SAND2024-16614PE

Sandia National Laboratories

Exceptional service in the national interest

DIGITAL ASSURANCE FOR HIGH CONSEQUENCE SYSTEMS

MISSION CAMPAIGN

DAHCS MICOP | DEC 11th, 2024

Shelley Leger
Technical Lead

Will Zortman
Campaign Manager

U.S. DEPARTMENT OF ENERGY NNSA National Nuclear Security Administration

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 U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525

SAND2024-16614PE

OUTLINE

- Sandia LDRD
- Digital Assurance for High Consequence Systems (DAHCS) Mission Campaign (MC)
- DAHCS Research (LDRD) Call
- How to get involved
- Questions/Discussion

SANDIA'S LABORATORY DIRECTED RESEARCH AND DEVELOPMENT (LDRD) PROGRAM

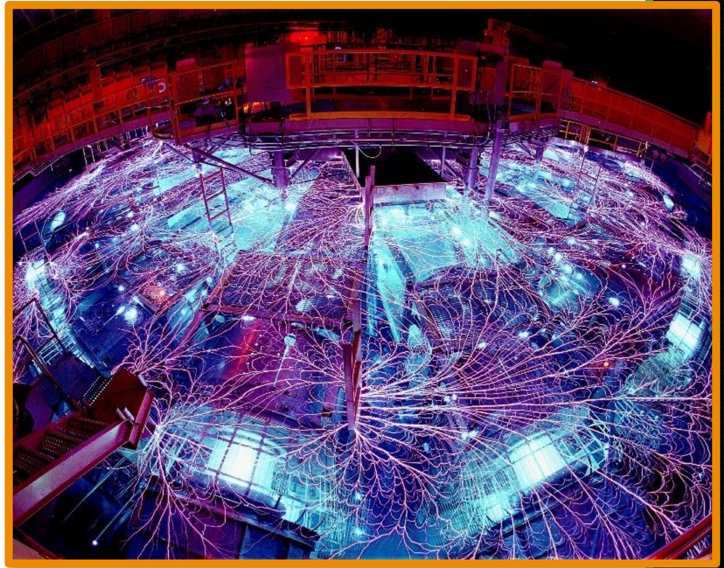

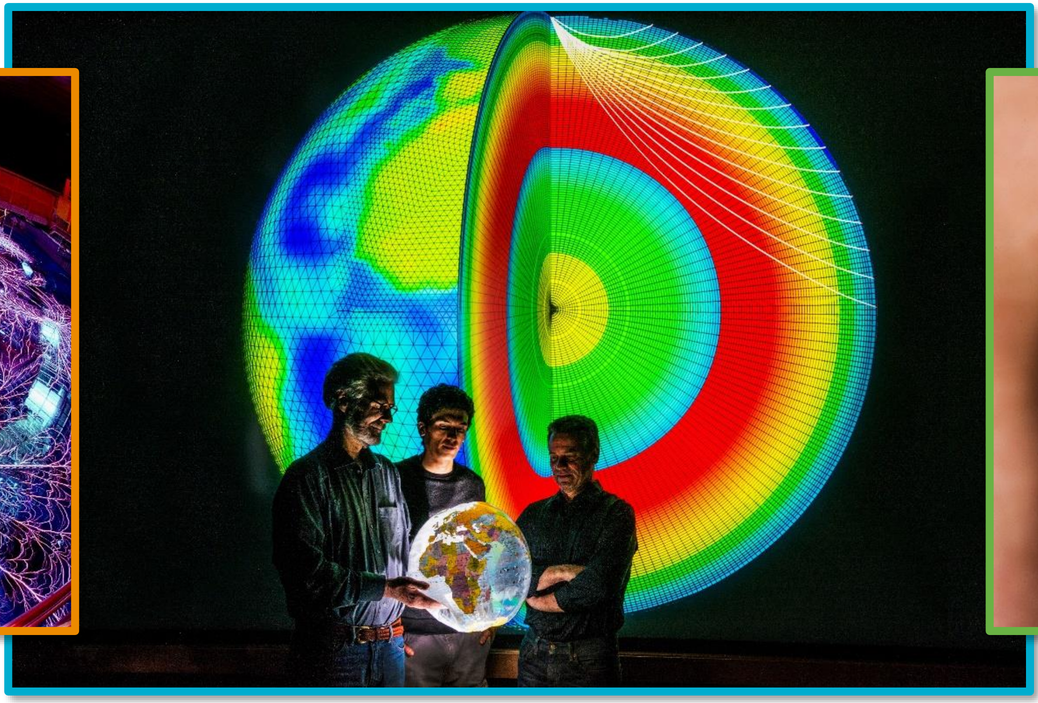

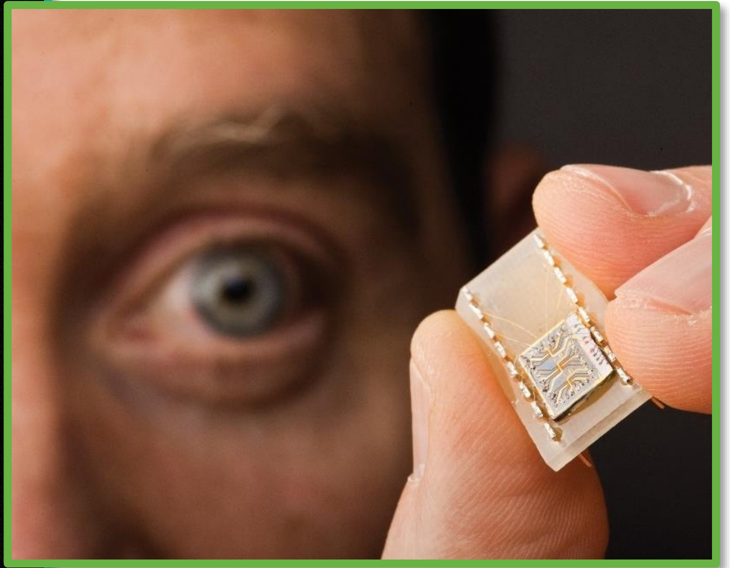

Enable national security missions

Attract, develop and retain a world-class technical workforce

Develop innovative solutions and novel tools

As Sandia's sole source of discretionary R&D funding, the LDRD program provides the flexibility to anticipate and respond quickly to future mission needs and to explore potentially revolutionary advances in science and technology.

LDRD FUNDING AND UNIVERSITY COLLABORATIONS

SANDIA LDRD PROPOSALS

- Are driven by Sandia PIs
- Facilitate connections between Sandia and academia
- Seed project collaborations (as opposed to maintaining collaborations)

UNIVERSITY COLLABORATIONS

- Can be funded through the core project
- Can be funded through supplemental funding from the DAHCS MC
- Some non-DAHCS LDRDs may be managed by the Sandia University Partnerships Network

Sandia PIs submit proposals, but they can collaborate with faculty on the proposal within certain bounds.

Funding can be used to support faculty research, student research assistants, and project-related travel – and allowable purchases if necessary.

Communicating project needs and purchases early will enable Sandia to determine allowability.

U.S. citizenship is preferred for faculty and is expected for students working on Sandia LDRDs.

PARTNERING WITH SANDIA

- Sandia LDRD funding is ***not* a grant** – it **is a research contract** managed by Sandia with project-defined deliverables.
- Universities invoice projects at least monthly and adhere to contracted terms and conditions (e.g., pre-publication review).
- University accomplishments and project results must be received **by the end of August** each year.
- LDRD funding does not carry over Sandia's fiscal year (FY) boundary and must be costed **by September 30** each year.
- Sandia PIs engage regularly with faculty and student/postdocs working on projects; Sandia PIs will report on results and accomplishments.
- Sandia may share some sensitive information, so universities should use caution in discussing Sandia project information.

