorate the SoI ability to produce desired results
- **Deny**: block access (physical); claim non-performance (opposite of non-repudiation)
- **Distort**: modify desired form (physical or virtual (data, information))
- **Deceive**: cause the SoI to perceive and thus respond to something not true thus having it produce desired results under false pretenses
- **Dated**: the SoI does not provide the features and functions available from newer alternatives; or, the SoI does not fulfill current stakeholder desires

The effect and consequence degree depends on context. Abstract effect and consequence degrees are *low, medium*, and *high* with many nuances such as annoyance, distraction, disturbance, degradation, delay, damage, disabling, destruction, or devastation. The impact implications are difficult to discern with a high degree of accuracy and certainty. Often, what seems like a trivial impact has tremendous consequences, as Benjamin Franklin said "the kingdom was lost… and all for the want of a horseshoe-nail." The impact may be temporary loss of use to a production database, the effect may be a short product shipment delay, but the consequence is a devastating market share loss due to earlier product availability from the competition.

LOSS-DRIVEN DOMAINS

Every engineered system has a purpose to fulfill its mission such as satisfying stakeholder desires. SoI efficacy is its capacity to fulfill its mission. LDSE provides for features and functions to safeguard the SoI, preserve its efficacy, and enable the SoI to fulfill its mission. Table 1 describes the current set of LDSE domains.

Reliability describes a system or component's ability to function under stated conditions for a specified period (IEEE 1990). Reliability as a measure is a failure probability. Concepts related to reliability include consistency, repeatability, durability, dependability, trustworthy, reproducibility, and lacking unintended variation. Reliability engineering includes design features helping the engineered system provide consistent and repeatable results.

Sustainable design, as defined by the US General Services Administration website, seeks to reduce negative impacts on the environment. Sustainability engineering designs or operates a system so they use energy and resources at a rate not compromising the natural environment or future generation ability to meet their own needs (Vallero and Brasier 2008). Sustainability measures include maximizing renewable resource use and minimizing depletable resource use.

Survivability, defined by Dictionary.com, is the ability to continue in existence or use. System survivability is the system's ability to minimize a finite disturbance impacts on value delivery (Richards et al. 2007, slide 10). The system achieves survivability through either satisfying a minimally acceptable value delivery level during and after a finite disturbance, or reducing a disturbance's likelihood or magnitude (Richards et al 2007, slide 10). An *a posteriori* survivability measure is survival rate. An indirect measure is on survivability contributors (fault-tolerance) and inferring a survivability level. An *a priori* survivability measure is the degree to which it is compatible with the current order.

Risk management predicts the loss probability (occurrence) and the loss degree (severity) across all system characteristic aspects. Loss may be real (physical) or virtual (data). There may be asset access loss, asset use loss, or asset loss. The *risk posture* captures stakeholder loss tolerance (*risk tolerance*).

Many notable engineers advocate for *proactive resilience*. "Resilience Engineering looks for ways to enhance the ability of organizations to monitor and revise risk models, to create processes that are robust yet flexible, and to use resources proactively in the face of disruptions (Dekker et al 2008)." "In a world of finite resources, of irreducible uncertainty, and of multiple conflicting goals, safety is created through proactive resilient processes rather than through reactive barriers and defenses (Woods and Hollnagel 2006)." LDSE captures the ***proactive*** (before something occurs), ***reactive*** (after something occurs), ***active*** (dynamic adjustment), and ***passive*** (static) spirit across *resistance* and *resilience* concepts.

A system is **resistant** if it produces desired results at or above a minimal efficiency threshold while *preventing the effects of an adversity*; resistance *retains* desired state, function, resources, environment, content, and value-delivery. Resistance enables the SoI to fight through the attack by preventing adverse effect(s). Prevention may *avoid* or *withstand*. There may be *active* resistance or *passive* resistance; when under missile attack, a military airplane may maneuver out of the way and deploy anti-missile devices, both are active avoidance. The airplane's fuselage may resist flak penetration from anti-aircraft fire, a passive resistance or withstand.

A system is **resilient** if it produces desired results at or above a minimal efficiency threshold while *undergoing the effects of an adversity*; resilience *regains* desired state, function, resources, environment, content, and value-delivery (note: *regain* does not necessarily mean return to original). Resilience enables the SoI to fight through the attack by dealing with an adverse effect via *withstand* or *recover*. Withstand minimizes the adversity effects or contains the adverse effect. Recover is to achieve value-delivery even if doing so with alternative means and performing at diminished efficiency.

Agility implies dynamic adaptation versus a static adaptation where the latter includes fault-tolerance in redundancies; if the primary hydraulic system fails, the system uses the built-in secondary hydraulic system. If the secondary hydraulic system fails and we somehow install a cable system on-the-fly to maintain control, this is dynamic adaptation or *agile*. An agile-system or an agile-workflow adapts to sustain value-delivery in predictable and unpredictable change (Dove 2014). This implies the ability to change SoI characteristics such as structure, state, function, or resource consumption.

To be ***safe***, according to Merriam-Webster's online dictionary, is to be free from harm or risk; or to be unhurt. To be ***secure***, according to Merriam-Webster's online dictionary, is to be free from danger or free from risk of loss. Engineers often use the terms interchangeably though we intuitively have distinctions in mind. For example, we think of a seatbelt more in safety terms and a door lock more in security terms. For harmonizing LDSE domains, safety predominantly addresses *accidental loss* and security predominantly addresses *malicious loss*.

Table 2: Thoughts Toward LDSE Principles (Table 2 *continues on next page*)

Domain	Notional Principles
Context	Express meaning and value in a proposition context (Frege 1884) Context shapes expressing stakeholder desired results Context shapes expressing loss and loss tolerance
Reliability	Continuous monitoring: ongoing observation to raise awareness Failure resistant: avoid SoI failure Accuracy: continual validation (do the right thing), continual verification (do the thing right) Consistency: features and functions producing repeatable results Dependability: features and functions produce desired results when needed
Sustainability	Resource management: minimize resource consumption; minimize depletable resource use, maximize renewable resource use Earth: minimize physical waste; minimize contamination Air: minimize air emissions Water: minimize waste release to water Mind: minimize cognitive workload; minimize psychological trauma
Survivability	Current order: remain compatible with the current order Maximize viability Maximize relevance
Risk Management	Formalize stakeholder risk tolerance Maximize organizational efficacy; minimize threat efficacy Minimize loss (negative risk side); maximize opportunity (positive risk side) Accept risk when benefits are greater than cost; accept only necessary risk Ignoring risk implicitly accepts risk, conscious choice above omission by oversight Manage uncertainty; intelligent decision-making considers risk Risk management facilitates continual adaptation Continual adaptation requires continual risk management
Resistance	Retain effectiveness, efficiency, elegance, efficacy Retain state, function, resource, content, environment, value delivery
Resilience	Regain effectiveness, efficiency, elegance, efficacy Regain state, function, resource, content, environment, value delivery

LDSE DOMAIN HARMONIZATION

Harmony is an emergent order;

Table 2: Thoughts Toward LDSE Principles (continued)

Domain	Notional Principles
Agile	Adapt to predictable change Adapt to unpredictable change Adapt predictably (deterministic); playbooks Adapt unpredictably (non-deterministic) or flexibly; emergent behavior Actions include planned and emergent dynamic composition to perform the following (Willett *et al.* 2016): • **Monitor**: ongoing observation with intent to raise anomaly awareness (anomaly is deviation from expected) • **Detect**: become aware of anomaly • **Characterize**: categorize anomaly for faster processing • **Notify**: inform most relevant support tier for the anomaly • **Triage**: prioritize addressing anomalies • **Escalate**: inform most relevant specialization group • **Isolate**: contain adversity or adverse effects • **Restore**: alternative means to produce desired results; regain value-delivery • **Root cause analysis**: distinguish symptom from problem • **Recover**: resolve the problem; regain loss • **Feedback**: systemic adjustment due to lessons learned
Security (Willett 2008)	**Confidentiality**: ensure only authorized disclosure **Integrity**: ensure only authorized modification **Availability**: ensure ready for use; ensure no service denial **Possession**: ensure physical retention; ensure no physical loss or theft **Authenticity**: ensure conformance with reality; ensure no deceit **Utility**: ensure fit for purpose **Privacy**: right no observation, the right to forgetting **Non-repudiation**: ensure accountability for actions; ensure non-deniability **Authorized use**: ensure only authorized [cost-incurring] service use
Safety	Minimize unintentional harm; minimize intentional harm Sacrifice property before life Sacrifice non-human life before human life Safeguard SoI's state, function, resource, content, environment, value delivery Safeguard other SoI's Hierarchy on *harm* degree choices in preference order (priority): • **Avoid** rather than deflect (no contact) • **Deflect** rather than damage (light contact, redirecting force) • **Damage** rather than destroy (medium contact) • **Destroy** rather than kill (hard contact) • **Kill** only as a last resort

"harmony resides in a reality to be created each and every time (Sundararajan 2013, p.2)." Harmony is not uniformity; rather, "harmony is a relational term which entails diversity and difference (Sundararajan 2013, p.2)." Harmony is a holistic perception, an overall sense of things rather than focusing on any particular thing (Lu 2004). Harmony is a dynamic equilibrium (The Doctrine of the Mean 1971). The following narrative harmonizes LDSE domains with respect to a system providing value-delivery. From this narrative we can *begin* discerning LDSE domain roles, fits, functions, and impacts on each other, the system to which they apply, and establish a framework to discern their holistic relationships, find their dynamic equilibrium, and their emergent order.

Expressing value varies among stakeholders; there are differences in *stakeholder currency*; stakeholder currency to a politician is votes, a scientist is *knowledge*, a general is *lives*, and a banker is *money*. The system's main *goal* is providing value-delivery in stakeholder relevant terms. Two macro-level system sub-goals are remaining viable and relevant. Measurable *objectives* sustaining viability and relevance are for the system to be effective, efficient, and elegant. Measurable sub-objectives to these include **reliability** (consistent, dependable), **sustainability** (renewable), and **survivability** (compatible with the current order); and, there are other sub-objectives at this layer (future discussion).

Methods include tactics, techniques, and procedures (TTP's) to achieve the objectives. **Risk management** is a method to predict the *loss probability* and the *loss severity* to help the stakeholders determine their risk tolerance in turn driving what to do about the risk. An adversity poses a loss risk to one or more system characteristics. If the loss occurs, it occurs to some degree of adverse effect. **Resistance** methods attempt to *retain* system characteristics (avoid or withstand adverse effects). **Resilience** methods attempt to *regain* system characteristics (withstand or recover from adverse effects). Resistance and resilience forms vary among **agile** (dynamic, composable), static (passive, playbook), proactive (preemptive), and reactive (responsive).

Methods invoke products and services (solutions) as part of their processes. With respect to LDSE, these solutions are safeguards addressing **safety** (accidental loss) and **security** (malicious loss). Safety and security products and services provide the solution space helping ensure viability and relevance so the system continues to provide value-delivery.

From this narrative, we see reliability, sustainability, and survivability as measurable objectives. Risk management, resistance, resilience, and agile are methods to achieve the objectives. Safety and security provide solutions the methods invoke. The LDSE domains are necessary but not sufficient to sustain value-delivery. LDSE is part of a larger construct (future discussion) for the system to achieve and sustain value delivery.

Toward LDSE Principles

Table 2 provides thoughts toward LDSE principles; incomplete and for discussion.

Thoughts toward refining LDSE principles include resilience types and ethics. Resilience types:

- **Innate**: born with; applies to natural living systems; not contrived by humans
- **Inherent**: essential, intrinsic; applies to non-living systems, contrived by humans; resilience emerges via the normal SoI features and functions
- **Planned**: a contrived SoI part, intentional design such as redundant component. Redundant feature or function. Invoke something known (playbook)
- **Emergent**: an agile behavior invokes planned features and functions in a manner producing new behaviors; composing a new ability producing desired results

Harm may be necessary for strategic success: the pawn to save the king, and win the battle or sacrifice the data server to learn more about adversary strategy. Intentional harm will at times be necessary to resolve moral dilemmas; autonomous vehicle *must* choose to hit four school children on the left, a woman pushing a baby carriage on the right, or crash into the barrier straight ahead thus causing harm to itself and its contents. Accepting this takes us down the path that a SoI perpetrating some harm is necessary. Now comes the extremely difficult question to the acceptable degree of harm and in what form the harm remains acceptable. This will vary according to context such as cultural differences in morality, and acceptable behavior and consequences. The final version of safety principles must capture these concepts.

