

CYBER SECURITY CONSEQUENCE PRIORITIZED DESIGN VIA SIMULATION

Andrew S. Hahn*, Adam J. Beauchaine, Titus Gray
Sandia National Laboratories, PO Box 5800 MS 0748, Albuquerque, 87185-0748,
**ashahn@sandia.gov*

Abstract

Cybersecurity of operational technologies has been exceedingly difficult to demonstrate with repeatable and meaningful accuracy. For the nuclear industry this uncertainty in cybersecurity effectiveness and its assurance increases regulatory hurdles and has the potential to result in costly implementations. Robust, repeatable, and systematic cybersecurity analysis essential for assisting engineering, design and development as well as to inform sufficient and efficient cybersecurity protections. Many cybersecurity analyses evaluate the system and implementation in the context of attack pathways, vectors and vulnerabilities (i.e., the “penetrate and patch” approach), lacking determinism and limited inclusion of diverse novel attack methods. The Advanced Reactor Cyber Analysis and Development Environment (ARCADE) is intended to provide repeatable and systematic cybersecurity analysis, as well as a development platform to evaluate the efficacy of cybersecurity implementations and methodologies. The ARCADE platform allows plug-and-play operation for common industry simulation tools (e.g., Flownex, Simulink), while providing a generic API for custom AR simulation environments. The integration of high-fidelity physics with software defined network emulations, enables consequence-focused design, a key principle in Cyber Informed Engineering. This paper will describe the development of ARCADE and the use of this platform in evaluating and providing assurance of cybersecurity engineered controls of an AR control system. Finally, future work and development of ARCADE to address the nuclear industry’s needs will be discussed.

1. INTRODUCTION

Operational Technology (OT) cybersecurity is a maturing field which has primarily relied on Information Technology (IT) practices to secure critical infrastructure. This leaves significant gaps in how risk is treated and evaluated, as IT consequences pale in comparison to OT consequences. Additionally, the cyber risk space is practically infinite as cyber threats match pace with security designs. What is needed is an engineering approach which is based on solid principals and backed by analytical methods which enables verifiable security by design. The Tiered Cybersecurity Analysis (TCA) seeks to answer this need, but it requires a system which can provide analysis driven prioritization of system function cyber risk. The Advanced Reactor Cyber Analysis and Development Environment (ARCADE) can provide this prioritization, which will serve the TCA as well as Cyber Informed Engineering (CIE) as a key tool for cyber design decisions. This paper will document the results of using ARCADE for this prioritization work as well as brief discussions of the TCA and ARCADE’s construction.

2. TIERED CYBERSECURITY ANALYSIS

Ultimately the goal of the TCA is to reduce the cost and complexity of cybersecurity implementation on OT systems while maintaining the highest security profile reasonably achievable. The TCA supplies a framework which enables risk informed and analytically supported engineering practices to be implemented in the cybersecurity domain. Utilizing risk grading through ARCADE can enable better design decision making, and right-sizing security implementations by focusing protection on systems which have demonstrable safety impact. The TCA and the automated tools to assist designers are currently being developed by Sandia to reduce their cybersecurity design burden and provide regulators with sufficient analytical evidence of design security.

The TCA is comprised of three layers outlined below in figure 1. Tier 1 (Design Analysis) explores the usage of secure by design (SeBD) features to mitigate attacks using physical plant design and seeks to measure what is possible on a given control system design. The analysis assumes that the adversary is omnipotent, omniscient, and all powerful in the digital domain. Tier 2 (Access Prevention) addresses digital domain risk pathways through Defensive Cybersecurity Architecture (DCSA) to deny adversary access to exploitable plant functions. It is assumed in Tier 2 that the adversary is bounded by the architecture of the network. Tier 3 (Denial of Task) involves the prevention of specific tasks and workflows an adversary needs to complete in order to produce a favourable outcome. Tier 3 may be viewed as a more “classical” system security, as it assumes the

adversary is limited by network and system architectures. These mitigation efforts are focused on attack vectors not adequately addressed in the preceding tiers, including trust management of plant hardware and supply chain monitoring. We further detail and provide example scenario information for each of these tiers in the subsequent sections.

