Federal Register Notice: 89 FR 51554, Federal Register :: Networking and Information Technology Research and Development Request for Information on Digital Twins Research and Development, June 18, 2024.

Request for Information on the National Digital Twins R&D Strategic Plan

SRI

SRI Digital Twins Research and Development

SRI is pleased to respond to the National Science Foundation's (NSF) Notice Networking and Information Technology Research and Development Request for Information on Digital Twins Research and Development. As an independent, nonprofit research institute headquartered in Menlo Park, CA with offices across the country, SRI has had the privilege of supporting the federal government for many decades. In the most basic of terms, SRI develops, delivers, and integrates solutions to our society's most pressing problems.

SRI has been working on various aspects of Digital Twins (DTs) in multiple domains, which include engineering design and manufacturing, materials, computer hardware, supply chains, climate modeling, biology, human behavioral modeling, computer human teaming, and social systems. Our history of collaboration across technical and scientific disciplines has provided us unique insights into the broad challenges that need to be addressed by a large government program. The focus of this response is not aimed at narrow or singular ideas but instead at cross cutting technologies.

Terminology: Death of Prototyping

We would like to propose to expand the scope of DTs to include physical systems that do not exist but are in the process of being designed. The focus of *Foundational Research Gaps and Future Directions for DigitalTwins (2024*) and NSF's call is on "operational" use of DTs where the physical artifact already exist. However, the case where the physical product is being designed and built represents a major opportunity that is currently handicapped by our inability to manage complexity. It is also believed that developments in this area will having cross cutting benefits for society in general.

The ultimate goal of DTs in the "design" context is to eliminate physical prototyping and use reasoning over DTs to find unanticipated problems. We realize that this is a far-reaching goal which may never be entirely achieved, but it provides insights in the fundamental technologies that need to be developed that benefit both use cases of DTs. The anticipated impact will be reduced cost and time (>2x) to build first time quality (and qualified) systems. It will also democratize design and manufacturing and put complex products within reach to small and medium sized companies that do not have the expertise or labs to currently compete with the large defense or commercial contractors. Additional motivations described further below.

The ”design” DT use case raises two major questions:

a. How do you create a DT when the physical entity does not exist? Typically, a lot is already known about the components that make up a system (e.g., simulation and math models, empirical models, the laws a physics, engineering and scientific principals), the main challenge is how to ingest and compose this information automatically and incrementally. For example, how would you embed the laws of physics in a neural network to reduce training time for neural nets[1]? If you have two DTs representing subsystems, how would you compose them? This also benefits the “operation” DT use case, because many partial models already exist.

b. How do you find unanticipated problems or emergent behavior in a DT that is composed of multiple subsystems represented as DTs? The aim is to go beyond current uses of DTs that aim to optimize or predict some future state of known quantities. The main reason why we still build physical prototypes is because we lack the computational ability to find unanticipated problems.

The subsequent examples are rooted in design and manufacturing of complex engineered systems. SRI has extensively studied these and gained insights that could be applied to other domains (e.g., medicine, the influence of the biome on the rest of the body). To some degree, engineering design represents the “simplest” case even though it is far from solved despite the claims from certain commercial software vendors.

Example 1. Emergent Behavior: Huygens’ synchronizing clocks

In 1665, Dutch clock builder and scientist Christiaan Huygens observed that when two pendulum clocks are mounted on a common beam, their pendulum swings would eventually synchronize. It took roughly 300 years to explain the phenomena mathematically[2]. Suppose you represent the clocks as DTs, and then “compose them”, how would you discover this emergent behavior if you did not à priori know about this behavior? How would you structure the DTs so that you can compose the models and then find this behavior? There are several branches of mathematics that may show a way forward such as Algebra of Operads and Sheaf theory, however they quickly breakdown when physics gets involved. A lot more fundamental work is needed.

Example 2. Unanticipated Problems: V-22 case study

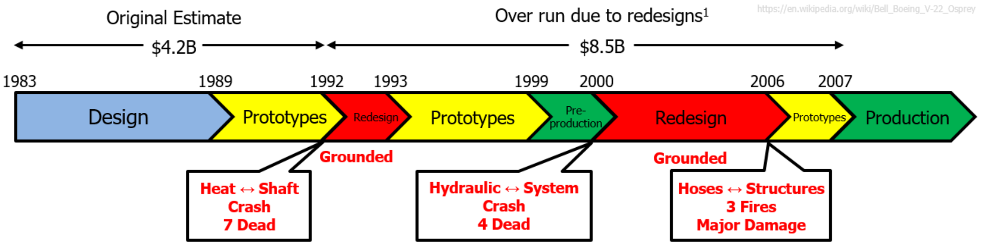

https://en.wikipedia.org/wiki/Bell_Boeing_V-22_Osprey

V22 Production Timeline

The V-22 Osprey tilt rotor aircraft combines vertical takeoff capabilities with fixed wing horizontal flight. The DoD launched the program in 1983 with an original budget of $4.2B, and first delivery expected a decade later. What followed were a series of failures during prototype and pre-production testing causing multiple fatalities, two major redesigns and a 15 year, $8.5B

---

[1] Raissi: Physics Informed Neural Networks, 2019.