CONCLUSION

LDSE domains work in harmony providing a comprehensive approach to identify and integrate loss-driven requirements in a holistic solution design addressing all system state, behavior, resources, content, environment, and value characteristics. LDSE facilitates producing the risk posture reflecting stakeholder loss tolerance. LDSE complements opportunity-driven systems engineering as iterative methods to sustain viability and relevance to achieve the main value-delivery goal (Willett 2020). ■

REFERENCES

- Confucius. 1971. *Confucian Analects: The Great Learning, and the Doctrine of the Mean*. Edited by J. Legge. New York, US-NY: Dover Publications.
- Dekker, S. A., E. Hollnagel, D. D. Wood, and R. Cook. 2008. "Resilience Engineering: New Directions for Measuring and Maintaining Safety in Complex Systems." Technical Report, School of Aviation (Lund, SE).
- Dove, R., and R. LaBarge. 2014. "Fundamentals of Agile Systems Engineering—Part 1 and 2." Paper presented at the 24th Annual International Symposium of INCOSE, Las Vegas, US-NV, 30 June-3 July.
- Frege, G. 1980. *The Foundations of Arithmetic*. Evanston, US-IL: Northwestern University Press.
- Institute of Electrical and Electroncs Engineers (IEEE) Computer Society, Standards Coordinating Committee. 1990. *IEEE Standard Computer Dictionary: A compilation of IEEE Standard Computer Glossaries*. New York, US-NY: Institute of Electrical and Electronics Engineers.
- Lu , R. R. 2004. *Zhung-guo gu-dai xiang-dui guan-xi si-wei tan-tao* [Investigations of the idea of relativity in ancient China]. Taipei, CN: Shang ding wen hua.
- Richards, M. G., D. H. Rhodes, D. E. Hastings, and A. L. Weigel. 2009. "Defining Survivability for Engineering Systems." Paper presented at the annual Conference on Systems Engineering Research, Hoboken, US-NJ, March.
- Sundararajan, L. 2013. "The Chinese Notions of Harmony, with Special Focus on Implications for Cross Cultural and Global Psychology." *The Humanistic Psychologist* 41: 1–10.
- Vallero, D., and C. Brasier. 2008. *Sustainable Design: The Science of Sustainability and Green Engineering*. Hoboken, US-NJ: Wiley.
- Willett, K. D. 2008. Information Assurance Architecture. Boston, US-MA: Auerbach Publications.
- ———. 2020. "Systems Engineering the Conditions of the Possibility." Paper presented at the 30th Annual International Symposium of INCOSE, virtual event, 20-22 July.
- Willett, K. D., R. Dove, R. Cloutier, and M. Blackburn. 2016. "On System Dynamics Modeling of Human-Intensive Workflow Improvement—Case Study in Cybersecurity Adaptive Knowledge Encoding." Paper presented at the 26th Annual International Symposium of INCOSE, Edinburgh, GB-SCT.
- Woods, D. D. 2006. Resilience Engineering: Concepts and Precepts. Edited by E. Hollnagel. Farnham, GB: Ashgate Publishing.

ABOUT THE AUTHOR

[Editor: Author biography was current when the paper was initially published in 2020.]

Dr. Keith D. Willett is a senior strategist and enterprise security architect for the United States Department of Defense with over 35 years' experience in technology. He is a co-chair for the INCOSE working groups on *Systems Security Engineering* and *Agile Systems and Agile Systems Engineering*; plus, an active member in working groups for *Resilient Systems* and *Systems Science*. He is the lead for INCOSE's *Future of Systems Engineering* (FuSE) agility project with intent toward systems engineering methods developing and operating inherently adaptable systems.

Versatile Test Reactor Open Digital Engineering Ecosystem

Christopher Ritter,a* christopher.ritter@inl.gov; **Jeren Browning,a** jeren.browning@inl.gov; **Peter Suyderhoud,a** peter.suyderhoud@inl.gov; **Ross Hays,a** ross.hays@inl.gov; **AnnMarie Marshall,a** ann.marshall@inl.gov; **Kevin Han,c** kevin_han@ncsu.edu; **Taylor Ashbocker,a** taylor.ashbocker@inl.gov; **John Darrington,b** john.darrington@inl.gov and **Lee Nelson,a** lee.nelson@inl.gov

[a] *Idaho National Laboratory, Department of Digital Systems & Engineering, Versatile Test Reactor Program, Idaho Falls, ID 83401*
[b] *Pivotal Solutions, 3 Grace Avenue, Suite 162 Great Neck, NY 11021*
[c] *North Carolina State University, Department of Civil and Environmental Engineering, and Department of Computer Science, University of Illinois at Urbana-Champaign, 205 N. Mathews Ave., Urbana, IL, 61801*

■ **ABSTRACT**
Modern design of nuclear facilities represents unique challenges: enabling the design of complex advanced concepts, supporting geographically dispersed teams, and supporting first-of-a-kind system development. Errors made early in design can introduce silent errors. These errors can cascade causing unknown risk of complex engineering programs. The Versatile Test Reactor (VTR) Program uses digital-engineering principles for design, procurement, construction, and operation to reduce risk and improve efficiencies. Digital engineering is an integrated, model-based approach which connects proven digital tools such as building information management (BIM), project controls, and systems-engineering software tools into a cohesive environment.

The VTR team hypothesizes using these principals can lead to similar risk and cost reductions and schedule efficiencies observed in other engineering industries. This research investigates the use of a digital engineering ecosystem in the design of a 300-MWt sodium-cooled fast reactor. This ecosystem was deployed to over 200 engineers and used to deliver the conceptual design of the VTR. We conclude that initial results show significant reductions in user latency (1000x at peak use), the possibility of direct finite-element-analysis (FEA) integrations to computer-aided design (CAD) tools, and nuclear reactor system design descriptions (SDDs) that we can fully link throughout design in data-driven requirements-management software. These early results led to the VTR maintaining milestone performance during the COVID-19 pandemic.

■ **KEYWORDS:** digital engineering, digital twin, digital thread, versatile test reactor, Idaho National Laboratory

I. INTRODUCTION

The VTR program has referenced the Department of Defense (DoD) digital-engineering strategy (https://fas.org/man/eprint/digeng-2018.pdf) as a key upper-level strategy for VTR implementation. The DoD breaks digital engineering into five key functions:

1. Transform to end-to-end digital representation through models.
2. Structure information into an authoritative source-of-truth.
3. Integrate technological innovation and technology advancements.
4. Provide an information technology (IT) infrastructure and environment to support the process.
5. Change the culture and workforce through change management, communications, training, and strategy.

This digital transformation approach proved to reduce the schedule by approximately 10 years on new DoD aircraft (https://www.foxnews.com/tech/air-force-flies-6th-gen-stealth-fighter-super-fast-with-digital-engineering), increase performance by 25% in construction (https://www.mortenson.com/vdc/study), and avoid $1 billion in cost through advanced digital twins (https://www.ge.com/digital/blog/industrial-digital-twins-real-products-driving-1b-loss-avoidance). To deliver on this strategy and attempt to receive similar benefits for an advanced Gen-IV nuclear reactor, the VTR Program has decomposed this strategy into key implementation-focus areas (see Figure 1. Digital Engineering Strategy).

The VTR program has implemented each function of the DoD strategy through the following five key areas:

1. Data-Driven Tools: For data capture in design engineers deploy BIM, project portfolio management, and data systems-engineering tools.
2. Digital Thread: An open-source data model (ontological)-focused software platform to connect nuclear data.
3. Technological Innovation: Partnership with universities to provide integration with FEA and artificial intelligence/machine learning (AI/ML).

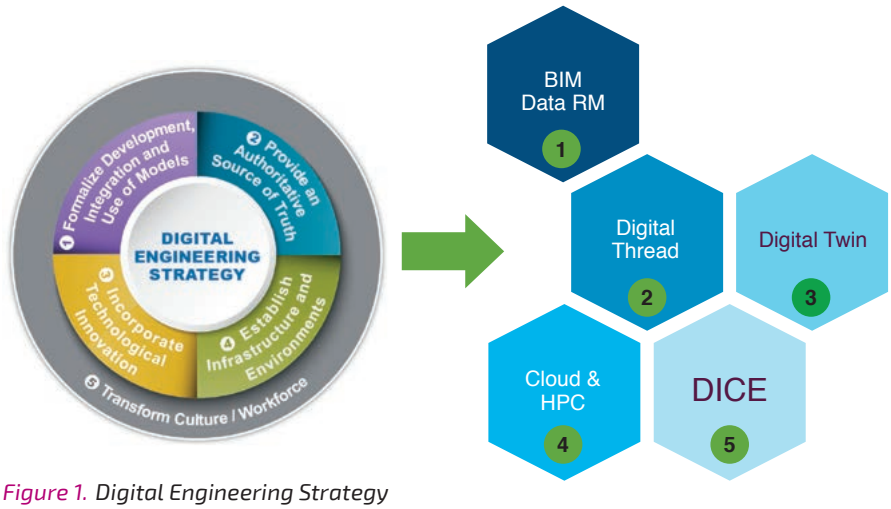

Figure 1. Digital Engineering Strategy

4. Cloud and High-Performance Computing (HPC): An environment and infrastructure to host data-driven tools throughout the VTR team.
5. Digital Innovation Center of Excellence (DICE): A community of practice to transform laboratory culture.

Idaho National Laboratory (INL) provides the tools, digital thread, and computing for the VTR program with laboratory and industry partners, contractors, and supporting universities interfacing with this environment.

II. DIGITAL ENGINEERING FUNCTIONS

Data-Driven Tools

The VTR Program uses models and data-driven tools across the engineering process to facilitate design. Currently, key tools for the program are BIM for two- and three-dimensional (2D and 3D) CAD/equipment, systems-engineering, and project portfolio-management software.

The AVEVA solution manages BIM for the capture of engineering design data. BIM has several capabilities, including the capture of 2D piping and instrumentation diagrams (P&ID) and 3D mechanical, civil, and structural diagrams. AVEVA's suite of software includes a web-based collaboration platform, AVEVA Net, to enable project stakeholders to collaborate and review digital-asset information across the lifecycle in real time. AVEVA defines a data model, with mapping to the digital thread authoritative source of truth.

The team deployed the IBM Jazz engineering life-cycle management (ELM) software for the development of systems-engineering artifacts. The Jazz ELM includes a suite of integrated tools to manage requirements, define test cases (through verification and validation), model physical and logical architectures, and manage system changes. These tools are web-based and require no software installation on client machines. The Jazz ELM uses an open standard—open services for life-cycle collaboration (OSLC)—to enable links across the systems-engineering tool suite. These links enable a named association (satisfied by) between systems-engineering artifacts, to currently 5225 requirements artifacts are currently allocated, including all system-level requirements. This allows connection within the Jazz software system itself. To connect to the authoritative source of truth engineers developed a software adapter and compatible ontological linkage.

The VTR Program uses the Oracle Primavera P6 enterprise project portfolio-management (EPPM) software. The P6 EPPM platform is used to develop the integrated master schedule, to integrate planned and actual costs, and to define overall project health. P6 provides a web-services application program interface (API) which the team integrated as an adapter to the authoritative source of truth.

Digital Thread

The authoritative source of truth consists of two key technologies: a database to centrally store all information across the program and a formalized data model or ontology to organize the information. VTR uses an open-architecture data framework, Deep Lynx, originally developed on the VTR to centrally integrate and store the source-of-truth data. All data in Deep Lynx map to the open-source Data Integration Aggregated Model and Ontology for Nuclear Deployment (DIAMOND). DIAMOND contains classes, their properties, and relationships that represent various concepts and data within the nuclear domain (Figure 2–Deep Lynx Architecture).

Deep Lynx is a data warehouse, unique in that it stores its data in a graph-like format which maps to a user-provided ontology or taxonomy. Currently, engineering teams operate in siloed tools and disparate teams, where connections across design, procurement, construction, and operating systems undergo manual translation or over brittle point-to-point integrations. The manual nature of data exchange increases the risk of silent errors in the reactor design, with each silent error cascading across the

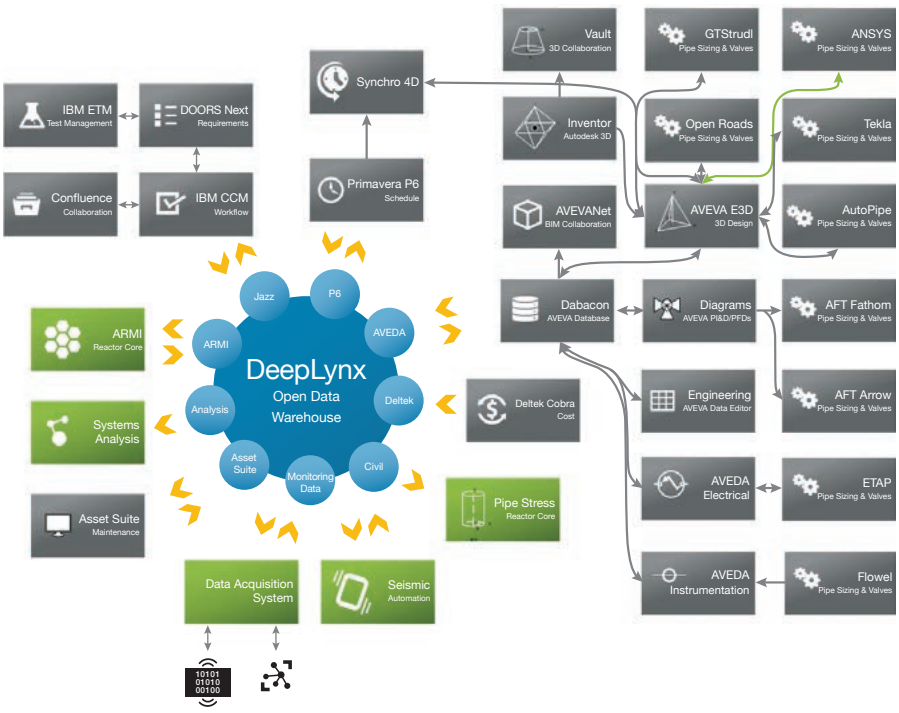

Figure 2. Deep Lynx architecture

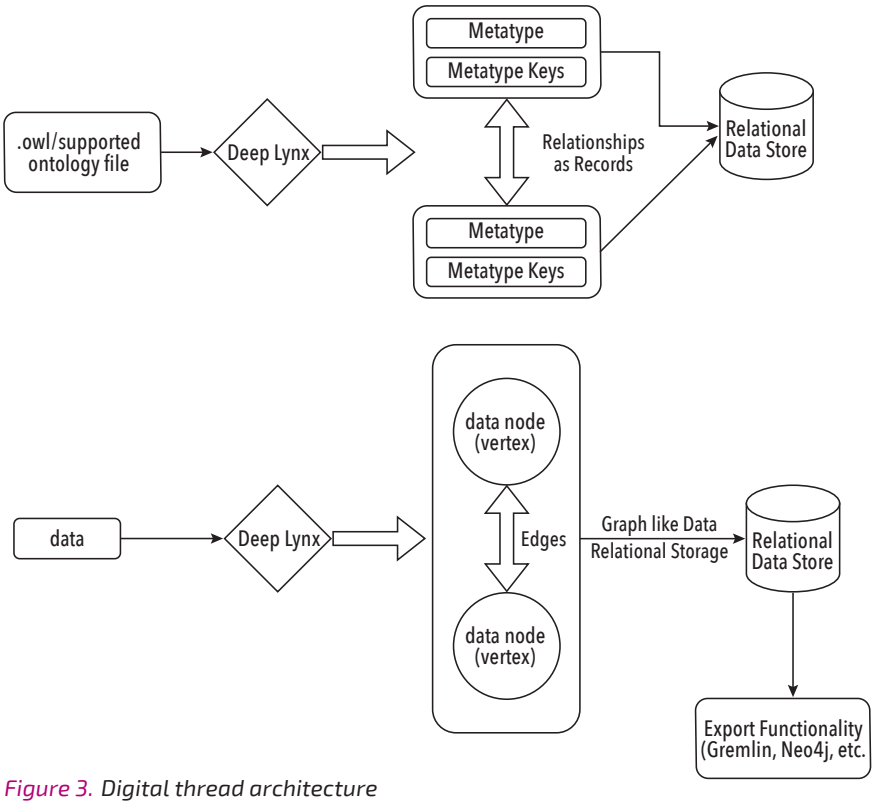

Figure 3. Digital thread architecture

design. Deep Lynx allows for an integrated platform during design and operations of megaprojects. Deep Lynx allows the user to safely and accurately aggregate data from various programs and data sources, store those data in a user-provided organization such as the DIAMOND ontology, and then generate reports and visualize ingested data.