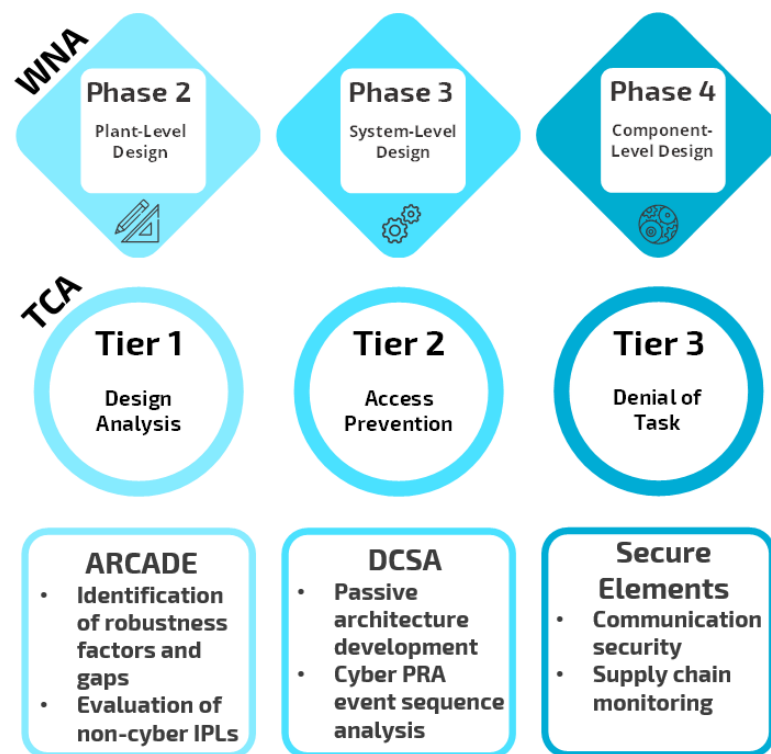


Figure 1: TCA Tier Outlines with Corresponding Implementation Efforts

2.1. Tier 1 (Design Analysis)

Tier 1 analyzes and evaluates AR plant design to identify elements that may be considered “secure by design.” To apply this label, plant design elements must be capable of fully eliminating physical risk from an attack within the digital domain. By allowing an adversary full, unrestricted access to a plant’s digital systems in tier 1 scenarios, plant designers may identify design features that would succeed in delaying or preventing all manner of digital attacks. These robustness factors inherently prevent the adversary from transitioning the state of the system from normal operational conditions to accident conditions. Given the lack of cybersecurity integration into many AR plant simulation models, work at this tier is primarily concerned with the development and subsequent fusion of technologies that allow for the incorporation of cybersecurity into plant design tools. Attacks that may be fully mitigated in this domain are removed from consideration in subsequent tier labels.

The goal of plant function evaluation is already common in AR design, but existing models often lack incorporation of cybersecurity attack scenarios into the design process. To give an example of such a scenario, consider an attack that succeeds in disabling a forced cooling system within in an AR, leading to a situation in which rising fuel temperatures present an un-mitigatable risk for a plant facility. Despite an assumed complete loss of control over the digital domain, a plant may still successfully prevent this attack through the usage of formally verifiable robustness factors. TRISO fuel provides a robustness factor in temperature control to offer increased resilience to digital domain attacks which increase fuel temperature. Should TRISO fuel be leveraged in conjunction with non-digital passive cooling systems such a Reactor Cavity Cooling System (RCCS), the proposed attack could be fully mitigated across all scenarios involving an attack on an AR cooling system, regardless of that system’s digital state.

The usage of RCCS in the previous example is assumed to be entirely disconnected from the digital domain of a plant system but does not qualify as an inherent safety feature and thus is not a robustness factor. Any components whose functions or workflows are rendered immutable by an attacker but are also not inherently

secure/safe may be considered a non-cyber Independent Protection Layer (IPL). These non-cyber IPLs are key design features which enable designers to cover gaps in the inherent robustness factors that would otherwise allow state transitions out of the safe operating envelop. ARCADE maps these robustness factors and verifies the protection from non-cyber IPLs to ensure that SeDB efforts are focused on gap in these features which adversaries may exploit. These gaps are not expected to be entirely eliminated, there will remain some which require Tier 2 pathway analysis to sufficiently protect the system.

2.2. Tier 2 (Access Prevention)

For cyber threats that cannot be sufficiently mitigated in Tier 1, Tier 2 analyzes system attack vectors provided by ARCADE and generates architectures designed at stopping digital domain attacks. Tier 1 should have ensured that a design has no single functions which could cause catastrophic failure, requiring adversaries to attack multiple systems. This tier assumes an attacker may complete their goals if given access to multiple plant components or functions from which multiple compromises could allow a transition out of the safe operating envelop. The adversary is assumed to be bound by the network architecture, e.g., if a given system component is non-networked, or in a separate security domain, it is assumed to be increasingly more difficult to compromise the required systems to cause consequence. Prior work has modeled attack behavior using probabilistic risk assessment trees [1], with results of this analysis being leveraged in the construction of Defensive Cyber Security Architectures (DCSA) for AR systems.

Core to the establishment of DCSA is the notion of security zones as defined in Nuclear Security Series (NSS) publication 17-T [2]. Zones may be defined as “a logical and/or physical grouping of digital assets that are assigned to the same computer security level and that share common computer security requirements owing to inherent properties of the systems or their connections to other systems” [3]. Figure 2 details the usage of zones as a DCSA enabler, with zones and levels being defined based on the safety classification and interoperation requirements of the plant functions they perform. Inter-zone communication is considered trustworthy for all endpoints within a given zone, with strict isolation controls for each zone topology. Privilege escalation is required to move data between zones, creating a model in which multiple zone compromises are necessary to perform an attack when leveraging DCSA.