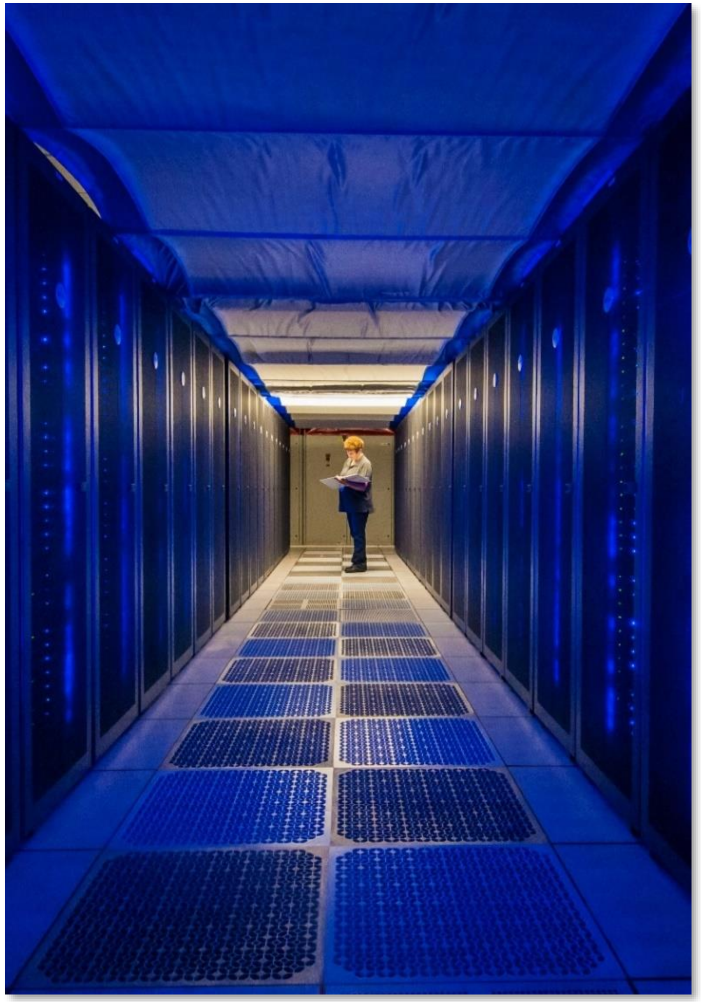

Funding runs from October 1 (with a completed contract) to September 30.

LDRD PROPOSAL SCHEDULE

COLLABORATION WINDOW

LDRD SCHEDULE

- DAHCS Assumptions Collaboration Day
- DEC. — Call Publicized
- JAN. — Info Sessions
- FEB. — Idea Deadline
- MAR. — DAHCS Principles Collaboration Day
- APR. — Ideas Invited to Full Proposal
- MAY — Proposal Deadline
- JUN.
- JUL. — Proposals Selected for Funding
- AUG. — Statement(s) of Work, etc.
- SEPT. — DAHCS Demos Collaboration Day
- OCT. — New Projects Begin Costing

Faculty should submit interests or ideas related to the call if they would like help connecting with potential Sandia proposers.

DAHCS MC OVERVIEW

BLUF

rigorous | rapid | cost-effective | generalizable | across system lifecycles

VISION

Disciplined systems engineering processes support systems-level tradeoffs against digital assurance.

Outcomes:

HIGH CONSEQUENCE DIGITAL ASSURANCE
from design to decommissioning

We have the foundation to efficiently and confidently:

- ✓ Characterize digital technologies
- ✓ Assess risks to our systems from digital technologies
- ✓ Select among options that appropriately manage and accept digital risk

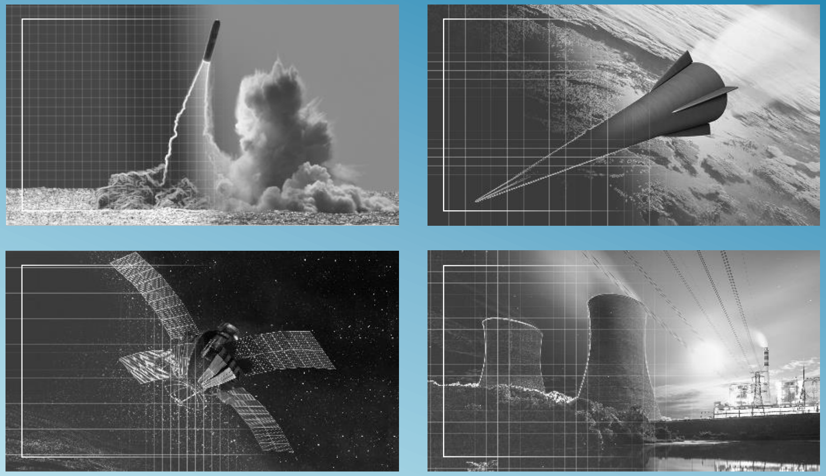

Goal: Create an ecosystem that gives us rapid confidence in our systems' digital assurance.

BLUF

Go "wild"! In high-risk research, the path may not be clear.

rigorous | rapid | cost-effective | generalizable | across system lifecycles

Outcomes:

HIGH CONSEQUENCE DIGITAL ASSURANCE
from design to decommissioning

We have the foundation to efficiently and confidently:

- ✓ Characterize digital technologies
- ✓ Assess risks to our systems from digital technologies
- ✓ Select among options that appropriately manage and accept digital risk

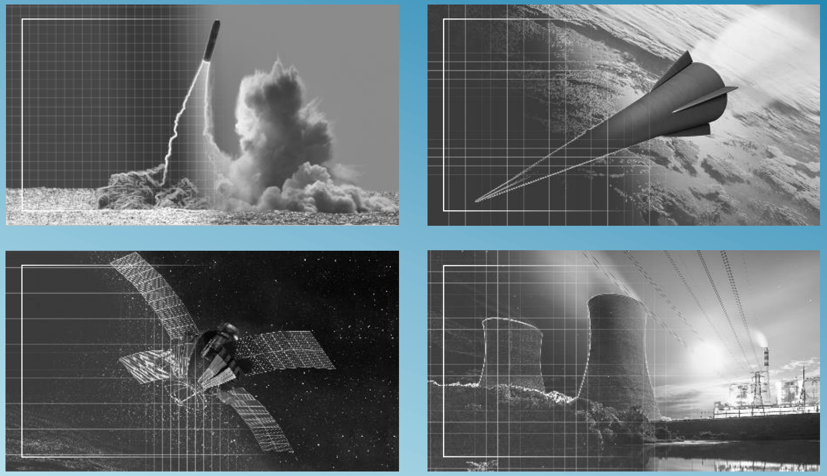

VISION

Disciplined systems engineering processes support systems-level tradeoffs against digital assurance.

Leap beyond existing research and create artifacts for larger systems.

Goal: Create an ecosystem that gives us rapid confidence in our systems' digital assurance.

DIGITAL ASSURANCE FOR HIGH CONSEQUENCE SYSTEMS (DAHCS)

HYPOTHESIS: DAHCS principles exist

VISION

Disciplined systems engineering processes support systems-level tradeoffs against digital assurance.

RESEARCH THRUSTS

Revolutionary DAHCS, Targeted Evaluation

- **SCALABLE ANALYSIS**
- **IMPACT ANALYSIS AMID UNCERTAINTY**
- **INTEGRATING WITH SYSTEMS ENGINEERING**

Goal: Create an ecosystem that gives us rapid confidence in our systems' digital assurance.

