

From Cold Start to Critical: Formal Synthesis of Hybrid Controllers

PI: Dane A. Sabo
dane.sabo@pitt.edu

Advisor: Dr. Daniel G. Cole
dgcole@pitt.edu

Track: PhD Mechanical Engineering

Monday 10th November, 2025

1 Goals and Outcomes

The goal of this research is to develop a methodology for creating autonomous hybrid control systems with mathematical guarantees of safe and correct behavior.

Nuclear power plants require the highest levels of control system reliability, where failures can result in significant economic losses, service interruptions, or radiological release. Currently, nuclear plant operations rely on extensively trained human operators who follow detailed written procedures and strict regulatory requirements to manage reactor control. These operators make critical decisions about when to switch between different control modes based on their interpretation of plant conditions and procedural guidance. However, this reliance on human operators prevents the introduction of autonomous control capabilities and creates a fundamental economic challenge for next-generation reactor designs. Emerging technologies like small modular reactors face significantly higher per-megawatt staffing costs than conventional plants, threatening their economic viability. What is needed is a way to create autonomous control systems that can safely manage complex operational sequences with the same level of assurance as human-operated systems, but without requiring constant human supervision.

To address this need, we will combine formal methods from computer science with control theory to build hybrid control systems that are correct by construction. Hybrid systems use discrete logic to switch between continuous control modes, similar to how operators change control strategies. Existing formal methods can generate provably correct switching logic from written requirements, but they cannot handle the continuous dynamics that occur during transitions between modes. Meanwhile, traditional control theory can verify continuous behavior but lacks tools for proving correctness of discrete switching decisions. By synthesizing discrete mode transitions directly from written operating procedures and verifying continuous behavior between transitions, we can create hybrid control systems with end-to-end correctness guarantees. If we can formalize existing procedures into logical specifications and verify that continuous dynamics satisfy transition requirements, then we can build autonomous controllers that are provably free from design defects. This approach will enable autonomous control in nuclear power plants while maintaining the high safety standards required by the industry. This work is conducted within the University of Pittsburgh Cyber Energy Center, which provides access to industry collaboration and Emerson control hardware, ensuring that solutions developed are aligned with practical implementation requirements.

If this research is successful, we will be able to do the following:

1. **Translate written procedures into verified control logic.** We will develop a methodology for converting existing written operating procedures into formal specifications that can be automatically synthesized into discrete control logic. This process will use structured intermediate representations to bridge natural language procedures and mathematical logic. Control system engineers will be able to generate verified mode-switching controllers directly from regulatory procedures without requiring expertise in formal methods, reducing the barrier to creating high-assurance control systems.
2. **Verify continuous control behavior across mode transitions.** We will establish methods for analyzing continuous control modes to ensure they satisfy the discrete transition requirements. Using a combination of classical control theory for linear systems and reachability analysis for nonlinear dynamics, we will verify that each continuous mode can safely reach its intended transitions. Engineers will be able to design continuous controllers using

standard practices while iterating to ensure broader system correctness, proving that mode transitions occur safely and at the right times.

3. **Demonstrate autonomous reactor startup control with safety guarantees.** We will apply this methodology to develop an autonomous controller for nuclear reactor startup procedures, implementing it on a small modular reactor simulation using industry-standard control hardware. This demonstration will prove correctness across multiple coordinated control modes from cold shutdown through criticality to power operation. We will provide evidence that autonomous hybrid control can be realized in the nuclear industry with current control equipment, establishing a path toward reducing operator staffing requirements while maintaining safety.

The innovation in this work is the unification of discrete synthesis and continuous verification to enable end-to-end correctness guarantees for hybrid systems. If successful, control engineers will be able to create autonomous controllers from existing procedures with mathematical proof of correct behavior. High-assurance autonomous control will become practical for safety-critical applications. This capability is essential for the economic viability of next-generation nuclear power. Small modular reactors represent a promising solution to growing energy demands, but their success depends on reducing per-megawatt operating costs through increased autonomy. This research will provide the tools to achieve that autonomy while maintaining the exceptional safety record required by the nuclear industry.

2 State of the Art and Limits of Current Practice

The principal aim of this research is to create autonomous reactor control systems that are tractably safe. But, to understand what exactly is being automated, it is important to understand how nuclear reactors are operated today. First, the reactor operator themselves is discussed. Then, operating procedures that we aim to leverage later are examined. Next, limitations of human-based operation are investigated, while finally we discuss current formal methods based approaches to building reactor control systems.

2.1 Current Reactor Procedures and Operation

Current generation nuclear power plants employ 3,600+ active NRC-licensed reactor operators in the United States. These operators are divided into Reactor Operators (ROs) who manipulate reactor controls and Senior Reactor Operators (SROs) who direct plant operations and serve as shift supervisors [1]. Staffing typically requires 2+ ROs with at least one SRO for current generation units. To become a reactor operator, an individual might spend up to six years to pass required training [?].

The role of human operators is paradoxically both critical and problematic. Operators hold legal authority under 10 CFR Part 55 to make critical decisions including departing from normal regulations during emergencies. The Three Mile Island (TMI) accident demonstrated how “combination of personnel error, design deficiencies, and component failures” led to partial meltdown when operators “misread confusing and contradictory readings and shut off the emergency water system” [2]. The President’s Commission on TMI identified a fundamental ambiguity: placing “responsibility and accountability for safe power plant operations...on the licensee in all circumstances” without formal verification that operators can fulfill this responsibility under all conditions [2]. This tension between operational flexibility and safety assurance remains unresolved in current practice.