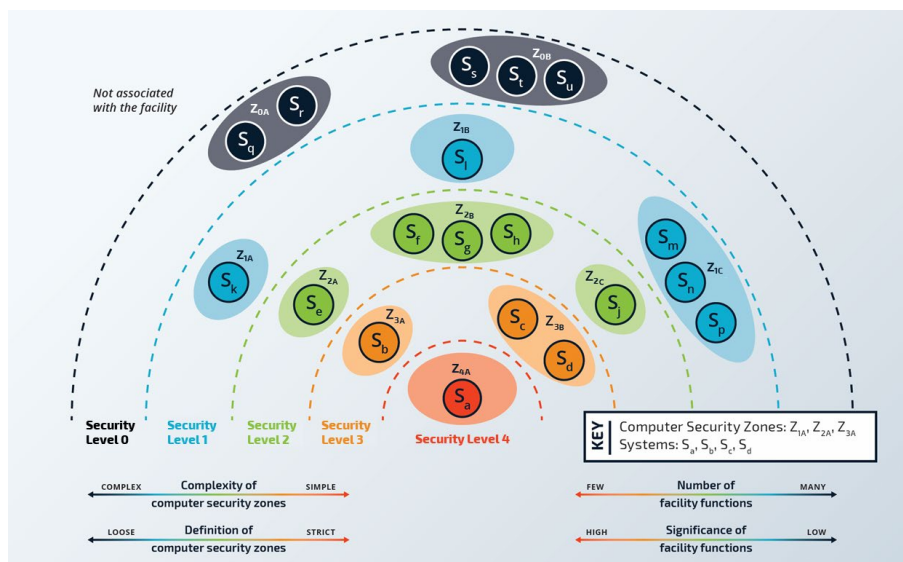


Figure 2: Tier 2 System Zone Diagram

2.3. Tier 3 (Denial of Task)

Tier 3 defines the cybersecurity monitoring and protection requirements of the system for threats that cannot be sufficiently addressed in the previous tiers. The purpose of this tier is to prevent the adversary from performing the necessary tasks required to cause a consequence to the system. Threat actors are assumed to leverage sophisticated strategies, including zero-day exploits of plant hardware to obtain a favorable result. Risk

may be defined similarly to the PRA strategy used in tier 2 or established using formal methods for hardware security. Prior work in this space has included the trust and supply chain management of AR hardware, including the usage of smart cards [4] and trusted platform module (TPM) for session key management. Through the layering of cryptosystems leveraging trusted hardware, tier 3 analysis aims to mitigate all remaining threats unaddressed in the previous 2 tiers.

3. ARCADE PLATFORM DESIGN & CONSTRUCTION

System designers can often struggle to understand and quantify security risks within OT environments for AR systems. This struggle comes primarily from existing analytical and risk-assessment tools lacking cyber event integration, leading mitigative strategies to be either speculative and based on theoretical research material, or, borrowed from other types of OT systems altogether [3]. While formal methods allow for modelling of plant safety conditions the complex nature of cyber threats extend beyond the typical safety analysis. Safety analysis is typically bounded by event probabilities and Subject Matter Expert (SME) contextual pathways to failure. Cyber threats can initiate multiple highly improbable Unsafe Control Actions (UCA) in combinations that are outside of the expected pathways analysed through the safety lens.

These challenges are continuously present for AR system designers, who need a manner of incorporating cybersecurity events into traditional analytical tools for AR design. They need a usable way of determining SeBD features in plant construction that may be leveraged in the prevention of cyber-attacks. This usability should be derived from compatibility with existing verified tools to allow for trustworthy results. Finally, they need a method of verification of system results, for determining if a bounded set of scenarios is appropriate for security assurance.

ARCADE aims to meet these needs for AR system designers. In this section, we detail the design and user workflows of ARCADE, and provide system architecture information as well as individual component descriptions. We additionally detail integration with existing tools and architectures, as well as result gathering techniques. Further results are included in section 5.

3.4. ARCADE Design

The Advanced Reactor Cyber Analysis and Development Environment (ARCADE) is the foundation of the first tier of TCA, with the goal of identifying, verifying, and mapping robustness factors and non-cyber IPL features in modern AR systems. ARCADE leverages compatibility with industry standard simulation tools such as Flownex® and Simulink in tandem with a virtualized, interactable environment, accessible via a generic API. In this subsection, we detail the system architecture of ARCADE, as well as individual component designs. At its core, ARCADE may be viewed as a fusion between physics simulation engines, and a cybersecurity sandbox architecture for deploying attacks. ARCADE allows for attackers to have full access to a plant's digital systems within this sandbox and is thus ideal for Tier 1 TCA attack simulations. A simplified architecture of ARCADE may be visualized below in Figure 3.