MISSIONS ARE AT RISK

BOEING 737 MAX 10

www.usatoday.com/story/todayinthesky/2013/04/15/reuters-lion-air-pilot-felt-jet-dragged-from-the-sky/2084899/

A "WICKED", SEEMINGLY IMPOSSIBLE PROBLEM

RAMPANT DISCONTINUITY

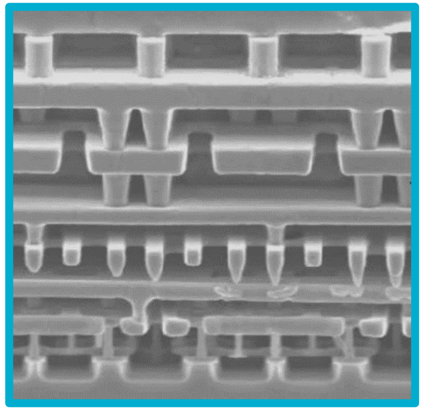

Billions of *interconnected transistors*

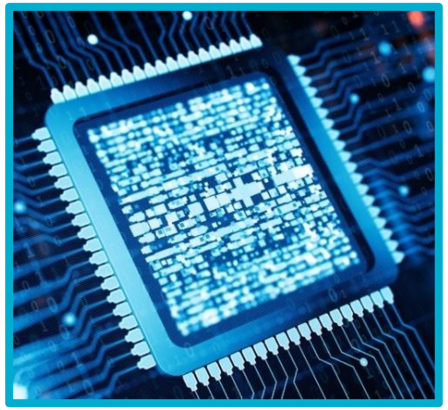

More states than particles in the observable universe

Overwhelming numbers of behaviors (outputs)

Tiny perturbations can dramatically change behaviors

What you want...

What you get...

A “WICKED” PROBLEM

RAMPANT DISCONTINUITY

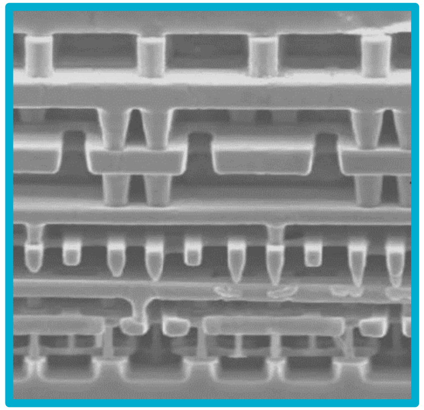

Billions of *interconnected transistors*

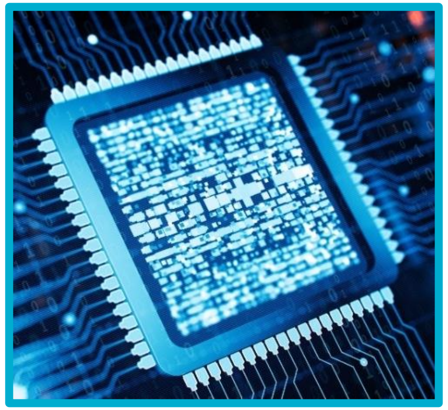

More states than particles in the observable universe

Overwhelming numbers of behaviors (outputs)

Tiny perturbations can dramatically change behaviors

ABSTRACTION

mathematica.stackexchange.com/questions/125043/how-to-make-a-spherical-cow

IMPLEMENTATION

mathematica.stackexchange.com/questions/125043/how-to-make-a-spherical-cow

The problem may seem impossible,
but standard scientific approaches can make it more tractable.

DEFINITIONS

High Consequence Systems (HCS)

Systems that serve very specific missions where **failure to function** can result in **unacceptable** consequences, e.g., grave damage to national security, catastrophic damage, or extensive loss of life.

Example "general" HCS characteristics:

- *Embedded cyber-physical controllers, often digitally simpler state machines*
- *Mission constraints (e.g., time, size, weight, power)*
- *Specific environments and purpose*
- *Long service life*
- *Different threat model*
- *Rigorous requirements*

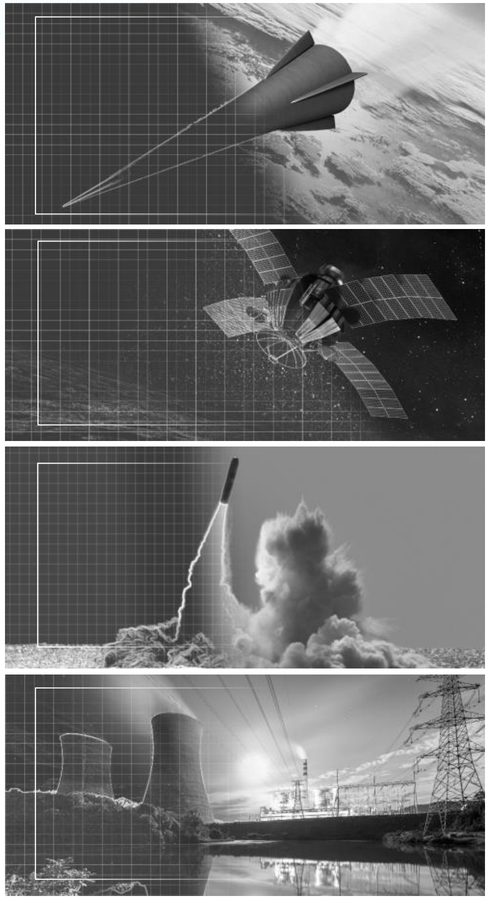

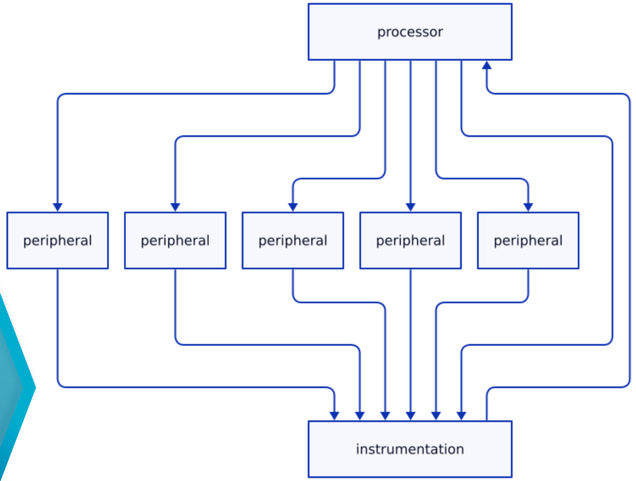

DAHCS MC Concern: **Embedded Cyber-Physical Controller Failure to Function** (availability/reliability) – due to either adversaries *or unexpected behaviors*.

DEFINITIONS

Digital Assurance

Practices, measures, and/or controls applied to digital technologies ...
within a high consequence system, or the system's design, production, or test capability,
in order to ensure **functional** (including performance), **reliability**, and **security**-related
requirements are met **while protecting against potential compromise or subversion ...**
from internal or external sources.

– modified from NNSA SD 452.4-1 Nuclear Enterprise Assurance (NEA) [1/27/2022]

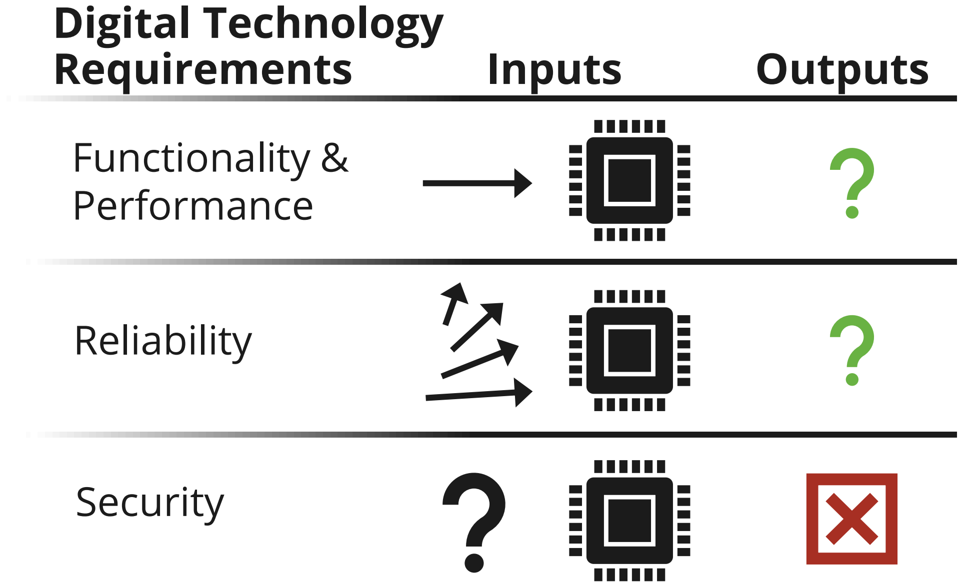

... while protecting against ...