The Deep Lynx software system can map to a provided W3C Web Ontology Language (OWL) file to instantiate a container with an ontology. A container separates between different sites or areas of work. Containers also allow Deep Lynx to manage user interaction with certain data, enforcing security controls at this level. Within a container, information from the ontology classes is persisted as metatypes for classes, metatype keys for class properties, and metatype relationships for relations between classes (see Figure 3–Digital Thread Architecture). This controls the database schema within a container. Data themselves persist as a set of nodes and vertices (edges) that map to the equivalent metatype and metatype relationship accordingly. Users can view these data can relationally, through GraphQL, or export them to view in any Gremlin-compatible graph database.

For an import's data to insert into Deep Lynx, each dataset must be associated with a Type Mapping. A type mapping defines how the incoming data maps to the defined ontology or taxonomy within Deep Lynx. Type mappings apply constantly to matching data automatically and require no user intervention apart from teaching the system through creation of new type mappings when there are new data structures submitted. Once an import's data completes mapping, the transformation and insertion process will begin automatically.

Technological Innovation

Managing design changes during construction for a nuclear energy facility requires the involvement of multiple design engineers to effectively communicate and collaborate changes that branch from a new change. To provide effective management and streamlined communication of design changes, engineers develop an automated bi-directional conversion of facility structure and piping systems. This interoperability solution between building information models (BIM—for example: Autodesk Revit, AVEVA E3D, and more) and physics-based model (i.e., finite-element model in advanced modeling and simulation [M&S] tools, such as ANSYS or Abaqus) allows quick conversion of as-built BIM models into analytical models. Any design changes during construction—hanger support location for a piping system due to site constraint—can reflect or show in an analytic model for quick assessment of risk. On the other hand, any new changes by designers in the analytical model— an increased number of hanger supports for the piping system—can be quickly translated into the drafting model (BIM) that the contractor use for construction. The process diagram in Figure 4–Versioning of design changes through bi-directional conversion of BIM and M&S models illustrates this quick assessment through bi-directional BIM-to-M&S conversion and its integration with Deep Lynx for versioning of design changes.

Deep Lynx, as a data warehouse, stores these BIM and analytical models and their versioning information through a simple relational database/structured query language (SQL), allowing pulling and pushing of the latest design changes. It can also allow tracking of changes by querying models at different times. The versioning as part of Deep Lynx has a synergistic opportunity for better managing building

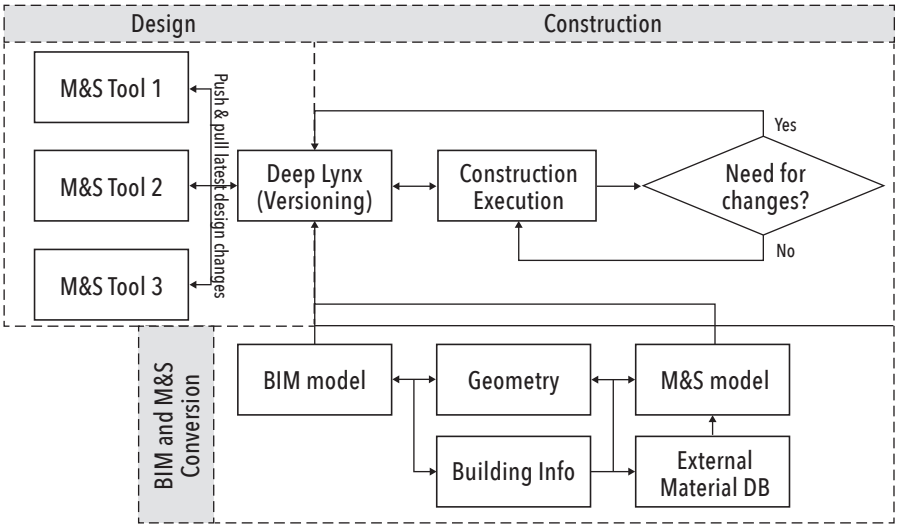

Figure 4. Versioning of design changes through bi-directional conversion of BIM and M&S models

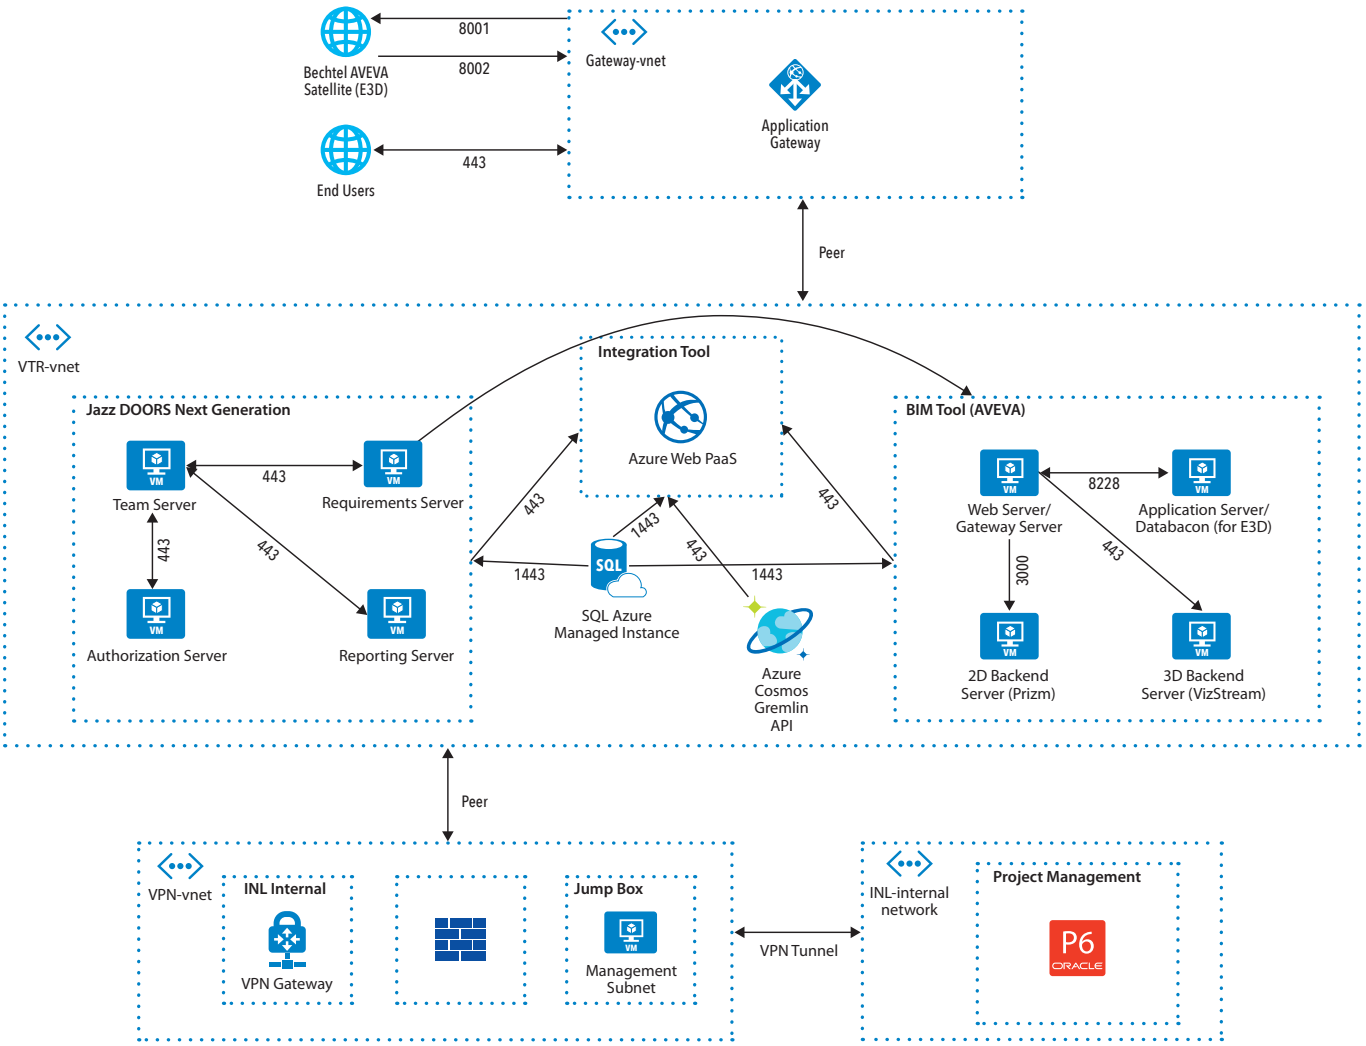

Figure 5. VTR network diagram

information throughout the design and construction phases. It can function as a central data hub for accessing design and construction models at different times for accurate access of information, minimizing the accumulation of human errors while managing the ever-increasing list of design changes. Last, keeping the most up-to-date building information through construction would be beneficial for the operational phase, especially given the digital-twin concept of operational monitoring and simulation.

Cloud and HPC

The Microsoft Azure for Government Cloud hosts the VTR digital-engineering software tools and ecosystems . The U.S. Department of Energy (DOE), Idaho Operations Office (ID) approved this cloud to host data up to official use only (OUO)/ export-controlled information (ECI).

The INL-managed cloud environment uses virtual networks to contain and separate different programs. The architecture includes a hub virtual network (VNET) which connects to the VTR internal spoke VNET. This hub VNET is then peered to an external spoke, which is peered to the internet. The INL Information Management networking teams manage configuration of this network; changes to the VTR spoke must receive approval though a change-request board.

An overview of the current VTR digital engineering network architecture follows as Figure 5 –VTR Network Diagram:

This environment dramatically reduced peak latency of the requirements-management tool, Dynamic Object-Oriented Requirements System—Next Generations (DOORS-NG), during peak use. Latency between the database server and application was 3,000 ms during peak times before cloud migration, with latency reduced to 2–3 ms duration once cloud migration was complete. This represents a 1000× decrease in peak latency. Additionally, end-user latency was reduced from to approximately 6 ms to 60 ms, depending on geographic location.

An automation is in development between Deep Lynx's cloud and the INL's HPC environments. This automation uses a queue model to allow computationally expensive models (multiphysics) to send to HPC for processing but uses the cloud to distribute the results geographically among the team.

DICE: Culture and Workforce Transformation

DICE serves as a virtual center to formalize and coordinate digital-engineering, digital-twinning, and digital-transformation activities across next-generation energy systems (https://dice.inl.gov). DICE is a laboratory-wide service that provides to leadership a strategy focused on energy-system needs, recognition through research accomplishments, coordination to share community best practices across the laboratory, outreach to universities, industry partners, and other national laboratories, and enhancement training and education materials on digital

engineering and twinning. The VTR shares best practices with the DICE community and encourages process transformation to support digital engineering.

III. CONCLUSION

The digital-engineering strategy implemented for the VTR program led to sustained milestone performance throughout the COVID-19 pandemic. The ecosystem of digital tools allowed for contractors, laboratories, and universities to seamlessly switch to digital collaboration as the primary means of communication. To date, this strategy led to all system-level requirements to be persisted at the object level, latency reduced by 1000× during peak use, and proof of CAD-to-FEA integration automation. Additionally, changes to any of the more than 5000 requirements are automatically tracked and updated across the program, ensuring that potential silent errors are avoided. The VTR program's technologies have seen expanded use on other DOE nuclear-reactor programs, National Nuclear Security Administration (NNSA)-safeguards digital twins, and commercial nuclear programs. As the VTR program matures planned continued expansion of the Deep Lynx ecosystem and automation will realize cost reductions, schedule improvements, and engineering performance gains. ■

ABOUT THE AUTHORS

[Editor: Author biographies were current when the paper was initially published in 2022.]

Christopher Ritter is the director of the Digital Innovation Center of Excellence (DICE) and department manager of digital engineering at INL. Chris leads digital engineering across the Versatile Test Reactor, National Nuclear Security Administration Safeguards Digital Twin, and National Reactor Innovation Center (NRIC) programs. Before coming to INL, he was director of software development at SPEC Innovations and served as the chief architect of Innoslate, a popular model-based systems engineering tool.