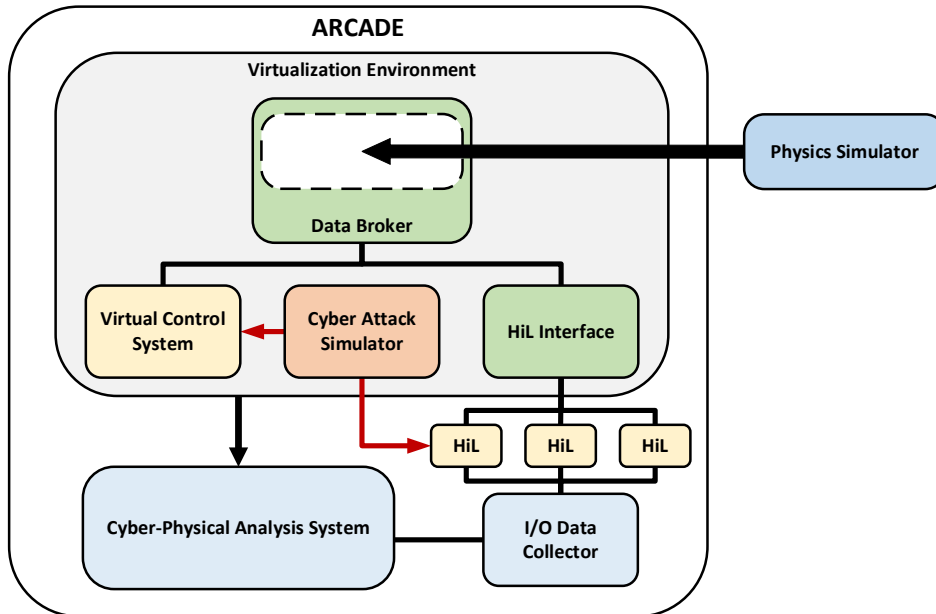


Figure 3: Simplified ARCADE Architecture

The core functionality of ARCADE allows for the gathering of outputs and exposing inputs from a chosen physics simulation engine during runtime, and the sharing these I/O with downstream system components to perform cyber-attack simulations. ARCADE components are predominantly hosted on a virtualized OT topology known as the ARCADE Virtualization Environment. ARCADE additionally supports hardware in the loop (HiL) devices as targets for simulation integration. Attacks are performed via a dedicated system known as the Cyber Attack Simulator and may be started dynamically at any point throughout a system workflow. The resulting changes in output due to any performed attacks are recorded using a data collection model and formatted for human readability. These components allow for far more in-depth investigations of potential cyber-attack impacts on AR systems. We describe each of these components and their interoperability inline.

3.4.1. Data Broker

Physics simulation engines such as Simulink and Flownex are currently being utilized by AR designers for real-time simulation of AR control systems to develop plants control functions. The Data Broker component of ARCADE seeks to bridge the I/O of these existing control system simulators with a fully emulated OT environment. To support these goals, the Data Broker consists of two core functions: the exchange of data from the physics simulator, and the propagation of this data connection across a simulation environment. Given the real-time analytical and attack deployment capabilities of ARCADE, both functions are run concurrently on a single multithreaded C program and designed with performance and scalability as primary concerns. Figure 4 displays an in-depth view of Data Broker functionality, and the methods used to interconnect different components.

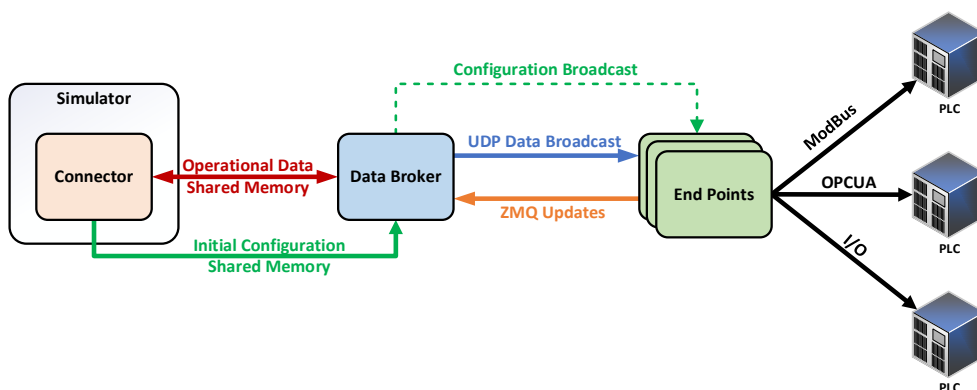


Figure 4: Data Broker Components and Interconnectivity

Modern physics simulation tools such as Flownex® allow for the integration of custom C# scripts with the simulation runtime. Simulink offers similar functionality with its ability to incorporate CMEX code into simulations. This real-time programmability is crucial for the data exchange function of the Data Broker. Using these custom scripts in the simulation engines allows the Data Broker and simulator to establish a shared memory object for the rapid exchange of live physics and control system actuation data. Due to kernel differences in memory management, as well as simulator compatibility, separate versions of the Data Broker were developed for Windows and Linux systems, with Windows being the current version used for testing.