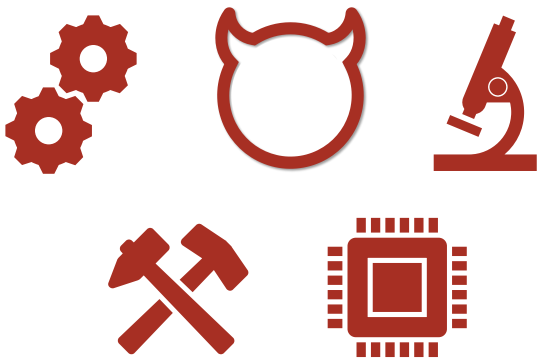

DAHCS MC Concern:
up to one **Insider Threat (single entity)** –
human or digital (e.g., compromised chip or development tool).
Not excluding unexpected behaviors sans adversary.

T&E SCENARIOS

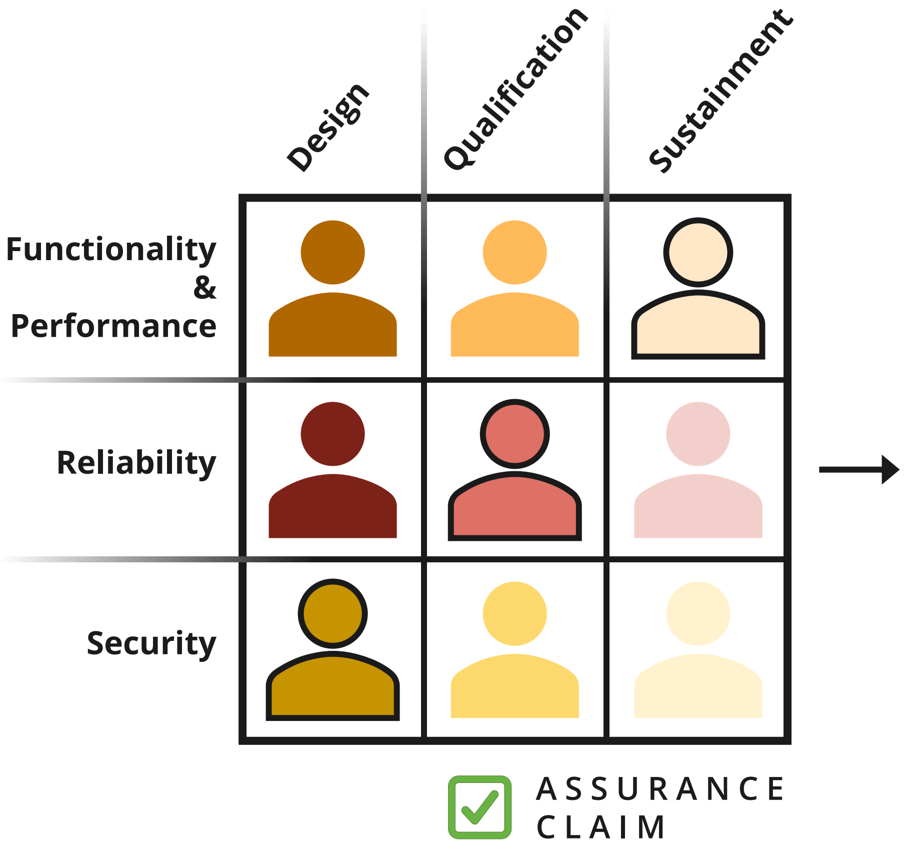

A. **Rapid Reassessment**
   Provide, within two weeks, an updated assurance determination and proposed actions given a technical surprise (e.g., a new threat, a failed test)
B. **Rapid Build**
   Build, within six months, a new controller with requirements altered from a prior design but with as much digital assurance as possible within the timeframe
C. **100% Solution**
   Aim to build, at whatever cost, an entirely cyber-secure, digitally assured controller (we assume this is impossible, but we aim for it)

WHAT IS AN ASSURANCE CASE?

ASSURANCE CLAIM:

Consolidated judgment(s) of adequacy

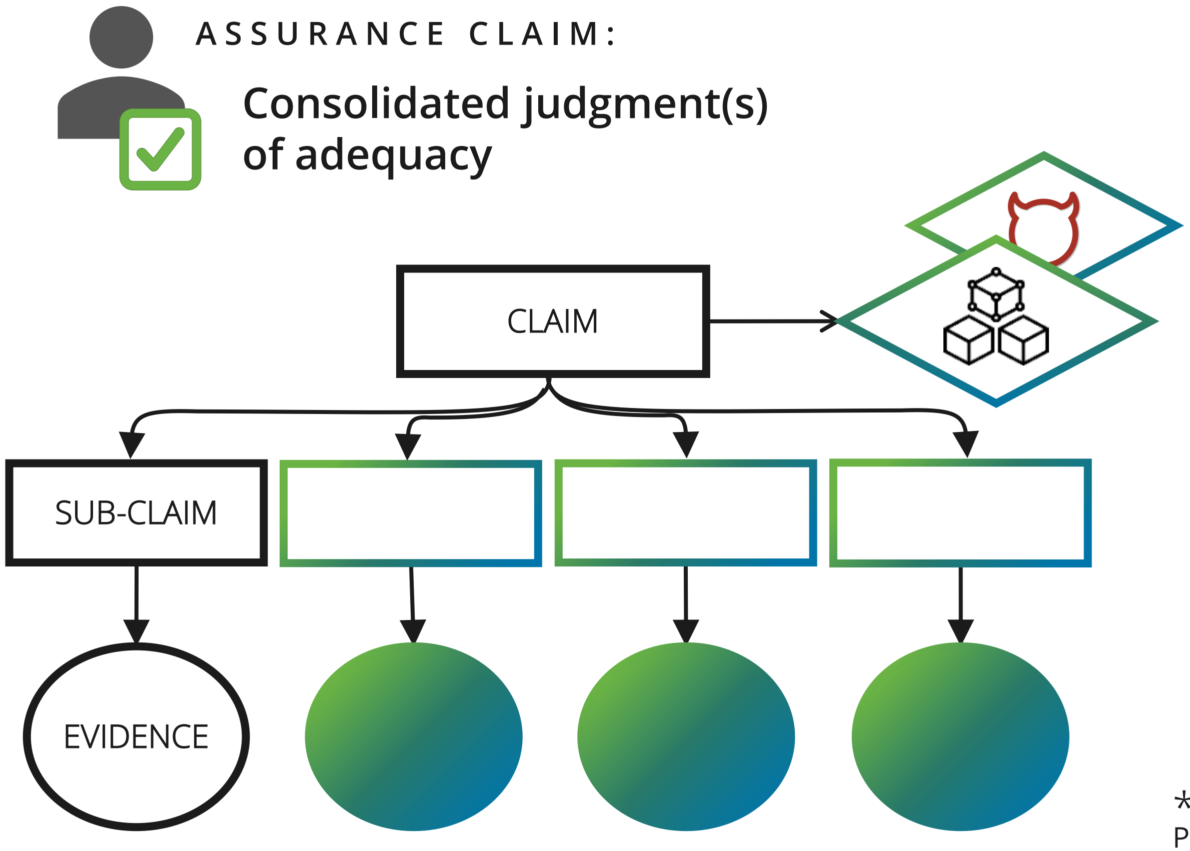

- T&E will use assurance cases* to measure MC progress and focus research
  - across three scenarios and identified HCS *testbed* controllers
  - may test on hidden *validation testbed(s)*
- Assurance cases:
  - developed by safety communities (nuclear power, aerospace)
  - convenient formalism
  - provide structure and organization

*D. J. Rinehart, J. C. Knight and J. Rowanhill, "Current Practices in Constructing and Evaluating Assurance Cases With Applications to Aviation," NASA, 2015.