[2] https://royalsocietypublishing.org/doi/10.1098/rsos.170777

overrun[3]. The cost of correcting design problems increases exponentially as a product matures in its development cycle[4], which means the earlier a problem can be found the better. The first major unanticipated problem was an engine fire that in turn melted the carbon fiber shaft that connects both rotors to enable single engine operation in case of engine failure. Even though engine failures were anticipated as part of the design, the side effect of melting the shaft was unanticipated. Most of the other major failures were also unanticipated. The V-22 is not unique, most other recent defense programs suffered from similar delays and cost overruns[5]. A RAND[6] report concluded that the growth in complexity and our inability to manage it is one of the main drivers in cost and time overruns.

A way to deal with complex problems is to subdivide them into simpler problems or disciplines. Medicine is subdivided in orthopedics, cardiology, neurology, gastroenterology, and so on. Likewise, airplane design is subdivided into many subdisciplines: aero, structures, propulsion, hydraulics, electrical, stability and controls and many more. Each of the disciplines painstakingly create math or empirical models in 100s of disparate software systems to capture their designs, often representing 10s-1000s of FTEs of effort, each.

Integration, meaning determining whether everything will fit and work together, is predominantly done manually during integration design reviews and by performing limited simulation and bench studies. Combining these disparate models into a full DT representation of the product is currently impossible because the underlying software cannot handle the complexity and scale, while the amount of effort required to combine the models is prohibitive. A Boeing 777 contains about 5M parts, while current geometry Computer Aided Design (CAD) systems can only work with ~250k objects during a session[7]. There are many other challenges, but the consequence is that complex product design is not integrated and heavily relies on physical prototyping and expertise that are only in the hands of a few companies. However, the growth in complexity has outpaced our ability to find unanticipated problems using our current tools and methodologies. Similarly, if an ailment in a human crosses multiple disciplines it often becomes exponentially difficult to diagnose because of lack of "integration" between the disciplines. Medicine by and large lack an integrated approach to diagnostics[8].

The remainder of this RFI response focuses on Artificial Intelligence.

The Vision, Part 1: "DT Compiler"

Creating DTs ab initio is too labor intensive and expensive (current practice). The vision is to develop a "DT compiler" that will automate ingestion of the available information[9], and build

---

[3] GAO-09-692T: Published: Jun 23, 2009.

[4] https://www.qualitydigest.com/feb00/html/design.html

[5] Norm Augustine's 16th law

[6] RAND: https://www.rand.org/pubs/monographs/MG696.html

[7] *Conversation with Burhop @ Siemens*

[8] Beauchamp et al, "Integrative diagnostics: the time is now—a report from the International Society for Strategic Studies in Radiology", National Library of Medicine, Dec 2023.

[9] DARPA Perceptual Task Guidance is focused on heterogeneous ingestion and composition of information

DTs with several orders of magnitude reduction of effort. This vision is inspired by DARPA's Computable Models program which demonstrated a 100-1000x reduction in effort to create physics simulators with equal or better accuracy using the notion of a compiler[10].

The major challenges to be addressed are:

- **Ingestion:** There is lots of data in many forms, how can they be ingested to build DTs? SRI insights: DARPA AI Research Assistant (AIRA); Perceptual Task Guidance (PTG) and Computational Cultural Understanding (CCU) – which ingest and fuse data from heterogenous sources. Other insights: DARPA Big Mechanisms that read cancer texts and generated a system of ordinary differential equations to model cancer pathways, i.e., Large Language Models are not always the answer.
- **Embedding Priors / Small Data:** We know the laws of physics and other scientific principles; how can they be built into the structure of neural nets? This eis specially critical for areas with little data to reduce training effort (e.g., engineering, medicine). SRI insight: DARPA Physics Informed AI (PIA)
- **Composition of DTs:** Physical objects interact with each other in non-intuitive ways across multiple physics. How can DTs' representations encapsulate this? SRI insights: DARPA CompMods with algebra of operads

The vision, part 2: Finding Unanticipated Interactions

The vision is that once you have a DT representation of a system, to search through all the subsystems represented by DTs and find unanticipated problems that would occur during operations. For example, if you route wires too close to a turbo charger, they may melt if they are too close because turbo-chargers get hot. This requires spatial awareness between the DTs, the side effects that each DT has (heat, vibrations, noise, EM fields) and the specifications (e.g., material properties, in this case insulation melting temperature).