Jeren Browning is the digital thread group lead for the digital engineering department at INL. His expertise includes digital thread and systems integration technologies, digital twin development, full stack development, and systems architecture. He currently acts as the principal investigator on an autonomous digital twin research effort.

Kevin Han is an assistant professor of civil engineering at North Carolina State University. He received a PhD in Civil Engineering and Master of Computer Science at the University of Illinois at Urbana-Champaign. His research focuses on improving construction performance monitoring and management through advances in visual sensing and building information modeling.

Peter Suyderhoud is the model-based systems engineering group lead in the digital engineering department at INL. His focus is applying systems engineering principles to large, complex energy facilities using modern digital engineering platforms. Previously he worked in the EPC industry on several novel Department of Energy facilities, including a first of a kind nuclear waste processing plant.

Lee Nelson is the director of acquisition and logistics on the VTR Program. Lee has over 30 years' experience in project engineering, project management, and chemical engineering. On the VTR, Lee established the digital engineering sub-project and currently advises on all aspects of practical engineering application. Recently, Lee served as the project engineer on the DOE TREAT program which delivered a year ahead of schedule and $20 million under budget.

John Darrington is a digital twin research scientist working in the digital engineering department at INL. Specializing in software engineering and design, John continuously works to introduce and manage new technology and strategies in the digital engineering world. John has worked with the digital engineering department for the last 3 years and most recently has spearheaded efforts on an internal data warehouse, Deep Lynx.

AnnMarie currently works as the management of configuration management in support the VTR program at INL. She participated in authoring and submitting their critical decision (CD) documentation and establishing and managing their CM Program. In addition, AnnMarie established, created, and managed VTR the project records repository. Prior to working at BEA, AnnMarie was the Manager for Configuration Management for New Plant Projects at GE-Hitachi Nuclear Energy (GEH) where she participated in multiple projects related to the ESBWR, ABWR, and PRISM product lines.

Ross Hays has been researching nuclear fuel cycle deployments, biomass feedstock logistics, and digital engineering applications at Idaho National Laboratory for the past seven years. Most recently he became the principal investigator for digital engineering integrations with the Natrium Advanced Reactor Demonstration Program Project.

Taylor Ashbocker is the principal cloud architect in INL's digital engineering department. He is responsible for managing cloud infrastructure while developing compliance and protection of cloud resources. Before joining INL in February 2019 he was a senior development operations engineer at First American Financial, managing critical business infrastructure in Amazon Web Services.

Systematic Identification and Analysis of Hazards for Automated Systems

Lina Putze, lina.putze@dlr.de; and **Eckard Böde,** eckard.boede@dlr.de

■ **ABSTRACT**
The introduction of automation into technical systems promises many benefits, including performance increase, improved resource economy, and fewer harmful accidents. In particular, in the automotive sector, automated driving is seen as one key element in Vision Zero by eliminating common accident causes such as driving under the influence, reckless behavior, or distracted drivers. However, this is contrasted by new failure modes and hazards from the latest technologies. In this article, we address the problems of finding common sources of criticality for specific application classes and identifying and quantitatively assessing new sources of harm within particular automated driving systems.
■ **KEYWORDS:** automated driving; hazard analysis; risk assessment; criticality; SOTIF; scenario identification; open context

INTRODUCTION – THE PROBLEMS OF IDENTIFYING RISKS FOR AUTOMATED DRIVING

Accidents due to speeding, distraction, or driving under the influence of alcohol – human misbehavior, intended or unintended, is an important factor in accident statistics. Self-driving vehicles are supposed to increase road safety by reducing the "human" risk factor. Although hazards associated with humans, like a collision due to a distracted driver, might be mitigated, the new technologies come with unknown risks and failure modes. The research topic, *Automation Risks*, focuses on identifying and assessing hazards and scenarios likely to trigger critical situations in the interaction of automated driving systems with their environment. In this article, we will focus on investigating automated driving systems since the methods presented have been developed in close collaboration with partners from the automotive industry. Nonetheless, we are actively adapting to other domains, like the maritime industry.

The safety of road vehicles is a well-known issue in the automotive industry. Due to the rising complexity of interacting safety-critical components, even conventional driving systems need to undergo a systematic safety process corresponding to ISO26262:2018 (ISO2018]). To keep development costs and efforts to a minimum, it is essential to include safety considerations from the beginning of the concept phase and throughout the entire development process because integrating changes in the system during early design phases is significantly easier. Knowledge about the common sources of criticality, for example, from accident databases, is an essential prerequisite for these first safety considerations. Moreover, a comprehensive safety concept requires a systematic identification and analysis of system-specific sources of harm. In the automotive domain, several methods exist for a so-called hazard and risk analysis (HARA), which is well-established in developing road vehicles.

Common hazard and risk analysis methods emphasize functional safety, which focuses on identifying and mitigating possible hazards caused by malfunctioning behavior of safety-related electrical and electronic systems. Assistance systems currently on the market, like adaptive cruise control, lane-keeping assistance, and combinations thereof, still require a human driver to monitor the vehicle and the environment and intervene when necessary. Nonetheless, many of those systems already take over parts of the driving tasks by providing braking, acceleration, and steering support while relying on sensor data that captures the internal and external environment. This comes with new potential sources of harm that take root in the system's specification. Let us consider an automatic emergency braking function (AEB). Despite the absence of faults and malfunctions, such hazards might occur due to incorrect interpretation of sensor input. For example, a poster on the roadside with a picture of a pedestrian crossing the road could be perceived as a natural person resulting in a breaking maneuver that could trigger a collision. This demonstrates that additional examination beyond the functional safety of the system is needed. We need to ensure that the system is robust concerning incorrect or unexpected sensor input, can comprehend situations correctly, and plans and acts responsibly based on these perceptions. These issues concerning the safety of the intended functionality (SOTIF) are addressed by ISO 21448:2022 (ISO 2022).

As assistance systems still have the driver as a redundant and immediately available fallback, such systems only

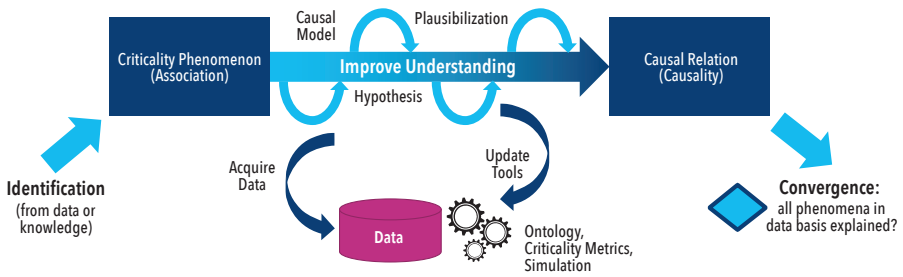

Figure 1. Basic concept of the criticality analysis

require evidence that the safety concept is fail-safe because the system does not provoke any additional risks, for example, by unintended interventions. In contrast to well-established systems, conditionally or highly automated driving functions like a traffic jam chauffeur temporarily release the driver from monitoring the environment for a certain time. This important step in the Levels of Driving Automation comes with additional safety difficulties as the abandonment of the driver as supervising instance involves the loss of a comprehensive and immediately available fallback. Therefore, it is necessary to prove that the system takes all the actions required to mitigate critical situations and that these actions are always carried out correctly and with the right timing: an operational-safe concept is required.

This is particularly problematic since automated driving systems driving on public roads face the challenge of operating safely in an open context. This arbitrarily complex, infinitely dimensional environment includes myriad factors that might lead to harm. Thus, it is infeasible to describe all relevant scenarios explicitly and specify the intended behavior. Moreover, hazards cannot be sufficiently reconstructed from existing real-world data. While there is extensive data for conventional driving systems, the challenges for automated driving systems differ from those for the human driver. For example, falling leaves in autumn are not generally a problem for the human eye, but if they hit the lens of a camera, object detection is not feasible anymore. Therefore, we cannot solely rely on data considering conventional systems and need extensive data that reflects the impact of automated systems on criticality.

To address these outlined issues, our research into *Automation Risks* is based on two main pillars: First, there is the criticality analysis which aims at finding common factors associated with criticality. Its focus is not on a specific system but on abstract application classes, such as the function of a highway chauffeur. Hence, the scope is in a pre-development phase where working groups comprising representatives from regulation authorities, standardization bodies, and industry define standard guidelines that every manufacturer of such a system must meet. In this setting, the criticality analysis will be a systematic approach to identify potential sources of criticality and specify a complete, well-defined set of criticality phenomena to be used as the basis for a homologation concept. Second, we work on a methodology that can be employed to perform a comprehensive hazard and risk analysis for specific highly automated systems that accompany the development process. This automation risks method aims to identify specific scenarios for further verification and validation and define safety goals as a basis for a fail-operational safety concept. The method intends to integrate functional safety (ISO 26262:2018 (ISO 2018)) and SOTIF (ISO 21448:2022 (ISO2022)) concerns.

STRUCTURING THE OPEN CONTEXT – CRITICALITY ANALYSIS

The first method we present is the criticality analysis. Its purpose is to investigate and structure the open context that constitutes the environment of automated vehicles. This includes not only the problem of identifying factors, parameters, and scenarios that have an essential impact on criticality but also abstracting these artifacts and mapping them on a finite set of criticality phenomena. This abstraction structures the criticality-inducing factors into comprehensive but manageable lists that can serve as a foundation for systematic verification and validation processes that enable a homologation for classes of automated systems. Furthermore, it helps to understand the underlying causalities to derive generic safety principles and mechanisms that avoid or mitigate the effects of critical situations.

Therefore, criticality analysis relies on a combined approach of expert-based and data-driven methods that precedes the design phase of specific systems. For example, it can be applied to urban traffic to set up a foundation for developing automated systems in this domain. In addition, it can support the operation and subsequent updates of corresponding systems in a DevOps process by continuously assessing changes in their domain. That might involve specific effects of amendments or enactments of laws and guidelines – a recent example would be the approval of e-scooters for German streets in 2019 – or even effects of climatic or societal changes. One of the fundamental principles of criticality analysis is that it does not only focus on the view of a single vehicle but also looks at the criticality of traffic. In this way, criticality analysis makes it possible to create generally accepted catalogs of criticality phenomena managed by regulation bodies and used by all manufacturers.

The basic approach of the criticality analysis is shown in Figure 1 and consists of three steps which we will present individually in the following.

1. *Identification and selection of criticality-triggering elements*: In the first step of the criticality analysis, candidates for criticality phenomena are selected for which a high correlation with a criticality increase is assumed. Expert knowledge, which is stored, for example, in the form of domain ontologies, test catalogs for vehicle approval, or accident databases, serves as a basis for the selection. Another source is data-driven approaches that systematically evaluate data from driving tests on test fields or in real traffic and data from specific computer simulations.
2. *Plausibilization and elaboration of interactions between criticality phenomena*: In the next step, the individual selected candidates for criticality phenomena are further analyzed. To make their influence on criticality, measurable criticality metrics are employed that quantify specific aspects of criticality. A typical example of such a metric is the time to collision (TTC), indicating the minimal time until a collision occurs, provided no action is taken. To achieve a comprehensive causal understanding of how the different phenomena affect certain aspects of criticality, we model the underlying causal assumptions based on causal theory by Judea Pearl (Pearl 2009). This theory allows the qualitative and quantitative investigation of causal queries based on constructing a so-called causal graph that represents the causal relationships of the different factors on a certain abstraction level. *Figure 2* illustrates such a causal graph for the criticality phenomenon stationary occlusion of traffic participants.
3. *Consolidation and abstraction of criticality phenomena/convergence*: The last step of the criticality analysis maps the identified and relevant criticality

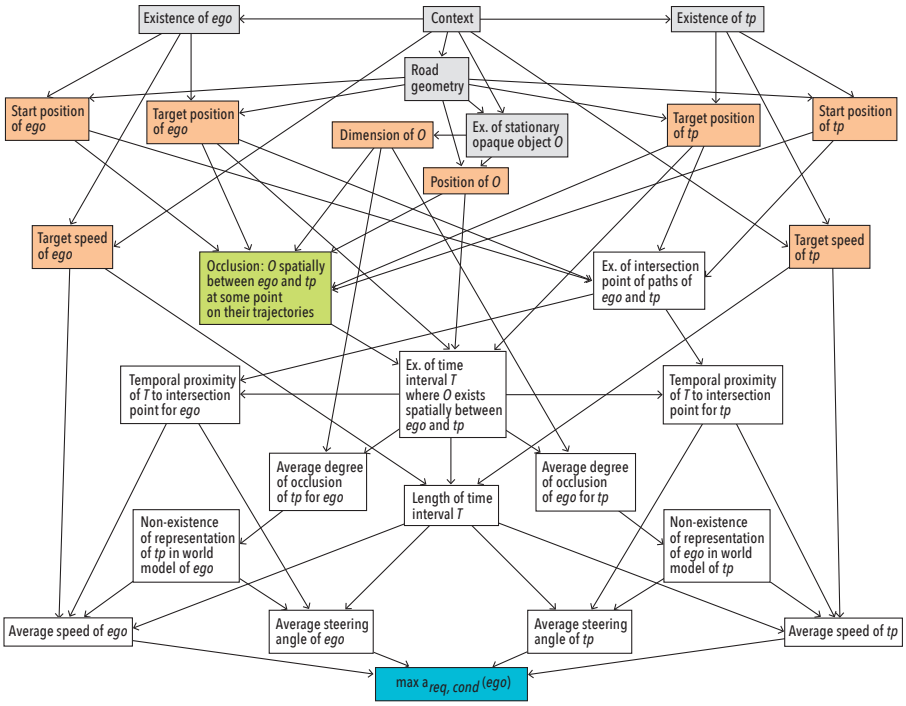

Figure 2. Causal graph for the criticality phenomenon "occlusion of a participant (tp)

phenomena to a manageable and finite set of classes of criticality phenomena. This assumes that such a manageable set must exist and that the number of criticality phenomena cannot be unlimited. If this were the case, the amount of data relevant to safe driving would surely exceed the processing capacity of human drivers. However, since we know that humans can drive a vehicle safely even in very complex situations, we can assume that there is a compact representation of the criticality phenomena. The procedure for generating the criticality classes takes each new criticality phenomenon from step 2 and compares it to the already identified classes of criticality phenomena. If these are similar, they are merged into a standard class. Otherwise, a new class is created. The process is continued until it is determined with sufficient statistical certainty that all new phenomena found in step 1 are only ever mapped to already known classes.