- ✓ rigorous
- ✓ rapid
- ✓ cost-effective
- ✓ generalizable
- ✓ across system lifecycles

Goal: create an ecosystem that gives us rapid confidence in our systems' digital assurance.

WHAT IS A TESTBED CONTROLLER?

Testbed Controllers:

- “Canonical HCS” testbed controllers, encouraging focus, integration, and generalizability
- Collected by our T&E team to test developed tools, techniques, and methods
- Providing coherent levels of abstractions of a system, so LDRDs can develop **novel strategies** to **support claims** about digital assurance across abstraction levels.

First DAHCS MC Testbed(s): **Software-based state machine application running natively on a microprocessor core on a simple system-on-chip (SoC) with internal & external I/O**

SIDLOC TESTBED
(HTTPS://SIDLOC.ORG)

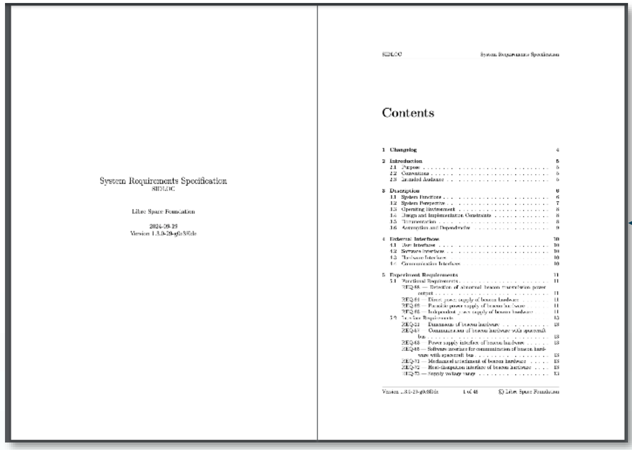

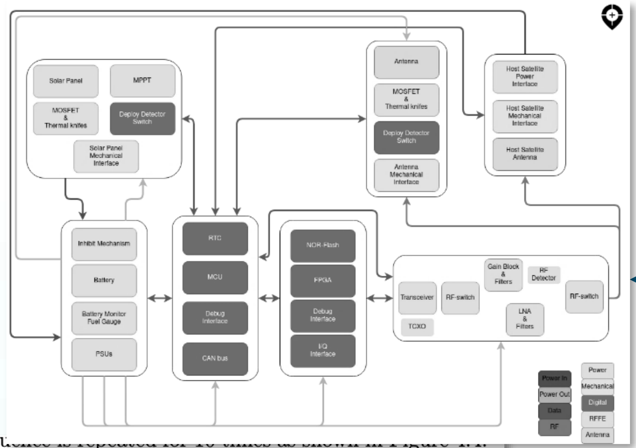

each data bit the Gold sequence is repeated for 10 times as shown in Figure ...
This results to a implementation loss of:

$$Loss_{impl} = 10 \times \log_{10}(10) = 10 \text{ dB}$$

The resulting effective data rate is:

$$R = \frac{\text{Chip Rate}}{\ldots} = \frac{1e6}{\ldots} \approx 48 \text{ bps}$$

```
module ctrl(input sys_clk,
            input rstn,
            input cs_n,
            input din_valid,
            input [7:0] din,
            output reg dout_valid,
            output reg [7:0] dout,
            output led_pin,
            output [1:0] gpio_pins);


  parameter REG_ID    = 0;
  parameter REG_SCRATCH = 1;
  parameter REG_LED = 2;
  parameter REG_GPIO = 3;

  parameter OP_IDLE = 0;
  parameter OP_READ = 1;
  parameter OP_WRITE = 2;


  /* Configuration register space*/
  reg [7:0] id = 8'h37;
  reg [7:0] scratch = 0;
  reg [7:0] led = 8'hFF; //LED is active low
```

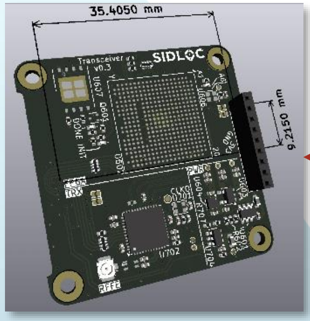

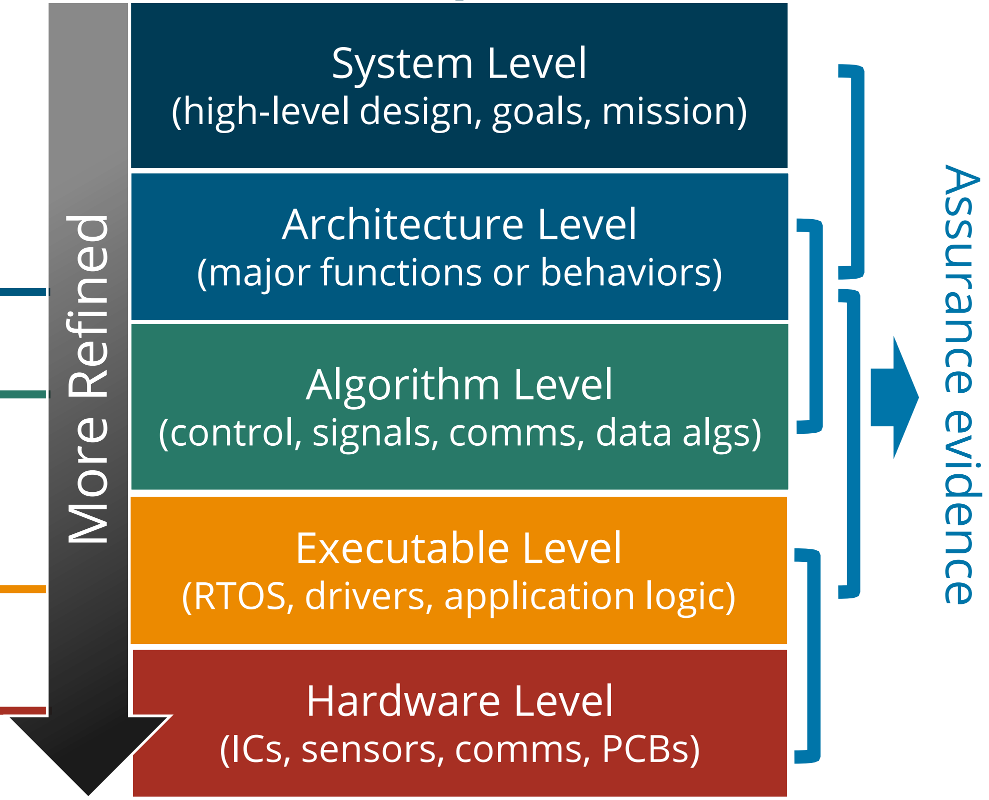

Common levels of abstraction
(design artifacts)

DAHCS MC FY26
CALL FOR LDRDS

FY26 RESEARCH FOCUS AREAS

Assuring Target Hardware and Configuration

Providing evidence that the physical hardware implements the expected digital abstraction

Intelligent Adversary & Hazard Modeling

Explicitly accounting for adversary goals, choices, and capabilities

Failure Consequence Characterization

Enabling end-to-end reasoning about consequences of failures

Digital Composition

Combining evidence across digital technologies… and techniques, abstractions, contexts, etc.

Revolutionary DAHCS

Novel, compelling approaches that meet MC goals in entirely new ways

Targeted Evaluation

Small-scale efforts to address missing T&E evaluation and integration needs

Vulnerability detection, IT systems, existing algorithmic scaling research, and systems of systems are out of scope unless pertaining to DAHCS principles.