The shared memory structure of the Data Broker and simulator currently requires them to inhabit the same computational environment, however, given scalability concerns, it is crucial that ARCADE is not limited by the performance of a single environment. The usage of “Endpoints” helps alleviate this burden of scalability by distributing the workload of updating and managing individual OT system components. Aside from an initial configuration broadcast, the Data Broker continuously sends updated simulation outputs to all connected endpoint devices. These messages are short in length and uniform in structure, allowing for rapid parsing and updating of OT devices. The updating functions of Endpoints are enabled by Modbus, OPC, and other industry standard protocols for OT system connectivity. Endpoints additionally return any requested data to the Data Broker in the form of a ZMQ socket connection. The combination of these components allows for the asynchronous transfer of incoming and outgoing Data Broker messages.

3.4.2. Virtualization Environment

ARCADE leverages a virtualization environment to provision operating environments for all system components. The usage of virtual machines for environment infrastructure further increases system scalability and reduces deployment costs. Minimega was selected as an environment provider for its rapid deployment capabilities and data capturing tools [5]. Minimega allows for the real time analysis and recording of network traffic, endpoint file activity, and scenarios that can be replayed through environment setup scripting. Minimega may also be used to design different network topological schemes and simulate different physical layer connection settings, both of which may be incorporated in future experimental work of ARCADE.

Minimega interacts directly with the KVM hypervisor and QEMU emulation platform, allowing for compatibility with most Debian-based Linux systems. Minimega requires no external software stack or complex initial configuration, allowing for fast, efficient deployment of a wide array of experimental testbed scenarios. OpenVSwitch is leveraged for an internal switching stack, and VMbetter allows for the export of virtual hosts into many common disk image formats. Protonuke, a simple layer 3 traffic generation module, allows for diverse network conditions for experimentation. These features allow for the ease of recording and exporting scientific results, which in turn allows for the replicability of experimental procedures.

The only current limitation on this strategy is the availability of system images and hardware emulation for certain devices. Because the Data Broker and Endpoints are designed using standard enterprise operating systems, this issue is entirely limited to OT component emulation (controller devices). In such instances, users will need to employ an HIL approach to system simulation. Minimega has supported HIL systems in the past and may still be used to provision the rest of the system in an HIL scenario.

3.4.3. Cyber Attack Simulator

Cyber-attacks in ARCADE are simulated via ManiPIO, a programmable scripting tool that leverages Python’s Modbus library to alter values on running controller systems [6]. ManiPIO is currently limited to the Modbus protocol, as well as network connectivity requirements to target controller systems. While the programmability of ManiPIO allows for the orchestration of complex cyber-attacks involving network exposed memory register values, certain attacks are beyond ManiPIO’s capabilities intentionally to prevent its use as a cyber threat. Such attacks involving non-network exposed data or alternative protocol schemes would require an extension of ManiPIO or direct implementation on an OT device. This work focuses on an initial set of cyber-attack data, implemented only using ManiPIO scripting.

When considering the goals of ARCADE as an evaluation system for Tier 1 of the TCA, determining system functions as being SeBD is a challenging task. The evaluative testing space includes all known cyber-attacks, which is an infinite language of programs and inputs to be run and verified. ARCADE requires a reduction in problem search space, so that formal verification of plant design features is practically possible. System Theoretic Process Analysis (STPA) provides such a reduction. By limiting analyzed attacks to STPA dynamic control actions, we may produce a list of unsafe control actions (UCAs) that may be evaluated to provide verification of system safety. These UCAs are grouped into 4 general categories as shown in Figure 5.

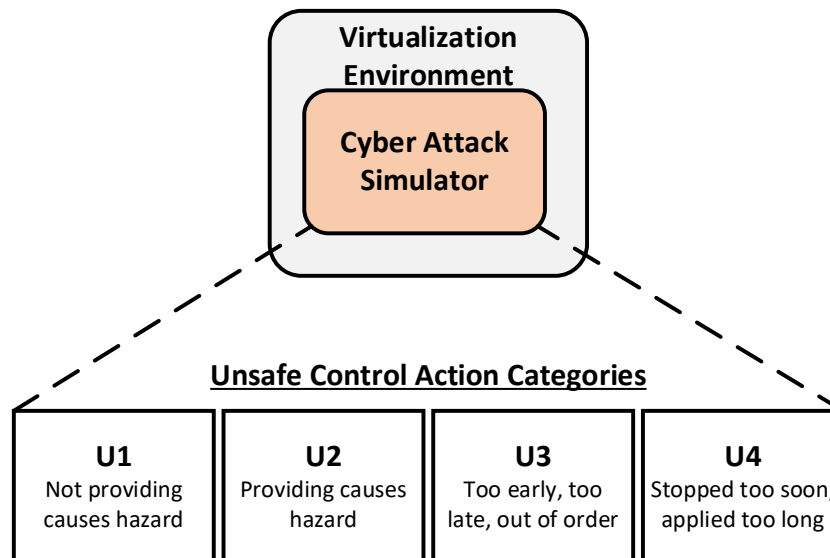


Figure 5: Unsafe Control Actions

U1 and U2 actions are trivial to implement via ManiPIO scripting, whereas U3 and U4 require additional considerations. Both require knowledge of “expected” system behavior and a method to apply controls to produce an “unexpected” result. The usage of system memory is likely to be required for both UCA evaluations to produce a baseline for system behavior comparisons and will also likely have to be implemented directly on an OT device. This paper is centered on results for U1 and U2 UCAs implemented via ManiPIO as U3 and U4 are current subjects of development.