During the execution of the method, individual parts, particularly in Step 2, are iterated repeatedly. This is done until the underlying mechanisms are sufficiently understood. A manageable finite set of abstracted criticality phenomena remains, covering all criticality-triggering causes for the investigated system class in a given environment. However, let us note that the method can be presented here only in a highly simplified form, and the figure notably omits details on where and how the feedback loops tie in with the process. For a comprehensive description of the methodology, please refer to Neurohr et al. 2021.

HAZARD AND RISK ANALYSIS FOR AUTOMATED SYSTEMS

The second method we elaborate on is the automation risks method (Kramer et al. 2020) which defines a comprehensive approach to the hazard and risk analysis of automated driving functions. It addresses both functional safety and SOTIF (safety of the intended functionality) by sustaining existing safety processes of the standards ISO 26262:2018 and ISO 21448:2022 and complementing them where necessary (ISO2018, ISO2022). The focus is on hazards that are inherent in the system but are triggered by external influences of the automated function, such as situations where the automated driving function does not react appropriately to its current environment. This includes non-detection or misclassification of objects, such as a bicyclist not detected or misclassified as a pedestrian, erroneous recognition of non-existing objects, and wrong predictions of future events, for example, due to wrong dynamic models. Therefore, the method builds on established analytical techniques for hazard analysis and risk assessment while it adds significant enhancements to enable the applicability to automated systems.

The proposed method is designed to accompany the entire development process. It is beneficial to initiate its application early during the concept phase so that safety considerations can be integrated into the system as early as possible. As shown in Figure 3, the method contains several feedback loops between the concept phase and development that enable the consideration of adjustments in the system, especially the integration and analysis of risk mitigation measures based on the previously gained knowledge, such as the implementation of redundancies or the definition of a higher safety distance.

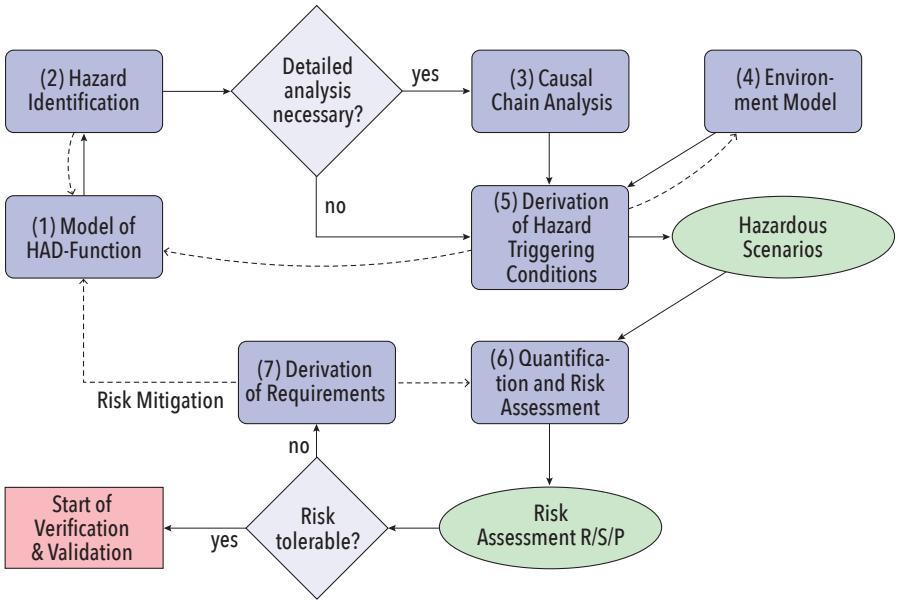

Figure 3. Overview of the automation risks method (Kramer et al. 2019)

ID	Basic Scenario	Basic Maneuver	Correct if (context)	Keyword	Incorrect Vehicle Behavior	Observable Effect(s) in Scenario	Additional Scenario Conditions (necessary for Top Level Event)	Potential Top Level Event
1	slower turn into path challenger	decelerate/ braking	front distance < safety distance	no	necessary breaking maneuver not initialed	ego continues with constant speed	challenger with significantly lower speed or critical Time-To-Collision	front/side collision with challenger
2				less	breaking maneuver not strong enough	Ego does not decelerate to prevent collision	challenger with significantly lower speed or critical Time-To-Collision	front/side collision with challenger

Figure 4. Table for identification of hazards on vehicle level (Kramer et al. 2019)

Functional Unit	Function			Key-word	Local Failure/ Functional Insufficiency	Basic Scenario	System Effect(s) in Scenario	Incorrect Vehicle Behavior	ID(s) of IVB	Possible System Cause(s)	Environ-mental Condition	Relevant for human driver?
	Input	Compu-tation	Output									
Sensors > Front camera > object recognition	camera image	segmen-tation	seg-mented camera image	no	segmented camera image not generated	slower turn into path challenger	challenger not detected by front camera > maneuver planning without information about the challenger	necessary braking maneuver not initiated	1	HW-failure, degradation or design fault	none	no statement
					no segments in camera image recognized	s/a	s/a	s/a	s/a	no night vision lacking sensibility at dark	darkness	likely (human vision also impaired by darkness)

Figure 5. Table for identification of hazards on component level (Kramer et al. 2019)

The approach involves two main parts: the identification of hazardous scenarios (Steps (1) – (5) in Figure 3. Overview of the automation risks method) and the quantification of corresponding risks (steps (6) and (7)).

The first part aims to identify hazards, understand the underlying causal relationships, and deduce scenarios that might trigger hazardous events. These hazardous scenarios serve as inputs to the following quantification part. They can also serve as a basis for comprehensive scenario-based testing within the verification and validation process and define a starting point for improvements in the system.

The investigation is based on an initial system description that involves at least an item definition and a functional architecture that describes an architectural model representing system functions, like sensor fusion or trajectory planning and their interactions. To identify hazards caused by incorrect behavior of the automated function, we employ a keyword-based brainstorming approach inspired by the hazard and operability study (HAZOP) (Ericson 2005, 365-381), a technique originated from the chemical industry. The main idea is to combine a set of basic scenarios with a set of basic maneuvers that the automated function could perform with a list of keywords to derive possible incorrect behavior of the automated system that might lead to harm. An example of such a table applied to a highway-chauffeur function is provided in Figure 4.

In the next step, we employ a second HAZOP-inspired approach to examine local failures and functional insufficiencies and their effects on the system and vehicle level by applying keywords to the individual functional units.

Based on the identified hazards, we aim to derive scenario properties that might provoke them. Therefore, we use a modified fault tree analysis (Ericson 2005, 183-222) which analyzes the causal chains starting from the top-level event of a hazard during a basic scenario.

A unique feature is that we denote environmental conditions in the tree wherever necessary for the propagation of a fault. We can derive the triggering

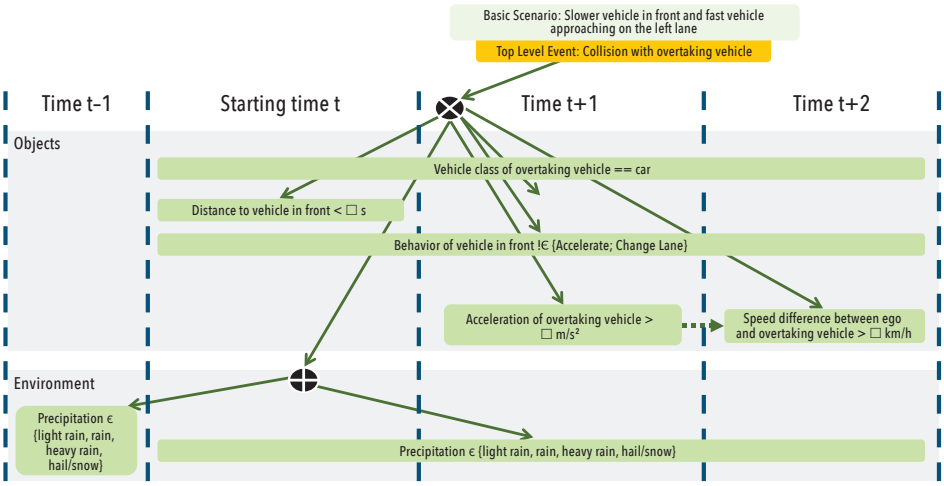

Figure 6. Exemplary part of an environmental fault tree reduced to the environmental conditions and chronologically ordered into discrete time steps

scenario properties by reducing the fault tree to these environmental conditions and identifying so-called minimal cut sets. An exemplary dependency graph is shown in Figure 6.

The quantification aims to derive a risk assessment that can be used to determine safety goals based on the afore-identified scenario properties. Therefore, it mainly builds on probability estimation. Relying on the probabilities of occurrence of the single environmental conditions and the conditional probabilities that an error propagates in the fault tree, we estimate the probability of a hazard occurring with the help of the single minimal cut sets representing the triggering scenario properties. This serves as a basis for the risk assessment according to the automotive safety integrity level (ASIL) of the ISO 26262:2018 (ISO 2018).

SUMMARY AND OUTLOOK

In this paper, we presented two methods that enable systematic investigation of criticality causes and their effects in the context of automated systems.

Criticality analysis aims at identifying a comprehensive list of all potential sources of criticality in a given application field which serves as input for certification authorities and test organizations to develop detailed homologation guidelines. The method is being developed in the VVMethoden project in close cooperation with representatives from the automotive industry.

The second approach describes an extension of a hazard and risk analysis in which functional safety is combined with SOTIF (safety of the intended functionality). This approach was developed in the PEGASUS project, where it was extensively tested using the example of a highway chauffeur function. A comprehensive description of the approach and the evaluation can be found in (Böde et al. 2019). Furthermore, we have investigated to what extent the approach can be adopted in other application domains. Vander Maelen describes the application of this method to a collision warning system in the maritime domain (Vander Maelen et al. 2019).

Currently, we are working on elaborating the methods, simplifying their application, and investigating other use cases. In two internal projects, we are investigating the suitability of these approaches for hazard detection in automated road traffic (https://verkehrsforschung.dlr.de/de/projekte/kokovi) and for automated ship navigation in port areas (https://verkehrsforschung.dlr.de/de/projekte/das-projekt-futureports-fuer-hochautomatisierte-digitalisierte-und-intermodal-vernetzte). ■

ACKNOWLEDGMENTS

This paper reports on collaborative work by many colleagues from DLR and partners in the projects PEGASUS (https://www.pegasusprojekt.de/en/home) and VVMethoden (https://www.vvm-projekt.de/en/). These projects received funding from the German Federal Ministry for Economic Affairs and Climate Action.

In particular, we want to thank our DLR (formerly OFFIS) colleagues Matthias Büker, Birte Neurohr, Christian Neurohr, Sebastian Vander Maelen, and Lukas Westhofen, the scientific leaders Werner Damm and Martin Fränzle, as well as our industrial cooperation partners Martin Butz (Bosch), Martin Bollmann (ZF), Ulrich Eberle (Stellantis) and Roland Galbas (Bosch) for their substantial contributions.

REFERENCES

- Böde, E., M. Büker, W. Damm, M. Fränzle, B. Kramer, C. Neurohr, and S. Vander Maelen. 2019. "Identifikation und Quantifizierung von Automationsrisiken für hochautomatisierte Fahrfunktionen." Technical report, OFFIS e.V.
- Ericson, C. A. 2005. *Hazard Analysis Techniques for System Safety.* John Wiley & Sons, Inc.
- Kramer, B., C. Neurohr, M. Büker, E. Böde, M. Fränzle, and W. Damm. 2020. "Identification and quantification of hazardous scenarios for automated driving." *International Symposium on Model-Based Safety and Assessment*: 163–178.
- Neurohr, C., L. Westhofen, M. Butz, M. H. Bollmann, U. Eberle, and R. Galbas. 2021. "Criticality analysis for the verification and validation of automated vehicles." *IEEE Access 9*: 18016-18041. doi:10.1109/ACCESS. 2021.3053159.
- Pearl, J. 2009. *Causality* (2nd ed.). Cambridge: Cambridge University Press. doi:10.1017/CBO9780511803161
- Vander Maelen, S., M. Büker, B. Kramer, E. Böde, S. Gerwinn, G. Hake, and A. Hahn. 2019. "An Approach for Safety Assessment of Highly Automated Systems Applied to a Maritime Traffic Alert and Collision Avoidance System." *2019 4th International Conference on System Reliability and Safety (ICSRS)*: 494-503, doi:10.1109/ICSRS48664.2019.8987712.
- ISO (International Organization for Standardization). 2022. ISO 21448:2022. Road vehicles — Safety of the intended functionality. Geneva, CH: ISO.
- ———. 2018. ISO 26262:2018. Road vehicles – Functional safety. Geneva, CH: ISO.
- SAE (Society of Automotive Engineers). 2021. SAE J3016:2021. Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles. Geneva, CH: SAE/ISO.

ABBREVIATIONS OF STANDARDS, CORRELATED WITH REFERENCE-LIST CITATIONS

ISO 21448:2022	(ISO 2022)
ISO 26262:2018	(ISO 2018)
SAE J3016:2021	(SAE 2021)

ABOUT THE AUTHORS

[Editor: Author biographies were current when the paper was initially published in 2022.]