THRUSTS	FY25	FY26	FY27	FY28	FY29	FY30	FY31

SCALABLE ANALYSIS

***End State**: Analyses cover large discontinuous state spaces across modern digital technologies*

- Assuring Target Hardware & Configuration
- Behavior Coverage
- Force-Multiplying Expertise

IMPACT ANALYSIS AMID UNCERTAINTY

***End State**: Assessments are rigorously threat- and uncertainty-informed*

- Intelligent Adversary & Hazard Modeling
- Model Inference Given Partial Information
- Failure Consequence Characterization

INTEGRATING WITH SYSTEMS ENGINEERING

***End State**: Humans make system-level trade-offs about digital assurance*

- Digital Composition
- System Assurability Tradeoff Analysis
- Evidence Communication for Decision Support

Revolutionary DAHCS / Targeted Evaluation

DAHCS TESTBEDS

◇ = integrated T&E activities

VALIDATION 1: CRAWL

VALIDATION 2: WALK/RUN

DEVELOPMENT TESTBEDS

IMPACTS

Transition Partners

End State of Roadmap: *Rigorous, efficient solutions for digitally assured high consequence systems*

SCALABLE ANALYSIS

Goal: *Dramatically scale end-to-end DAHCS, seeking at least two orders of magnitude* improvement in time/cost or complexity of handled digital technologies*

- discovering the limits of hardware, state, and input complexity that we can reasonably analyze within given design and resource tradeoffs
- characterizing tradeoffs needed to achieve given levels of digital assurance
- extending and generalizing existing capabilities

End State*: Analyses cover large discontinuous state spaces across modern digital technologies*

* when a baseline exists

Assuring Target Hardware and Configuration

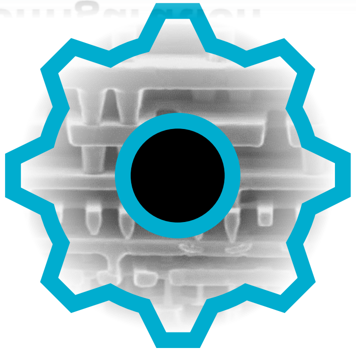

Behavior Coverage

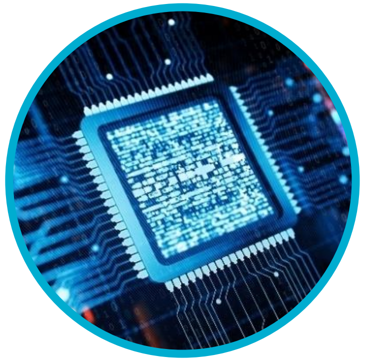

Force-Multiplying Expertise

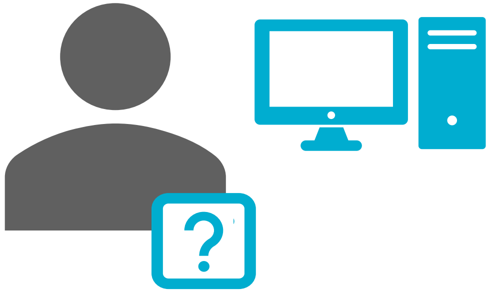

SCALABLE ANALYSIS

Assuring Target Hardware and Configuration

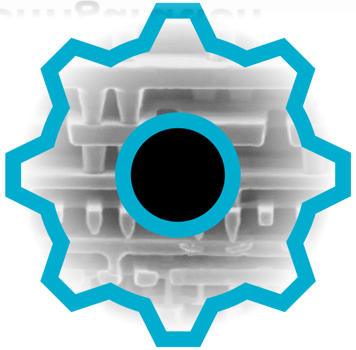

Providing evidence that the physical hardware implements the expected digital abstraction (hardware *logic* is covered in Behavior Coverage)

- revolutionizing ways to obtain digital assurance anywhere along the hardware lifecycle path through holistic, "wild" ideas
- addressing custom ASICs, COTS microelectronics (e.g., FPGAs, CPUs, GPUs, etc.), hybrid solutions, finished PCB assemblies, and critical aspects of a modern, multi-purpose systems-on-a-chip (SOC)

"**Hardware Assurance:** An evidence-supported *level of confidence* that a … device and its configuration do not contain unexpected characteristics or … behaviors …"

- Joint Federated Assurance Center (JFAC)

Goal: *Develop novel, highly scalable approaches, scalability enhancements, and strong measurements to provide a risk-informed level of assurance in the integrity and authenticity of target digital hardware and binary data*

IMPACT ANALYSIS AMID UNCERTAINTY

Goal: *Measure and increase confidence in an assurance case and its evidence, e.g., by identifying what additional information is needed to increase confidence by how much*

- focusing and evaluating assurance cases
- increasing our confidence in them using metrics that do not yet exist
- appropriately allocating our limited resources

End State*: Assessments are rigorously threat- and uncertainty-informed*

Intelligent Adversary and Hazard Modeling

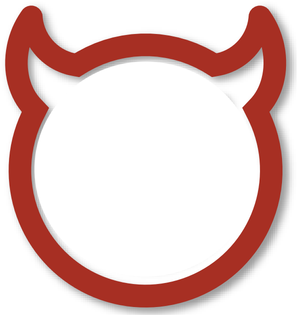

Model Inference Given Partial Information

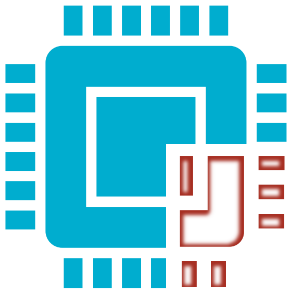

Failure Consequence Characterization

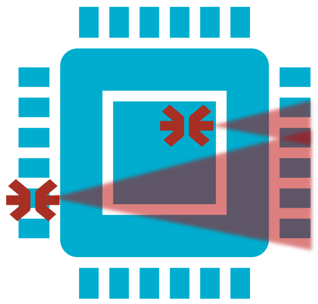

IMPACT ANALYSIS AMID UNCERTAINTY

Intelligent Adversary and Hazard Modeling

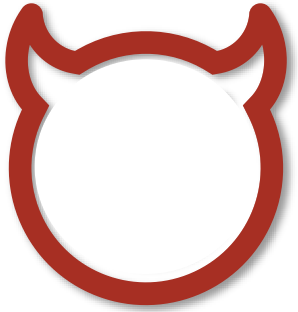

Explicitly accounting for adversary goals, choices, and capabilities

- systematically modeling intelligent adversaries and incorporating adversary models into well-characterized, repeatable, rapid, full-stack digital assurance capabilities
- systematically modeling internal or external hazards that cause system-relevant failures of digital technologies, including failures of assumed digital abstractions
- enabling threat-informed tradeoffs in digital assurance analysis and assurability, including enabling rapid re-evaluation when the threat evolves
- measuring the impact of uncertain or missing threat information on an assurance case

Goal: *Enable rigorously threat-informed digital assurance*

IMPACT ANALYSIS AMID UNCERTAINTY

Failure Consequence Characterization

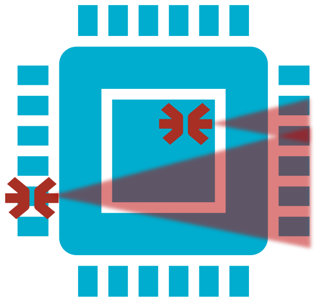

Enabling end-to-end reasoning about consequences of failures

- establishing missing links needed for end-to-end consequence analysis, e.g., by translating between many levels of abstraction
- categorizing or measuring impacts of aberrant behavior
- connecting low-level device effects to high-level system outcomes
- rapidly characterizing *direct*, *indirect*, and *aggregate* consequences
- determining what system-specific information or metrics are needed for robust consequence analysis