3.4.4. Cyber Physical Analysis System & Result Gathering

Even with reductions in problems search space, ARCADE requires a high amount of analytical testing before being able to produce actionable results. The Cyber Physical Analysis System is constructed using the Sandia developed Dakota parametric analysis program [7], and allows for rapid testing and verification of large input search spaces for cyber-attacks. Parallelism is leveraged to provide results in a time-efficient manner, and Dakota itself may orchestrate the evaluation process, and analyze the resulting reports. Figure 6 shows the data collection and evaluation workflow when using Dakota.

Data collection is currently implemented directly within the ARCADE Data Broker. As shared memory values are captured, a separate data collection thread appends selected result data from the physics simulator and the Endpoint responses to a CSV file. When a run is completed, this file is formatted to be human readable, and may be exported to Dakota for further analysis. The flexible nature of CSV formatted files allows for the analysis and comparison of individual runs, as well as unionization and concatenation of result data, if required for post-run analysis. Testing procedures involve similarity analysis between nominal cyber-attack datasets. Additional proposed testing procedures are further detailed in section 4.

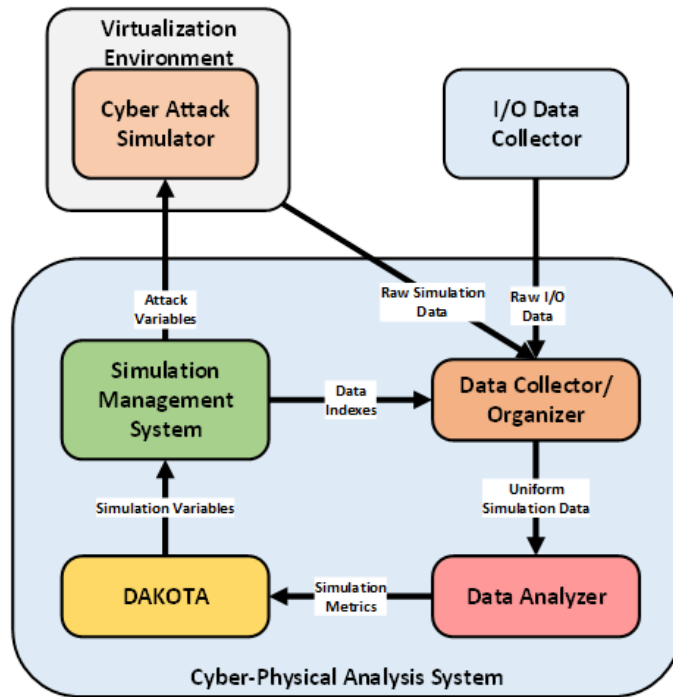


Figure 6: Cyber Physical Analysis System Diagram

4. PROBLEM SPACE ANALYSIS & PROPOSED INTEGRATIONS

ARCADE comprises a set of tools capable of outputting large amounts of experimental results data on a virtualized OT platform. Such data is amenable to a wide range of analytical and machine learning (ML) methods for data classification. Due to the size of the problem space as discussed in section 3.2.7, machine learning methods are a useful tool for increasing analysis efficiency. In this section, we briefly detail potential implementations for Analytical and ML methods for run classification, as well as future proposed integrations of ARCADE with additional tools.

While Sandia's Dakota platform allows for performant analytical analysis of large datasets, given the size of the problem search space, the implementation of ML methods is a natural fit for producing faster results that may be verified against the classical methods of Dakota. The CSV formatted data files produced by ARCADE produce a naturally labelled dataset for each run. Because labels of SeBD must be applied to complete datasets, some transformative procedure will be necessary for the Data Collector outputs to compose each as a single vector for a classification procedure. Example transformative procedures that will likely be applied include statistical feature extraction, time series aggregation/Fourier transformation procedures, and the usage of an autoencoder. A simple figure detailing this transformative process is displayed in Figure 7. Alternatively, an unsupervised framework such as clustering may be applied to input datasets, and resulting output features may be considered for classification.




Figure 7: Example ML Analytical Workflow

Future integration efforts within ARCADE will expand beyond ML applications. The real time results of ARCADE cyber attack simulation allow for streamlined integration with a range of Sandia created analytical tools such as the ETE and Dante. The Sandia Equipment Test Environment (ETE) is an experimental environment developed for experimentation with OT hardware in the loop which will provide critical support for Tier 3 TCA activities. Dante is a Sandia developed high fidelity physical security analysis suite, which can simulate physical security threats to facilities and determine optimal defender response [8]. ARCADE is a natural candidate that may expand the capabilities of both existing tools, and additionally developed systems in the future as needs arise.