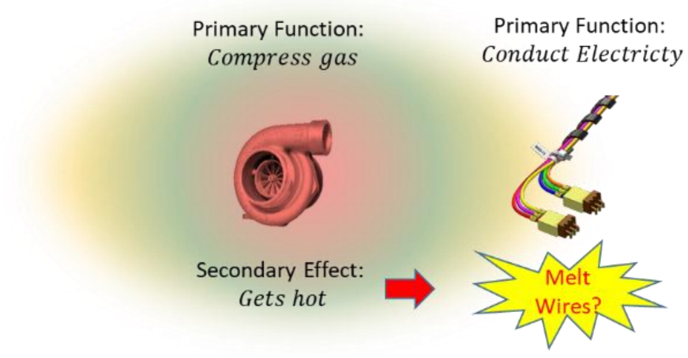

Unintended turbocharger interaction due to secondary effects.

---

[10] Pietrzyk et al., Automated upscaling via symbolic computing for thermal runaway analysis in Li-ion battery modules, JOCS, 2023

Some of the major challenges to address are:

- **Search**: Finding all potential interactions is currently prohibitively large. Using de Weck's formula[11], a Boeing 777 would have potentially 10^{10} interactions, which would require centuries to evaluate. On the other hand, AlphaGo beat a human even though the search space is 10^{170}. However, there are many ways to cull the search space because physics usually has a limited sphere of influence. The challenge is to autogenerate these "sphere of influences" from the DT descriptions.
- **Quantifying Interactions**: One direction of research is to develop the mathematics to describe emergent behavior as described above under Huygens. Another research direction is to leverage Qualitative Reasoning[12] (invented at PARC) which mimics how humans reason about physics without performing any detailed computations. The fall back is brute force simulation, which requires automatically generating new simulation tests without human interference. This requires additional technology development though inspiration may come from the gaming and animation industry with alternative methods to create simulators on demand (e.g., particle methods, cellular automata).
- **Ageing & uncertainty**: Nothing is manufactured perfectly; things age and wear out. This needs to be incorporated into DTs for either usage case (design & operations) to enable accurate predictions. SRI insight: DARPA META: Fault-Augmented Modelica Extension [13] which incorporated faults and wear in models

Facsimile of Intelligent Lifeforms (FOIL)

SRI is developing the Facsimile of Intelligent Lifeforms (FOIL) system to create and run cognitive simulacra of humans (digital twins) teaming with each other and AI agents in a task simulation environment. We will create twins for specific tasks and task-roles that can be adjusted to reflect variations in human factors such as trust and communicativity that are known to influence the effectiveness of human-AI teams (HATs). We will develop techniques to promote the fidelity of twins to humans by training from observed human task performance. More importantly, we will develop techniques that minimize the amount of human data required for training twins: the purpose of running simulations with twins is to reduce the need for human trials, which is defeated if numerous trials are needed to create twins in the first place.

We will validate our techniques for twin training with minimal data in a restricted setting and then show how it generalizes to other settings. The restricted setting will be a text-only version of the collaborative board game Pandemic. In Pandemic, between two and six players are assigned different roles, abilities, and resources, and must work together as a disease-fighting team to stop the spread of deadly viruses. It is known for its challenging game play and the need for players to work as a team to avoid quickly losing the game. Although the domain is relatively simple, it

---

[11] De Weck @ MIT: https://www.ssse.ch/sites/default/files/evt_files/SWISSED-2018-de-Weck.pdf

[12] Williams & De Kleer, "Qualitative reasoning about physical systems", AI, 1991.

[13] Honda et al., "A Simulation and Modeling Based Reliability Requirement Assessment Methodology", ISIS, 2014

provides team-working complexity comparable to that required for real emergency response tasks. It is also a domain that already has digital environments and some AI players.

Coevolutionary Modelling of Emergent Teams (COMET)

SRI is interested in developing Coevolutionary Modelling of Emergent Teams (COMET) to build digital twins of a human-AI team (HAT). These digital twins abstract the operational environment while ensuring a rich interaction space between commander and an adaptive AI teammate. Figure 1 illustrates the capabilities of a COMET digital twin. A generative human simulation model is used to sample diverse behaviors. The HAT simulator directly models teams that emerge after adaptation (yellow rectangles) to the team-mate in a task (curved arrows). The underlying high-dimensional behavior embedding is trained to possess the property: the Euclidean distance between human and AI behaviors becomes smaller over the course of team adaptation. COMET enables realistic simulation of the HAT on novel tasks by simply sampling locally in this embedding space. COMET implements metrics to assess the real-world fidelity of the simulation and quantifies the nature of the HAT's interactions. The fidelity will be validated with multiple humans in-the-loop (HITL).

Thank you for the opportunity to respond to the National Science Foundation's (NSF) Notice Networking and Information Technology Research and Development Request for Information on Digital Twins Research and Development.

SRI POC: Christina Hildebidle, SRI