Lina Putze is a researcher at the DLR Institute on Systems Engineering for Future Mobility. Her background is in mathematics.

Eckard Böde is a manager of the research group System Concepts and Design Methods at the DLR Institute on Systems Engineering for Future Mobility. His background is in computer science and model-based safety assessment of cyber-physical systems.

IS2025 KEYNOTE SPEAKERS

Langdon Morris

Senior Partner at InnovationLabs. Award-winning innovator, futurist, and world-renowned strategy consultant.

Jon Reijneveld

Co-Founder and Chief Engineer at The Exploration Company (TEC)

Robert Thirsk

Astronaut on two past missions: the shuttle Columbia and an expedition aboard the International Space Station

William Donaldson

Associate Professor of Management at Christopher Newport University

#INCOSEIS 2025

Ottawa

www.incose.org/symp2025

Enhancing Early Systems R&D Capabilities with Systems—Theoretic Process Analysis

Adam D. Williams, adwilli@sandia.gov

SAND2023-TBD. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the US Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525. *This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the US Department of Energy or the United States Government.*

■ **ABSTRACT**

Systems engineering today faces a wide array of challenges, ranging from new operational environments to disruptive technological — necessitating approaches to improve research and development (R&D) efforts. Yet, emphasizing the Aristotelian argument that the "whole is greater than the sum of its parts" seems to offer a conceptual foundation creating new R&D solutions. Invoking systems theoretic concepts of emergence and hierarchy and analytic characteristics of traceability, rigor, and comprehensiveness is potentially beneficial for guiding R&D strategy and development to bridge the gap between theoretical problem spaces and engineering-based solutions. In response, this article describes systems–theoretic process analysis (STPA) as an example of one such approach to aid in early-systems R&D discussions. STPA—a 'top-down' process that abstracts real complex system operations into hierarchical control structures, functional control loops, and control actions—uses control loop logic to analyze how control actions (designed for desired system behaviors) may become violated and drive the complex system toward states of higher risk. By analyzing how needed controls are not provided (or out of sequence or stopped too soon) and unneeded controls are provided (or engaged too long), STPA can help early-system R&D discussions by exploring how requirements and desired actions interact to either mitigate or potentially increase states of risk that can lead to unacceptable losses. This article will demonstrate STPA's benefit for early-system R&D strategy and development discussion by describing such diverse use cases as cyber security, nuclear fuel transportation, and US electric grid performance. Together, the traceability, rigor, and comprehensiveness of STPA serve as useful tools for improving R&D strategy and development discussions. Leveraging STPA as well as related systems engineering techniques can be helpful in early R&D planning and strategy development to better triangulate deeper theoretical meaning or evaluate empirical results to better inform systems engineering solutions.

INTRODUCTION

Systems engineering today faces a wide array of challenges, ranging from increasingly complex operational environments to new and novel interdependencies to dynamic (r)evolutionary technological changes to fluidly shifting roles of human actors. In response, research, and development (R&D) efforts have focused on developing solutions to address these challenges. Yet, R&D projects aimed to address interdependencies may struggle with managing uncertainty in the analysis, possibly resulting in overly narrow hypotheses or suffering "unintended consequences." Or R&D efforts attempting to capture the pace of technological change may suffer from scope creep and get stuck in a seemingly never-ending cycle of revising research objectives to align with the newest technological breakthrough. When R&D efforts venture into this arena, identifying the appropriate level(s) of complexity to address, particularly when applied to real systems, is of utmost importance.

If "the action of working artfully to bring something about (INCOSE 2023)," then there is a need for adequately addressing these challenges in early-systems R&D. Systems theoretic concepts that underpin systems engineering approaches, founded

on the Aristotelian argument that the "whole is greater than the sum of its parts," offer a conceptual foundation for better understanding how to transition from theoretical problem spaces toward practice, engineering-based solutions. Revisiting two classic concepts from general systems theory, of hierarchy and emergence, is informative in this endeavor. If there are fundamental differences and relationships between levels of complexity within a system (Von Bertalanffy 1950), then hierarchy is a concept by which to identify what generates, separates, and links each level. Once identified, the dynamics between and within hierarchical levels can be described as higher-ranking components and influences constraining the range of possible behaviors of components and influences at lower levels leading to a structure designable toward optimized performance. Consider, for example, between digital valve controllers communicating constraints on physical behaviors within nuclear power plant cooling systems.

Likewise, emergence describes the phenomenon wherein behaviors at an observed level of complexity are irreducible to and cannot be explained by the behavior or design of its subordinate components (Von Bertalanffy 1950). Once irreducibility is acknowledged, invoking emergence helps capture how observed system behaviors are, at least in part, driven by interactions within conditions, settings, and circumstances of system operations. For example, consider how risk for transporting spent nuclear fuel internationally relates to successfully executing combinations of technical, administrative, and procedural requirements. Taken together, hierarchy and emergence suggest that systems can be designed to leverage interactions toward desired system performance. Introducing these concepts into R&D project discussions can orient efforts toward exploring both individual component reliability and collective component interactions each offering the potential to improve protection schemes and resilience for the US electric grid.

These systems theoretic concepts introduce additional characteristics potentially beneficial for connecting the theoretical problem space to engineering-based solutions. First, consider traceability as the ability to track behavior or status during movement through a process. A better ability to track changes in or responses to R&D decisions can better help develop projects. Second, the quality of being thorough and accurate, or improved rigor, can help ensure that R&D strategies are optimal reflections of the problem being investigated.

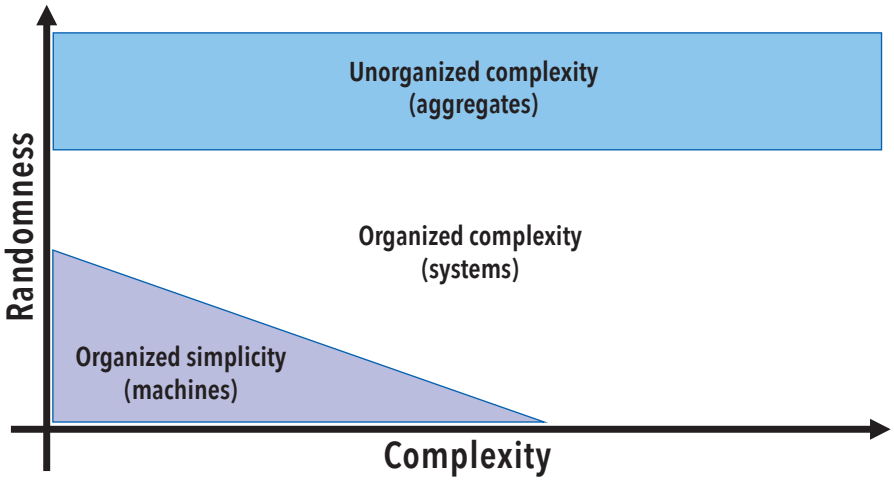

Figure 1. Comparison of zones of randomness and complexity, recreated from (Weinberg 1975)

Lastly, consider comprehensiveness as the ability to include all elements of a process, activity, or mechanism. Here, the extent to which the entire set of considerations are included will improve the impacts of R&D design and development.

Such characteristics seem helpful for adding structure to connect the theoretical problem space to potential engineering-based solutions, which also suggests similar applicability for scoping systems engineering R&D projects. In support of early systems R&D, this article offers systems-theoretic process analysis (STPA) as an example of an approach that has guided the development of R&D projects in such diverse use cases as cyber security, nuclear fuel transportation, and US electric grid performance. Leveraging STPA can be helpful in early R&D planning and strategy development to better triangulate deeper theoretical meaning or evaluate empirical results to better inform systems engineering solutions.

FROM THEORY TO PRACTICE

Expanding on the common Aristotelian argument, general systems theory provides the conceptual foundation for describing how observed performance is not always explainable by the behavior(s) of its constituent parts. For example, Figure 1 illustrates one way to explain system performance. Traditional R&D development practices are well-suited for using deterministic frameworks for addressing the zone of "simplicity" and stochastic approaches for addressing the zone of "unorganized complexity." Conversely, systems theoretic concepts and systems engineering techniques are uniquely suited to address the zone of "organized complexity", defined by Weaver (1948) as those "problems which involve dealing simultaneously with a sizable number of factors which are interrelated into an organic whole."

Beyond simply recombining components in attempts to describe real-world behaviors, systems engineering offers analysis technique and mental models to capture the non-statistical, non-random logic observed in the realm of organized complexity addressing the role of interactions, nth-order effects, and dynamism in understanding observed performance. Systems-theoretic process analysis (STPA) utilizes these concepts of emergence and hierarchy to provide traceability, rigor, and comprehensiveness in understanding of observed performance for complex engineering projects.

STPA is based on a causality model that defines safety of complex systems as the ability of a system to maintain a state that eliminates losses resulting from systems migrating into hazardous states and experiencing extreme external events (Leveson and Thomas 2018). Rather than emphasizing failure prevention, this framework analyzes safety as the avoidance of hazards and hazardous system states in terms of three fundamental-and controllable-concepts:

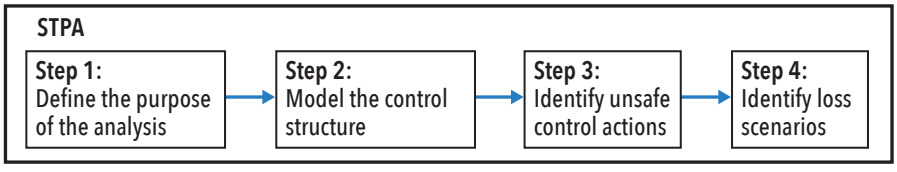

Figure 2. STPA process illustration, recreated from Leveson and Thomas (2018)

Table 1. Subset of UCAs for autonomous H-II transfer vehicle (HTV) operations, recreated from Appendix C in (Leveson and Thomas 2018)

Control Action (from ISS Crew)	Not providing causes hazard	Providing causes hazards	Too Early, Too Late, Order	Stopped Too Soon/ Applied Too Long
Abort	ISS crew does not provide Abort Command when emergency condition exists [H-1]	ISS crew provides Abort Command when HTV is captured [H-1] ISS crew provides Abort Command when ISS is in Abort path [H-1]	ISS crew provides Abort Command too late to avoid collision [H-1] ISS crew provides Abort Command too early before capture is released [H-1]	N/a
Capture	ISS crew does not perform Capture when HTV is in capture box in free drift [H-1]	ISS crew performs Capture when HTV is not in free drift [H-1] ISS crew performs Capture when HTV is aborting [H-1] ISS crew performs Capture with excessive/ insufficient movement (can impact HTV, cause collision course) [H-1]	ISS crew performs Capture too late, more than X minutes after HTV deactivated [H-1] ISS crew performs Capture too early before HTV deactivated [H-1]	ISS crew continues Performing Capture too long after emergency condition exists [H-1]

- **Constraints**, goals or set points by which higher levels within a hierarchy exhibit control of activities at lower levels based on the current understanding of the system being controlled
- **Control structures**, hierarchical model whereby the entire socio-technical system send commands and feedback signals to enforce constraints and avoid undesired system states
- **Process models**, abstracted representation of how a controller (for example, human or automation) processes and understand the process being controlled, including information regarding variable relationships, current system state, and the processes that can change the state of the system

Based on the logical analysis of this causality model, the goal of this analysis technique is to identify as many hazards as possible, thereby expanding the potential solution space to improve safety and providing decision-makers and designers with additional information to achieve desired complex behaviors (Leveson 2012). More specifically, STPA consists of four broad steps (Figure 2).

Where STPA was originally developed to support systems safety and hazard analysis, Step 1 offers a chance to explore other emergent systems properties including security, vulnerabilities, and risk. Based on the hierarchical control structure model of the system of Step 2, STPA uses control actions and feedback signals to illustrate communication between controllers whether physical, digital, or human and a controlled process (for example, normal nuclear power plant (NPP) operations). In this manner, STPA uses this set of desired control actions as a baseline for identifying a comprehensive set of logical violations for each. These logical violations are the analytic core of Step 3, and include:

- Necessary control commands *are not issued*;
- Unnecessary control actions (UCAs) *are issued*;
- Potentially correct control actions are *provided too early or late*; or,
- Potentially control actions are *stopped too soon (or too late)*.

The traditional language of STPA uses the phrase "unsafe" control actions — again, based on it's original development for systems safety. In this article, STPA is expanded to a broader set of potential emergent properties, suggesting the term "undesired" control action is more applicable.

The resulting undesired control action table, provided as an example in Table 1 illustrates how STPA can identify flawed interactions, mis-timed engineering activities, or incomplete communication structures, as well as component malfunction and hazards that occur when all components behave as expected. In general, identifying loss scenarios of Step 4 focuses on a more detailed description of why an undesired control action may happen.

STPA of Steps 1-3 provides a useful structure for guiding early R&D strategy and project development discussions. For example, STPA provides a high degree of traceability. More specifically, by linking component controls and constraints to system states of concern, STPA affords opportunities to not only identify where potential propagation of undesired behaviors could occur but can also map how potential design decisions matriculate throughout the hierarchical control structure model. Similarly, STPA is a rigorous process. By evaluating each component and their interactions in the hierarchical control structure model of Step 2 according to the logical violation categories the define undesired control actions of Step 3, STPA offers strict process for an exhaustive set of outputs. Lastly, STPA's logic paradigm suggests a comprehensive set of analytical outcomes. By defining desired system behaviors in terms of enforcing controls and constraints without prioritization, STPA inherently captures a wider range of realistic and plausible opportunities for undesired system performance.