5. RESULTS

Applying ARCADE to the X-Energy Xe-100 reactor Flownex® simulation, a set of individual UCAs and combination of UCAs were evaluated in a limited scope. The most mature ARCADE UCA categories (U1 and U2) were evaluated on the Turbine Control Valve (TCV) and the helium circulator controllers. Table 1 provides a matrix of the UCAs evaluated and a description of their implementation. The key evaluation metrics used were the maximum and average fuel temperatures, this will be expanded in the future as the evaluation systems in ARCADE continue development. The system was allowed to converge for 2000 simulation seconds before each UCA was activated for 225 simulation seconds. The UCA activation time was determined experimentally to allow all the solutions to reach maximum temperature and trend downward.

System	Plot Label	UCA Category	Description
TCV	UCA #1	U2	Valve set to %100 open
TCV	UCA #2	U1	Valve set to %0 open
CIRC	UCA #1	U1	Circulator A stopped
CIRC	UCA #2	U2	Circulator A set to %30 above nominal speed
CIRC	UCA #3	U2	Circulator A & B set to %30 above nominal speed
TCV & CIRC	UCA #1	U1 & U2	Circulator A set to %30 above nominal speed & TCV set to %100 open
TCV & CIRC	UCA #2	U2	Circulator A set to %30 above nominal speed & TCV set to %0 open
TCV & CIRC	UCA #2	U1 & U2	Circulator A & B set to %30 above nominal speed & TCV set to %0 open
TCV & CIRC	UCA #4	U1 & U2	Circulator A & B set to %30 above nominal speed & TCV set to %100 open

Table 1: UCAs evaluated with ARCADE.

The first set of UCAs evaluated were against the Turbine Control Valve which regulates the steam entering the generator turbine. This controller was selected for two reasons: the system is not unique to HTGRs, and to investigate secondary side dynamics effect on the fuel temperature. The effects of all the UCAs were negligible on the maximum fuel temperature and were far from the known maximum safe TRISO operating temperature of ~1695°C [9]. Though some effects were observed on the average fuel temperature, this was expected and contributes to the reactor's natural negative reactivity temperature coefficient which kept maximum temperatures low.

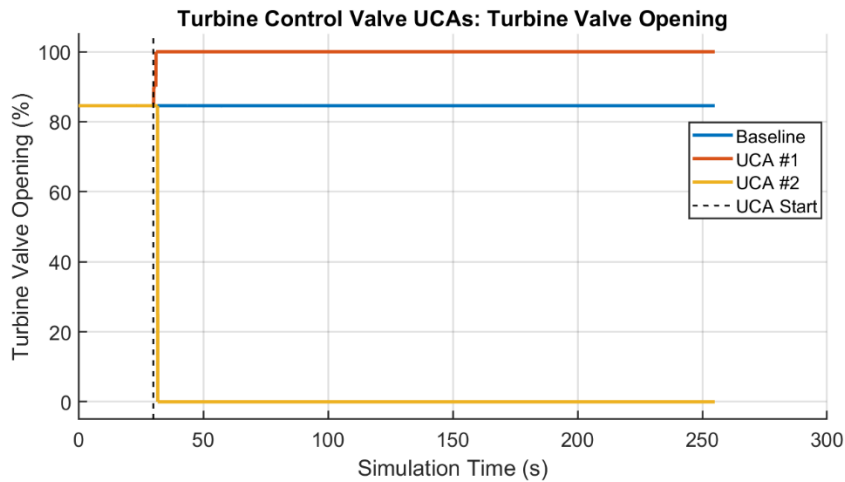


Figure 8: Turbine Control Valve opening during each TCV UCA transient.




Figure 9: Average fuel temperature during each TCV UCA transient.




Figure 10: Maximum fuel temperature during each TCV UCA transient.

The circulator controllers were selected as the second candidates for UCA evaluation as their manipulation was predicted to have a significant impact on primary loop temperatures. Circulator A was the primary target, the third and final UCA was applied to both circulator A and B. More investigation into this system is required as there are many more UCA combinations on these two systems to investigate. The preliminary results below are promising, showing that even in the event of both circulators being attacked the fuel remains safe.

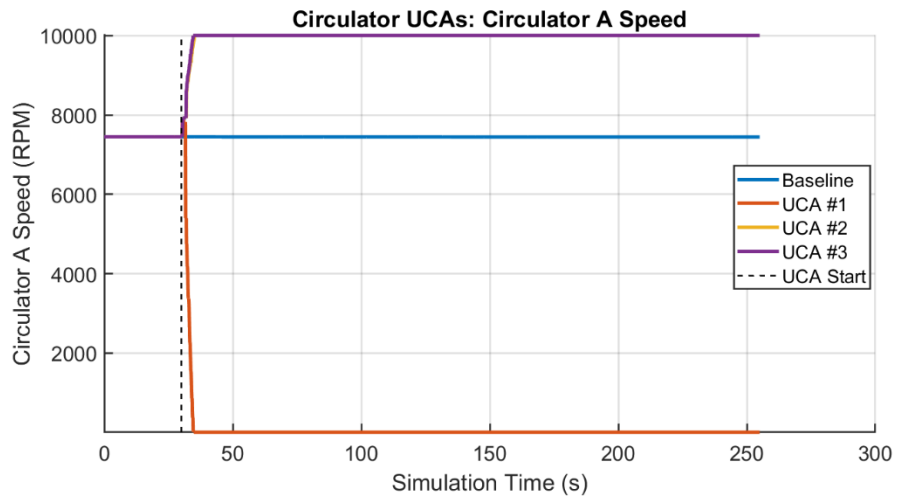


Figure 11: Circulator A speed during UCA transients.