Goal: *Develop tools and metrics of rigor for end-to-end reasoning about the impacts of digital technology failure mechanisms on high consequence systems*

INTEGRATING WITH SYSTEMS ENGINEERING

Goal: *Support systems-level decisions about digital assurance and residual risks, including making tradeoffs among digital technologies and digital design options*

- integrating and using digital assurance evidence within systems engineering approaches
- revealing and characterizing *emergent behaviors*
- specifying, understanding, and making effective system-level tradeoffs against digital assurance

End State*: Humans make system-level trade-offs about digital assurance*

Digital Composition

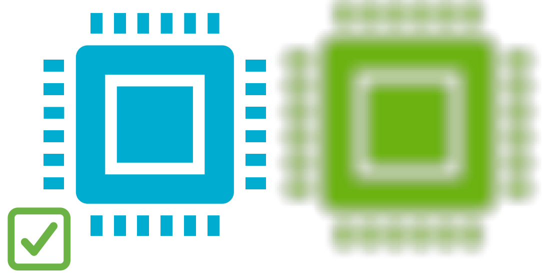

System Assurability Tradeoff Analysis

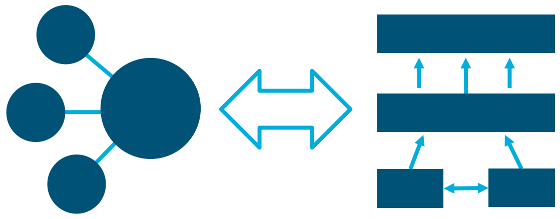

Evidence Communication for Decision Support

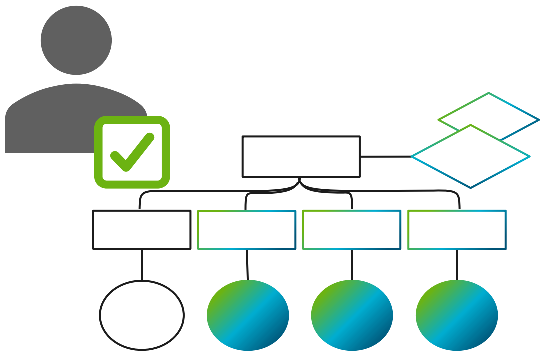

INTEGRATING WITH SYSTEMS ENGINEERING

Digital Composition

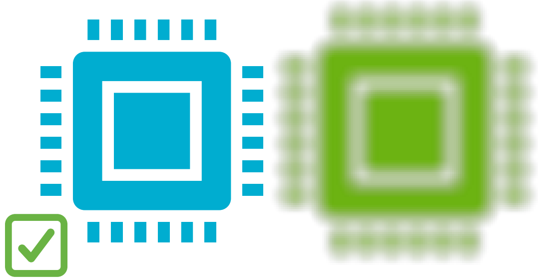

Combining evidence across digital technologies as well as analysis techniques, abstraction levels, and processing contexts

- enabling digital assurance assessments and requirements flow-down derivation across levels of abstraction
- revealing, characterizing, or mitigating *emergent behaviors*
- combining and comparing different types of assurance methods, evidence, and metrics across digital technologies and assessment contexts
- aggregating all assurance evidence into a system-level digital assurance case and a credible argument for a given level of assurance

Goal: *Create methods and metrics to rapidly combine evidence into a digital assurance case and compare options*

REVOLUTIONARY DAHCS

Providing end-to-end digital assurance of HCS through revolutionary approaches that explore ways to think entirely differently about DAHCS

Vulnerability detection, IT systems, systems of systems, and existing algorithmic scaling research are out of scope

Goal: *Approach DAHCS in entirely new ways that meet our needs of scalability, generalizability, integration, and rigor*

TARGETED EVALUATION

Providing rapid, proof-of-feasibility, or baselining of our testbed controllers and/or targeted integration of LDRDs:

1. filling small, applied research gaps for our T&E efforts, enabling the T&E team to demonstrate integration on our *validation testbed* controller(s), or
2. demonstrating integration of DAHCS MC LDRDs into an ecosystem in some other way

***Goal:** Demonstrate integration of DAHCS MC LDRDs into an ecosystem through small, applied research projects*

PROJECT SELECTION CRITERIA

Programmatic Alignment

- **Alignment**: Sandia-unique work advancing DAHCS MC vision and ecosystem, addresses priorities in call, takes risks to increase impact
- **Impact**: outcomes and deliverables impact Sandia, missions, and nation

Science/Innovation (TRL 1-5)

- **Merit**: novel, high-risk, clear, repeatable research advancing TRL or HRL*
- **Feasibility**: aggressive, clear, practical execution plan with **fail-fast** decisions
- **Qualifications & Budget**: reasonable budget, multi-disciplinary dream team

DAHCS MC Alignment

- **HCS Differentiation**: truly unique to HCS generally
- **Test & Evaluation**: strong plan for integrated ecosystem on DAHCS testbeds
- **MC Advances**: generalization, interoperability, scalability, rigor (at least one)

* TRL: https://esto.nasa.gov/trl/ HRL: https://www.osti.gov/biblio/1807329

SUMMARY

rigorous | rapid | cost-effective | generalizable | across system lifecycles

VISION
Disciplined systems engineering processes support systems-level tradeoffs against digital assurance.

Outcomes:

HIGH CONSEQUENCE DIGITAL ASSURANCE
from design to decommissioning

We have the foundation to efficiently and confidently:

- ✓ Characterize digital technologies
- ✓ Assess risks to our systems from digital technologies
- ✓ Select among options that appropriately manage and accept digital risk

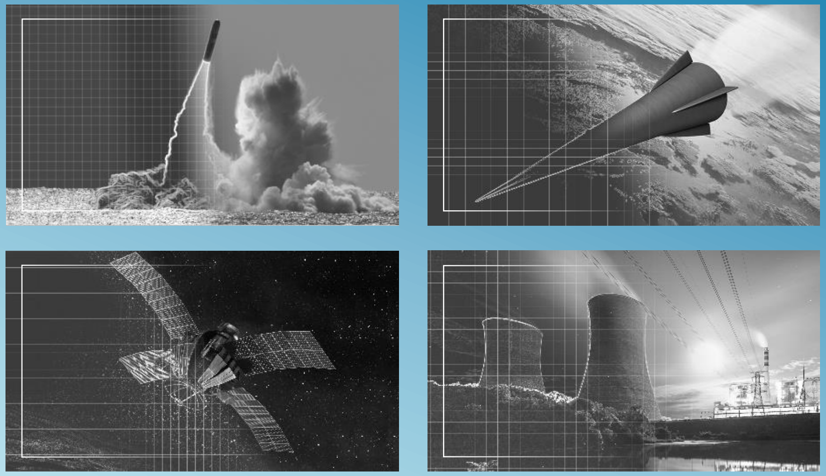

Goal: Create an ecosystem that gives us rapid confidence in our systems' digital assurance.

HOW TO ENGAGE

OPPORTUNITY: LIGHTNING TALKS BASED MATCH MAKING

Purpose: Expand DAHCS MC LDRD academic research partnerships

INFORMATION SESSION
- MiCoP-hosted DAHCS MC Research Call Overview (LDRD).
- Sandia will share slides and the internal research call information after the presentation.

SLIDE/PARAGRAPH DUE
- **Interested professors**, send a single slide or paragraph on your relevant idea to dahcs-micop@sandia.gov.
- Please incorporate a sentence or two outlining the **areas of expertise** in which you are interested in establishing a partnership.

LIGHTNING TALKS
- Professors [+ Sandia PIs] should be prepared to give a **3 min** Lightning Talk during their assigned time slot.
- There will also be time allocated to discuss potential collaborations with potential Sandia proposers.