Together, the traceability, rigor, and comprehensiveness of STPA serve as useful tools for bridging theoretical problem spaces with engineering-based solutions. Where STPA logically highlights how and where undesired control actions may manifest in a system, it provides opportunities to either define experiments to better understand the impacts of violated control actions on system behavior or guide development of novel solutions. Similarly, the hierarchical control structure model of STPA is uniquely suited to visualize — in a robust, yet clear manner — key interactions between components in a system that may (or may not) significantly impact overall system performance. Thus, the logical structure of STPA can be repurposed to help generate empirical designs, investigate theoretical underpinnings of performance, or explore

the efficacy of novel technologies to better inform systems engineering R&D efforts.

APPLYING PRACTICE TO RESEARCH

Though traditionally applied for evaluating safety operations in more mature or deployed systems, elements of STPA have successfully used to improve safety early in the system design life cycle (Fleming 2015) suggesting a similar ability to guide early R&D development and strategy discussions. For additional explanation, the following three use cases offer examples of STPA usage.

Use Case #1: Investigating the efficacy of evaluating hazards for digital instrumentation and control systems in nuclear power plants:

While the US Nuclear Regulatory Commission (NRC) mandates nuclear power plants prepare a cyber security plan, the lack of a consensus approach resulted in different plants taking different approaches to meeting this requirement. Traditional approaches focused on identifying critical digital assets and mapping them to safety-based risk assessment. Yet, the large number of probable cyber hazards, however, challenge the efficacy of deterministic approaches. This suggested a need to develop a risk-informed approach to explore possible hazards in digital components and systems in nuclear power plants (Williams and Clark 2019).

Invoking STPA, this project leveraged the concept of emergent systems behaviors as an organizing principle for better characterizing cyber security as an element of desired nuclear power plant operations that emerged from analog process components, digital systems, and operator actions. Accounting for the importance of these interdependencies between digital, physical, and human components within desired nuclear power plant operations is another key insight generated from STPA. Further, the (un)desired performance of nuclear power plants are described in terms of control action and feedback interrelationships between components in a control structure model.

The logical construct of STPA provided several key insights for this R&D project. First, the hierarchical control structure approach allowed for the creation of a hybrid model capturing both piping and instrumentation diagrams (P&ID) and digital network topologies shown in Figure 3. Though typically evaluated separately, merging these two descriptions of the nuclear power plant into a single diagram offers a more complete mapping of desired, and potential (un)desired, behaviors. Second, the logical foundation of undesired control actions is conceptually like basic events in fault trees, an insight gained from applying the STPA hazard analysis approach and comparing those results against more traditional fault

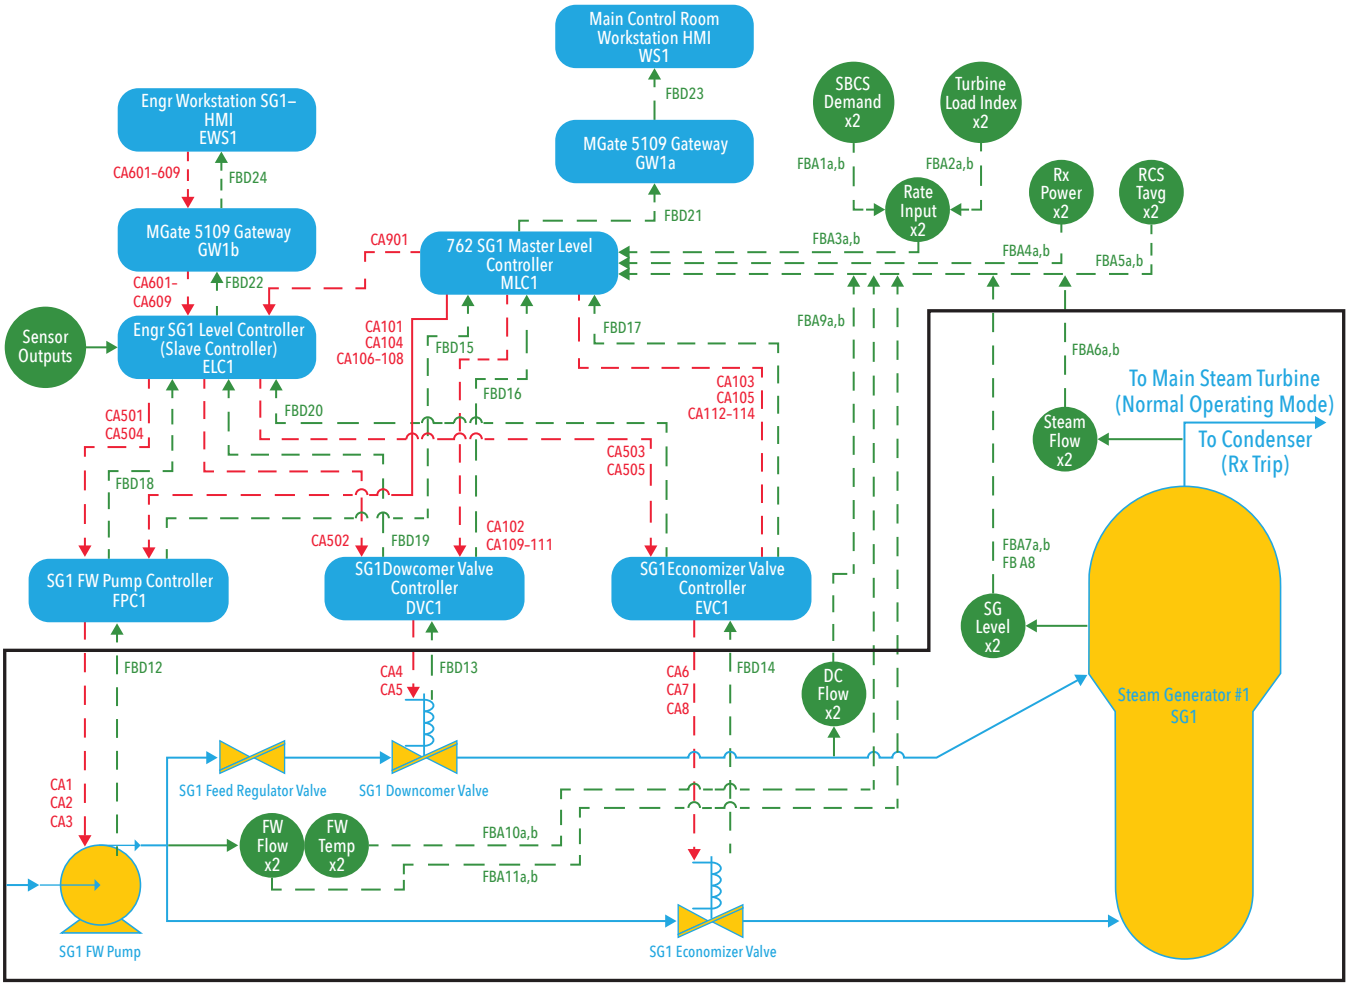

Figure 3. Notional main feedwater control digital and physical systems modeled as a hybrid STPA-related hierarchical control structure, from Williams and Clark (2019). Blue elements are controllers (in the STPA sense) which execute control actions (via red dashed lines) and green circles are sensors within the system which report various types of feedback (via green dashed lines—both of which are mapped onto a more traditional P&ID (the bottom portion that is blocked off). Describing notional nuclear facility system in this manner helped highlight several key features for improving cyber security.

Table 2. Summary of STPA-generated states of increased risk for a representative set of control actions for international SNF transportation, from Williams (2018). The middle column illustrates the benefit of invoking STPA early in R&D to help identify novel undesired versions of control actions—namely the 3S control actions not linked to a safeguards, safety, or security control action (last two rows). early in the (middle column). The right-hand column demonstrates the traceability from undesired versions of control actions to a range of states of increased risk (SIR)—where SIRs are not prioritized but labeled for categorial purposes.

Control Action	STPA Label 3S STPA Label	State of Increased Risk (SIR) [STPA hazard type]
Transmit GPS location of SNF cask	Safeguards Control Action1	SIR10 [NNP1,2]
	3S Control Action1	SIR10, SIR12 [NNP1,2]
Submit confirmation of removing SNF from inventory within 48 hours to IAEA	Safeguards Control Action2	SIR10, SIR11 [NNP] SIR10 [PNN2]
	3S Control Action2	SIR10, SIR11, SIR12 [NNP] SIR10, SIR12 [PNN2]
Physical assessment of cask contents in appropriately sealed facility	Safety Control Action1	SIR1, SIR2 [NNP2] SIR1, SIR2 [PNN1,2]
	3S Control Action3	SIR12 [NNP1] SIR1, SIR2 [NNP2] SIR1, SIR2, SIR5, SIR7 [PNN1,2]
Stop acceleration once at 55 mph	Safety Control Action2	SIR4 [NNP1]
	3S Control Action4	SIR4 [NNP1] SIR8 [Too early]
Engage rail car immobilization mechanism	Security Control Action1	SIR5, SIR6 [NNP] SIR5, SIR7 [PNN1]
	3S Control Action5	SIR5, SIR6 [NNP] SIR5, SIR7 [PNN1] SIR2 [PNN2]
Communicate the process for transferring armed security responsibility	Security Control Action2	SIR9 [NNP] SIR7, SIR9 [PNN1]
	3S Control Action6	SIR5, SIR9, SIR10 [NNP] SIR5, SIR7, SIR9 [PNN1]
Harmonize concepts of operations across safety, security, and safeguards	3S Control Action7	SIR3, SIR12 [NNP1] SIR1, SIR2 [NNP2] SIR1, SIR2, SIR5, SIR7 [PNN1,2]
Coordinate between safety, security, and safeguards during emergency plans	3S Control Action8	SIR3, SIR12 [NNP1] SIR1, SIR2 [NNP2] SIR1, SIR2, SIR5, SIR7 [PNN1,2]

STPA Hazard Types: NNP = "needed, not provided"; PNN = "provided, not needed"; Too early = "provided tool early"
Subscripts denote a particular conditional description for a violated control action aligned with a given state of increased risk

tree-based analysis for nuclear power plant safety. This resulted in the project incorporating undesired control actions into fundamentally new models called "systems-theoretic informed fault trees," or SIFTs. SIFTs utilize key systems theoretic concepts to expand upon traditional fault trees by incorporating (1) the uniqueness and complexity of digital components and (2) newly identified causes of hazards, including those from component interactions and that still result with no component failure occurring.

Including STPA-generated undesired control actions into SIFTs provides enhanced traceability in two manners. First, the inherent ability of STPA to identify potential propagation of undesired behaviors could occur, and associated potential to map how design decisions matriculate, is enhanced by the additional structure given to undesired control actions in the fault trees. The second element of traceability relates to mapping the fault tree cut set solutions into the following categories: random component mechanical failures, combinations of mechanical and undesired digital control actions, only undesired digital control actions. The fact that the SIFTs identify new types of cut sets is indicative of the rigor offered by invoking STPA. These categorical cut sets also speak to both the comprehensiveness of this R&D particularly considering that the cut set category "random component failures" matches solutions of traditional fault tree analysis. The end result of this STPA-inspired research project is the hazards and consequences analysis for digital systems (HAZCADS) analysis technique (EPRI 2018) currently being implemented to improve cyber security in US nuclear power plants.

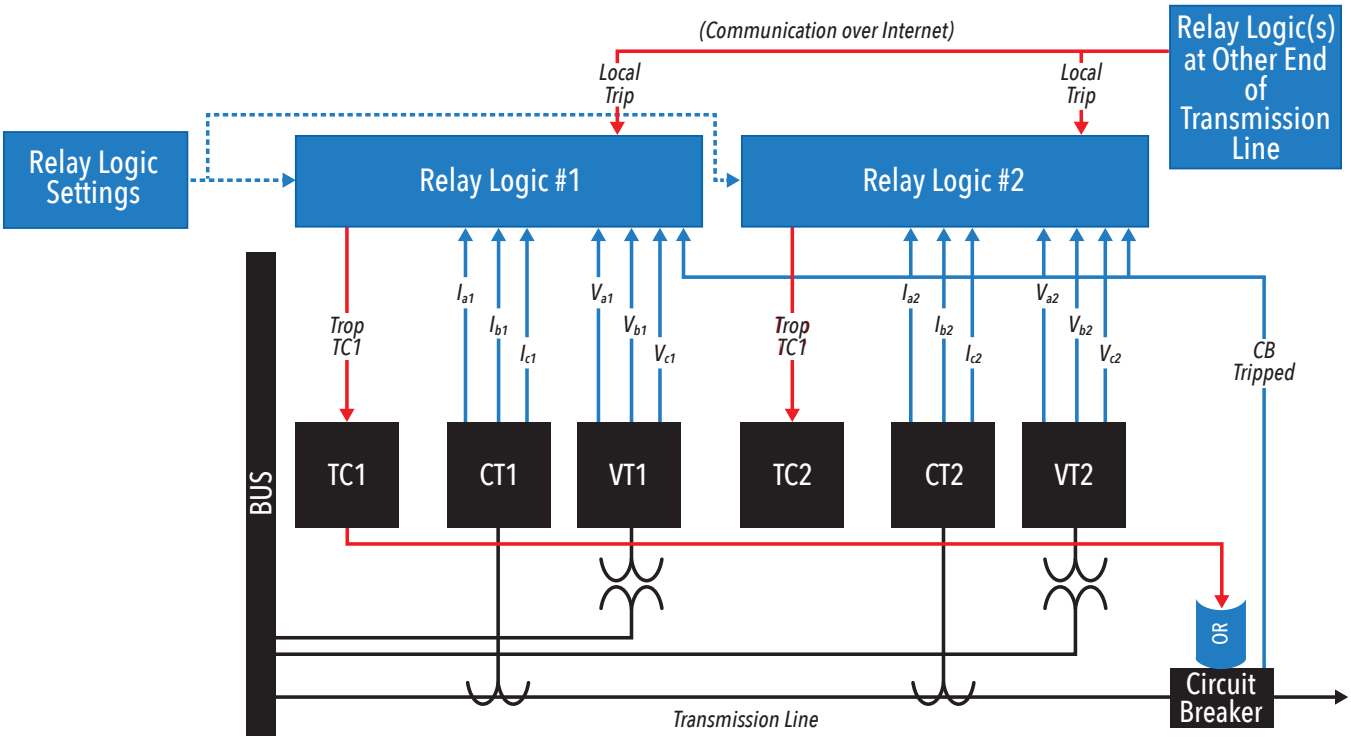

Figure 4. Notional .S grid transmission line protection systems modeled as a hybrid STPA-related hierarchical control structure. The blue boxes are modeled as STPA controllers (which execute control actions via the red arrows) and the black boxes are system elements and sensors feeding back signals on the state of the system (via the blue arrows)—both of which are overlaid on a black block diagram tracing electricity transmission. Describing the system in this manner helped identify key insights to improve grid protection and resilience.