Figure 12: Circulator B speed during UCA transients.




Figure 13: Average fuel temperature during each circulator UCA transient.




Figure 14: Maximum fuel temperature during each circulator UCA transient.

Finally, combinations of UCAs on both the TCV and circulators were evaluated. The first UCAs focused on circulator A, during which it was observed that circulator B decreased speed in response to the TCV being closed. This observation led to the creating what was expected to be the most impactful final UCA, which set both circulators to %30 above nominal speed and closed the TCV. This did cause the greatest effect on the maximum fuel temperature for the combine set of UCAs, but this was not significant to the safety of the fuel as shown in the below figures. There are some combinations of events which could not be simulated due to zero crossing events, ARCADE is helping identify and diagnose simulator robustness in extreme scenarios which will improve them for use in digital twin applications.




Figure 15: Average fuel temperature during each combine UCA transient.

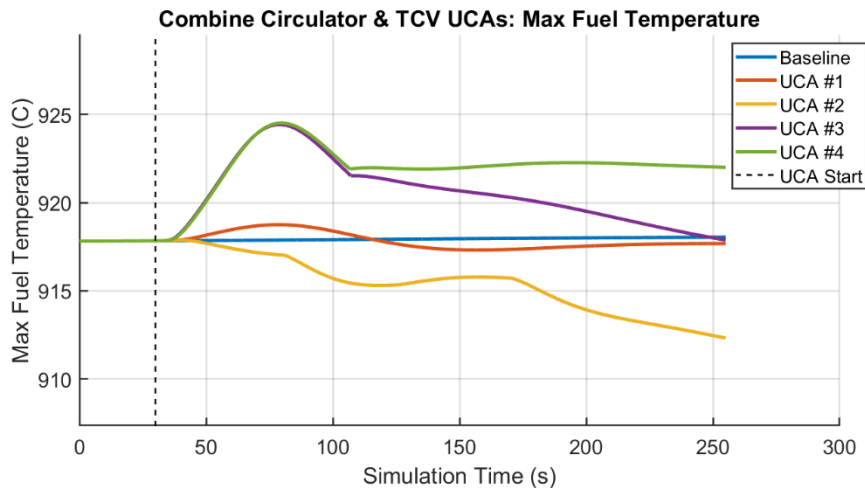


Figure 16: Maximum fuel temperature during each combine UCA transient.

It should be noted that this evaluation is far from complete, but these results demonstrate that it is possible to use ARCADE with high fidelity simulations to identify robustness factors in reactor designs. These preliminary results have also shown that the Xe-100 reactor design has significant inherent robustness features against cyber threats. Its critical to add that these UCA events were tested with the Reactor Protection System (RPS) and Investment Protection System (IPS) deactivated. Even under the exceptional stress of 3 concurrent UCAs with deactivated safety systems, the reactor maintained fuel safety with significant margins. Demonstrating this resiliency with ARCADE in the TCA framework is the first step to ensuring that the safety and security of Advanced Reactors is both thorough and risk based.

ACKNOWLEDGEMENTS

The authors would like to thank Matt Hertel for his support and assistance with experiment design, and Mike Rowland for his assistance on the operational theory and system design.

REFERENCES

1. Lee T. Maccarone, A.S.H., Michael T. Rowland. *Design of Defensive Cyber Security Architectures Using Event Trees*. in *American Nuclear Society Annual Meeting*. 2024. Las Vegas.
2. IAEA, *Computer Security Techniques for Nuclear Facilities*. 2021, Vienna: INTERNATIONAL ATOMIC ENERGY AGENCY.
3. Maccarone, L., A.S. Hahn, and M.T. Rowland, *System-Level Design Analysis for Advanced Reactor Cybersecurity*. 2023.
4. Karch, B. and M. Rowland, *Security Evaluation of Smart Cards and Secure Tokens: Benefits and Drawbacks for Reducing Supply Chain Risks of Nuclear Power Plants*. 2022: United States. p. Medium: ED; Size: 48 p.
5. Crussell, J., et al. *minimega v3.0*. 2015.
6. Hahn, A.S., *ManiPIO - Manipulate Process I/O for Industrial Control Systems*. 2021: United States. p. Medium: ED; Size: 7 p.
7. Adams, B.M., et al., *Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.15 User's Manual*. 2021: United States. p. Medium: ED; Size: 360 p.
8. Hart, B., et al. *Dante agent architecture for force-on-force wargame simulation and training*. in *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*. 2017.
9. Stempien, J.D., et al., *High-temperature safety testing of irradiated AGR-1 Triso fuel*. 2016, INL: Idaho Falls.