DEC 11th → Dec - JAN 23rd → Dec - JAN 23rd → Jan 24 – 30th → Jan 30th 11:30am – 1:00pm MST

DEVELOP SLIDE/PARAGRAPH
Develop a single slide or paragraph that will be used for a **3 min** Lightning Talk on **Jan 30th**

SANDIA SELECTION & COORDINATION
Idea submissions will be reviewed and professors [and Sandia PIs] assigned a small-group MS TEAMS meeting by topic.

Actions requested from Academic Partners:

All FY26 DAHCS LDRD Ideas are due **by the Sandia PIs** by end of day on February 10th to the Sandia LDRD Office. Academic connections should be made well before the 10th to be included in the Idea submission, though Sandia PIs may include or expand university partnerships for Ideas selected for the proposal phase.

FY26 RESEARCH FOCUS AREAS

Assuring Target Hardware and Configuration

Providing evidence that the physical hardware implements the expected digital abstraction

Intelligent Adversary & Hazard Modeling

Explicitly accounting for adversary goals, choices, and capabilities

Failure Consequence Characterization

Enabling end-to-end reasoning about consequences of failures

Digital Composition

Combining evidence across digital technologies… and techniques, abstractions, contexts, etc.

Revolutionary DAHCS

Novel, compelling approaches that meet MC goals in entirely new ways

Targeted Evaluation

Small-scale efforts to address missing T&E evaluation and integration needs

Vulnerability detection, IT systems, existing algorithmic scaling research, and systems of systems are out of scope unless pertaining to DAHCS principles.

Questions?

Contact: dahcs-micop@sandia.gov

To receive invites to regular MiCoP discussions, join the mailing list:

Sandia-external: email dahcs-micop@sandia.gov
Sandians: join DAHCS-CoP

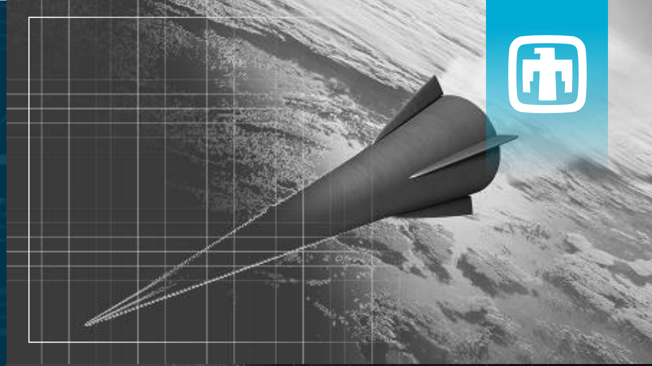

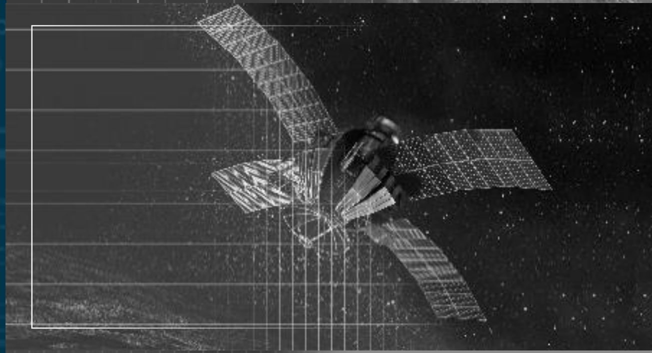

SCALABLE ANALYSIS

Behavior Coverage

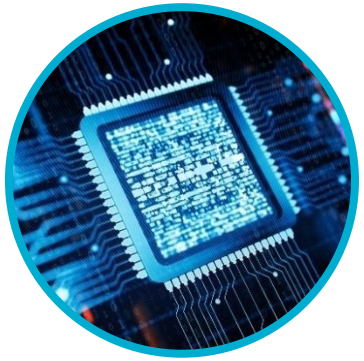

e.g.,

- appropriate abstractions and metrics
- hardware/software co-verification
- multi-fidelity and multi-abstraction analysis

Providing credible evidence that digital technologies (hardware, software, components) behave properly

- specifying requirements for digital technologies unambiguously, including sanity-checking requirements
- characterizing and demonstrating appropriate levels of digital assurance for a given type of requirement / claim
- explicitly expressing, measuring, or deriving assumptions
- rapidly assessing behaviors and updating behavior coverage as the system evolves
- creating analysis techniques that produce evidence of their own correctness, e.g., proof certificates
- dramatically extending the coverage or fidelity of existing analysis techniques

***Goal:** Create new capabilities to provide a risk-informed level of assurance that digital technologies meet their requirements*

SCALABLE ANALYSIS

Force-Multiplying Expertise

Scaling the expertise and human judgment needed for DAHCS by force-multiplying expertise, e.g.:

- optimizing resource allocation in resource-constrained DAHCS efforts that involve human expertise
- tailoring powerful existing computational methods that apply in other domains
- moving DAHCS expertise into automated analysis support and integration
- enhancing human expertise
- creating reusable, widely adoptable pieces of a DAHCS ecosystem, including supporting integration

Goal: *Create new methods that force-multiply expert intuition, analysis, and behaviors, e.g., by translating them into computation*

IMPACT ANALYSIS AMID UNCERTAINTY

Model Inference Given Partial Information

Overcoming obstacles to reasoning about a controller's implementation when relevant design or environment details are incomplete or unreliable, e.g., partial systems

- inspecting black-box digital technologies where we have limited or no insight
- translating from information *describing* digital technologies to a model useful for analysis
- inferring missing information based on interactions with characterized digital technologies
- measuring the impact of uncertain or missing information on an assurance case, including identifying which information or measurement is most needed

Goal: *Develop new approaches that automatically create, tailor, and validate models of digital technologies despite missing information*

INTEGRATING WITH SYSTEMS ENGINEERING

System Assurability Tradeoff Analysis

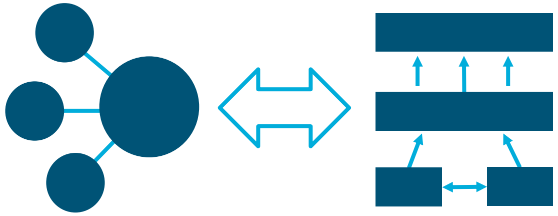

Directly comparing the impacts of implementation choices on digital assurance *as well as* other important characteristics

- directly comparing the impacts of implementation choices on digital assurance
- measuring or grading, at a system roll-up level, a system's digital assurance
- relating a system's digital assurance to other trade-offs like safety, reliability, size, weight, power, cost, or schedule
- relating digital margins to continuous "analog" margins
- relating quantitative metrics like probability to qualitative and opinion-based evaluations

Goal: *Create tools and metrics to explore tradeoffs between digital assurance and other system trade space options*

INTEGRATING WITH SYSTEMS ENGINEERING

Evidence Communication for Decision Support

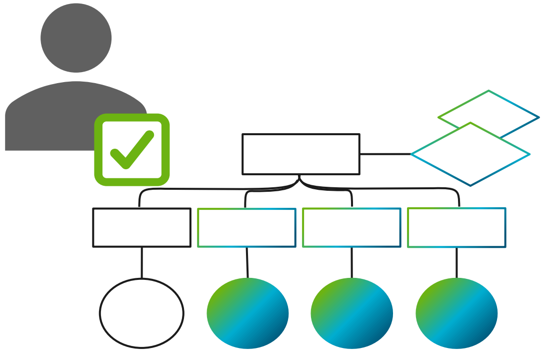

Supporting decision-makers with credible evidence about digital options and assurability tradeoffs

- characterizing and predicting factors that influence decision-making, including how decision-makers trust, interpret, and select among options and impacts
- presenting appropriate, clear evidence that explains complementary and competing alternatives
- selecting appropriate information and presentation options based on risk acceptance criteria
- supporting decision-makers within their own workflows across mission systems, requirement types, and risk management and system lifecycles

Goal: *Provide decision-makers with credible evidence about the functionality, reliability, and security of the system, enabling them to make well-informed tradeoff decisions*