Use Case #2: Examining the dynamics of safety, security, and international safeguards for international spent nuclear fuel transportation:

Real-world observations and expected operational realities illustrate increasingly complex challenges transporting spent nuclear fuel successfully and without incident. Yet, traditional analysis methods struggle to capture dynamics related to such anticipated challenges as overlaps in risk mitigation responsibilities, conflicting regulatory objectives, increases in transfers between transportation modes, and multiple geopolitical or maritime border crossings. In response, a research project was initiated to explore an analytical solution capable of evaluating challenges to safety (for example, preventing an accidental radiological release), security (for example, protecting against intentional malicious acts) and safeguards (for example, averting state-sponsored diversion of nuclear material) of spent nuclear fuel transportation (Williams 2018).

The core concepts of hierarchy and emergence inherent within STPA helped guide project planning discussions, particularly in terms providing a framework for appreciating the (in)direct relationships between hazards, threats, and risks to spent nuclear fuel transportation in complex globalized environments. STPA evaluates the ability for the spent nuclear fuel (SNF) transportation system to achieve its mission to physically move SNF from an origin facility to a destination facility without disruption of selected and approved routes, timelines, and operations. The underlying logic of STPA suggests that, if the system migrates into any of these potential states of increased risk, whether a safety, security, or international safeguards-focused risk, one additional external event could result in one of these unacceptable losses. For example, STPA argues that unauthorized access to the spent nuclear fuel during the transport results in a state of increased risk. The specific cause or contributing factors to the unauthorized access can range from the intentional use of explosives or a cask breach from an unintentional derailment. From the STPA perspective, the goal is not to prevent these causes but to design technical, administrative and systemic controls to keep the cask from experiencing unauthorized access and thereby entering the state of higher risk. Between combined hierarchical control structures and joint undesired control action analysis shown in Table 2, STPA guided the project discussion toward identifying a range of designed controls to mitigate the risks and unacceptable losses of international spent nuclear fuel transportation.

Introducing an STPA-based approach provided a high degree of traceability from undesired control actions to their associated states of increased risk and unacceptable system performance losses when evaluating safety, security, and international safeguards risk for international spent nuclear fuel transportation. Tracking propagation of undesired control actions also highlighted key areas of interdependence between safety, security, and international safeguards mitigations. For example, even though a security design decision can prevent unauthorized access to the cask, a violated security control could also result in an unplanned radiological release, a large safety hazard or a loss of continuity of knowledge that is a safeguards issue. In terms of rigor, applying STPA illustrated the importance of the "provided, not needed" control action violation where interdependence between safety, security, and international safeguard control actions existed including identifying states of increased risk missed by more traditional approaches. The ability for this research to identify additional states of increased risk

that not directly aligned with desired levels of safety, security, and international safeguards performance indicates a more comprehensive solution. STPA inspired insights and results from this research project have produced an analytical framework more aligned with the real-world, multi-modal, and multi-jurisdictional nature of ensuring adequate safety, security, and international safeguards during international spent nuclear fuel transportation.

Use Case #3: Exploring approaches to improve resilience-based and risk-informed decision making for project the US electric grid:

Recent events, including a 2014 shotgun attack on an electrical substation in California, the 2021 cyberattack on the Colonial pipeline in Texas, and the 2023 suspected domestic terrorist plot to attack substations in Maryland demonstrate the need to re-evaluate resilience analysis for the US electrical grid. This suggests a need to build scientific and logical arguments for analyzing potential vulnerabilities to craft more robust and comprehensive strategies to mitigate risk and increase resilience in the electric grid. Typical protections observed across the US electric grid are a mix of common baseline protections augmented by piecemeal, bespoke efforts that tend to be poorly coordinated. Further, many of these protections emphasize preserving individual grid components, which often results in shutting a component down, perhaps prematurely which may then cause downstream components to pass performance thresholds and cause rolling brown/black outs. The fact that the US electric grid consists of three major regions and more than 120,000 miles of lines operated by 500 companies is an additional challenge to improving resilience.

In response, asserting that risk, resilience, and vulnerabilities are emergent properties hypothesizes that STPA based thinking can evaluate how the electrical grid would recover following unknown, but anticipated, perturbations. Further, if vulnerabilities are conceptualized as opportunities to create undesired consequences, then resilience can be conceptualized as using control actions to ensure desired performance levels in STPA. The logic underpinning hierarchical control structure models can help illustrate the range of controls necessary to ensure generation, distribution, and transmission functions are maintained at desired performance levels, shown in Figure 4, as well as offer insights for characterizing spatial elements of risk and resilience. Similarly, the STPA undesired control action can help capture transient and dynamic interactions between resilience phases, with the last two undesired control action categories identifying temporal elements of risk and resilience.

The STPA basis for this research project allowed the inherent traceability to better specify connections between nodes in grid networks. In addition, the hierarchical control model provided the scaffolding on which to investigate both temporal and spatial elements of grid resilience. Even with the preliminary work in this project focusing on more simplistic representations, STPA demonstrated the ability to describe an exhaustive set of undesired control actions that directly challenge the resilience of the electric grid, which also suggests a similar level of rigor for higher fidelity grid descriptions. The STPA-generated hierarchical control structure approach also affords an opportunity to create template models for commonly occurring subsets of nodes within the US electric grid. By extension, undesired control actions associated with each template can be identified more quickly and novel interactions highlighted more efficiently as the template models are connected in ways to capture more comprehensively larger, more realistic grid (sub)systems. Ultimately, the traceability, rigor, and comprehensiveness of STPA will continue to drive the research to develop new and novel protection schemes to improve the resilience of the US electric grid.

CONCLUSIONS AND IMPLICATIONS

As demonstrated in the three use cases described in the previous section, STPA provides a logical foundation and analytical framework for connecting theoretical problem spaces to potential engineering-based solutions to aid in scoping systems engineering R&D projects. Hierarchical control structure models demonstrated traceability and comprehensiveness in evaluating cyber security for nuclear power plants and risk analysis for international spent nuclear fuel transportation. STPA derived undesired control actions illustrated rigor in identifying interdependent risks in international spent nuclear fuel transportation, as well as in characterizing temporal and spatial risk elements challenge US electric grid resilience.

Retuning again to the logical and theoretical foundations for STPA, invoking the phenomena of hierarchy and emergence provide useful guide rails for early-systems R&D discussions. STPA's hierarchical control structure models help capture — and simplify— the complexity in modern systems by abstracting them in a manner that both emphasizes the importance of component performance and the various interactions between them. Likewise, the emphasis on maintaining emergent system performance within a desired range offers a mechanism for exploring the flexibility of potential control actions or the range of possible redesigns to expand that desired operational space. Together, STPA can help inform either experimental design, hypothesis generation, system redesign options, or the characteristics necessary for novel, next-generation solutions.

Challenges experienced in bridging theoretical problem spaces with engineering-based solutions have both persisted and necessitated new approaches and potential solutions. As Von Bertalanffy eloquently stated in 1972:

> Modern technology and society have become so complex that the traditional branches of technology [and analysis] are no longer sufficient; approaches of a holistic or systems, and generalist and interdisciplinary nature become necessary (420).

The ability of STPA to provide holistic, generalist, and interdisciplinary solution to these challenges have also been demonstrated in a range of domains including aerospace (Fleming and Leveson 2014), medical (Pawlicki, et. al. 2016), automotive (Placke Duo 2015), port security (Williams 2015), and cyber security (Bakirtzis, et. al. 2017) in addition to the three use cases demonstrated in this article. To the extent that logical and analytical characteristics of other techniques are like STPA, then similar benefits for scoping and guiding R&D strategies can be anticipated. Ostensibly, this suggests that STPA based approaches, and perhaps broader systems engineering at large, provide a bridge for applying "artful will of bringing something to fruition" to improve early systems engineering R&D efforts. ■

REFERENCES

- Bakirtzis, G., B. T. Carter, C. H. Fleming, and C. R. Elks. 2017. "MISSION AWARE: Evidence-Based, Mission-Centric Cybersecurity Analysis." arXiv preprint arXiv:1712.01448.
- Electric Power Research Institute. 2021. *Hazard Analysis Methods for Digital Instrumentation and Control Systems Technical Report—Revision 1 (3002016698)*, Palo Alto, US-CA.
- Fleming, C. H. 2015. "Safety-driven early concept analysis and development." Dissertation, Massachusetts Institute of Technology (Cambridge, US-MA).
- Fleming, C. H., and N. G. Leveson. 2014. "Improving Hazard Analysis and Certification of integrated Modular Avionics." *Journal of Aerospace Information Systems*, 11(6).
- International Council on Systems Engineering (INCOSE). 2023. "About Systems Engineering." https://www.incose.org/about-systems-engineering.
- Leveson, N. G. 2012. *Engineering a Safer World: Systems Thinking Applied to Safety*. Cambridge, US-MA: MIT Press.
- Leveson, N., and J. P. Thomas. 2018. *STPA Handbook*. Cambridge, US-MA: Partnership for Systems Approaches to Safety and Security.
- Pawlicki, T., A. Samost, D. Brown, R. Manger, G.-Y. Kim, and N. Leveson. 2016. "Application of Systems and Control Theory-Based Hazard Analysis to Radiation Oncology." *Journal of Medical Physics*, 43 (3), 1514-1530.
- Placke, S., J. Thomas, and D. Suo. 2015. Integration of Multiple Active Safety Systems Using STPA. SAE Technical Paper 2015-01-0277. SAE.
- Von Bertalanffy, L. 1950. "An outline of general systems theory." *British Journal for the Philosophy of Science*, 1, 134-165.
- Von Bertalanffy, L. 1972. "The history and status of general systems theory." *Academy of Management Journal*, 15(4), 407-426.
- Weaver, W. 1948. "Science and Complexity." *American Scientist*, 36(4), 536-544.
- Weinberg, G. M. 1975. *An Introduction to General Systems Thinking*. New York, US-NY: Wiley
- Williams, A. D. 2015. "Beyond a series of security nets: applying STAMP & STPA to port security." *Journal of Transportation Security*, 8(3-4), 139-157.
- Williams, A. D. 2018. "Using Systems Theory to Address Complex Challenges to International Spent Nuclear Fuel Transportation." Paper presented at the 28th Annual International Symposium of INCOSE, Washington, US-DC, 7-12 July.
- Williams, A. D. and A. J. Clark. 2019. "Using Systems Theoretic Perspectives for Risk-Informed Cyber Hazard Analysis in Nuclear Power Facilities." Paper presented at the 29th Annual International Symposium of INCOSE, Orlando, US-FL, 20-25 July.

ABOUT THE AUTHOR

[Editor: Author biography was current when the paper was initially published in 2023.]

Dr. Adam D. Williams is a principal R&D systems engineer in the Center for Global Security and Cooperation at Sandia National Laboratories. Dr. Williams serves as principal investigator for R&D efforts in human/system interactions for complex systems, evaluating vulnerabilities in cyber-physical nuclear systems, managing complex risk in the nuclear fuel cycle, and exploring alternatives for future nonproliferation and arms control futures. He also leads initiatives for the US Department of Energy, National Nuclear Security Administration, the US Department of State and Laboratory Directed Research and Development (LDRD). He has a BS in mechanical engineering, MA in international affairs/national security, and PhD in human and systems engineering.

Need to Scale Active Learning in the Enterprise?

We Have the Science for That.

Empower your team with the essential digital engineering and machine learning skills to excel in today's competitive landscape. With a strong foundation in SE/MBSE taught by industry-leading experts, your team will be well-equipped to elevate performance and enable data-driven organizational transformation. Custom client programs and public courses available. Explore your options today.

Customizable Learning Programs For Your Organization & Teams

Customizable Programs for Groups and Enterprise Learning

New Courses in Artificial Intelligence & Machine Learning

Programs for Individuals

Short Courses and Certificates for Individuals

Caltech Center for Technology & Management Education

Get started:

ctme.caltech.edu

Connect with an advisor:

execed@caltech.edu

Aerospace • Agriculture • Automotive • Biotech • Chemical • Communications
Defense • Electronics • Energy • Government • High-Tech • Life Sciences
Medical Devices & Diagnostics • Precision Manufacturing • Scientific Research

INCOSE

International Council on Systems Engineering

A better world through a systems approach

PRESENTS

FuSE

Future of Systems Engineering

ITS MISSION

FuSE refines and evolves the SE Vision 2035 across competencies, research, tools & environment, practices, and applications

FuSE identifies critical gaps towards the vision realization and initiates & supports relevant actions

FuSE fosters involvement and collaboration within and outside of INCOSE

FuSE educates, shares success, and expands

CONNECT & GET INVOLVED

INCOSE.ORG/FUSE

SPONSORED BY

SPEC INNOVATIONS