Nuclear plant procedures exist in a hierarchy: normal operating procedures for routine operations, abnormal operating procedures for off-normal conditions, Emergency Operating Procedures (EOPs) for design-basis accidents, Severe Accident Management Guidelines (SAMGs) for beyond-design-basis events, and Extensive Damage Mitigation Guidelines (EDMGs) for catastrophic damage scenarios. These procedures must comply with 10 CFR 50.34(b)(6)(ii) and are developed using guidance from NUREG-0899 [3], but their development process relies fundamentally on expert judgment and simulator validation rather than formal verification. Procedures undergo technical evaluation, simulator validation testing, and biennial review as part of operator requalification under 10 CFR 55.59 [1]. Despite these rigorous development processes, procedures fundamentally lack formal verification of key safety properties. There is no mathematical proof that procedures cover all possible plant states, that required actions can be completed within available timeframes under all scenarios, or that transitions between procedure sets maintain safety invariants.

LIMITATION: *Procedures lack formal verification of correctness and completeness.* Current procedure development relies on expert judgment and simulator validation. No mathematical proof exists that procedures cover all possible plant states, that required actions can be completed within available timeframes, or that transitions between procedure sets maintain safety invariants. Paper-based procedures cannot ensure correct application, and even computer-based procedure systems lack the formal guarantees that automated reasoning could provide.

Nuclear plants operate with multiple control modes: automatic control where the reactor control system maintains target parameters through continuous rod adjustment, manual control where operators directly manipulate control rods, and various intermediate modes. In typical pressurized water reactor operation, the reactor control system automatically maintains a floating average temperature, compensating for changes in power demand with reactivity feedback loops alone. Safety systems instead operate with implemented automation. Reactor Protection Systems trip automatically on safety signals with millisecond response times, and engineered safety features actuate automatically on accident signals without operator action required.

The current division between automated and human-controlled functions reveals the fundamental challenge of hybrid control. Highly automated systems handle reactor protection like automatic trips on safety parameters, emergency core cooling actuation, containment isolation, and basic process control. Human operators, however, retain control of strategic decision-making such as power level changes, startup/shutdown sequences, mode transitions, and procedure implementation.

LIMITATION: *Current practice treats continuous plant dynamics and discrete control logic separately.* No application of hybrid control theory exists that could provide mathematical guarantees across mode transitions, verify timing properties formally, or optimize the automation-human interaction trade-off with provable safety bounds.

2.2 Human Factors in Nuclear Accidents

The persistent role of human error in nuclear safety incidents, despite decades of improvements in training and procedures, provides perhaps the most compelling motivation for formal automated control with mathematical safety guarantees.

Multiple independent analyses converge on a striking statistic: **70–80% of all nuclear power plant events are attributed to human error** versus approximately 20% to equipment failures [4, 5]. More significantly, the International Atomic Energy Agency concluded that “human error was the root cause of all severe accidents at nuclear power plants”—a categorical statement spanning

Three Mile Island, Chernobyl, and Fukushima Daiichi [6]. A detailed analysis of 190 events at Chinese nuclear power plants from 2007–2020 [7] found that 53% of events involved active errors while 92% were associated with latent errors (organizational and systemic weaknesses that create conditions for failure). The persistence of this 70–80% human error contribution despite four decades of continuous improvements in operator training, control room design, procedures, and human factors engineering. This suggests fundamental cognitive limitations rather than remediable deficiencies.

The Three Mile Island Unit 2 accident on March 28, 1979 remains the definitive case study in human factors failures in nuclear operations. The accident began at 4:00 AM with a routine feed-water pump trip, escalating when a pressure-operated relief valve (PORV) stuck open—draining reactor coolant—but control room instrumentation showed only whether the valve had been commanded to close, not whether it actually closed. When Emergency Core Cooling System pumps automatically activated as designed, operators made the fateful decision to shut them down based on their incorrect assessment of plant conditions. The result was a massive loss of coolant accident and the core quickly began to overheat. During the emergency, operators faced more than 100 simultaneous alarms, overwhelming their cognitive capacity [2]. The core suffered partial meltdown with **44% of the fuel melting** before the situation was stabilized.

Quantitative risk analysis revealed the magnitude of failure in existing safety assessment methods: the actual core damage probability was approximately 5% per year while Probabilistic Risk Assessment had predicted 0.01% per year—a **500-fold underestimation**. This dramatic failure demonstrated that human reliability could not be adequately assessed through expert judgment and historical data alone. quantify human error probabilities and performance shaping factors. The SPAR-H method represents current best practice, providing nominal Human Error Probabilities (HEPs) of **0.01 (1%) for diagnosis tasks** and **0.001 (0.1%) for action tasks** under optimal conditions [8].

However, these nominal error rates degrade dramatically under realistic accident conditions: inadequate available time increases HEP by **10-fold**, extreme stress by **5-fold**, high complexity by **5-fold**, missing procedures by **50-fold**, and poor ergonomics by **50-fold**. Under combined adverse conditions typical of severe accidents, human error probabilities can approach **0.1 to 1.0 (10% to 100%)**—essentially guaranteed failure for complex diagnosis tasks [9].

Rasmussen’s influential 1983 taxonomy [10] divides human errors into skill-based (highly practiced responses, HEP 10^{-3} to 10^{-4}), rule-based (following procedures, HEP 10^{-2} to 10^{-1}), and knowledge-based (novel problem solving, HEP 10^{-1} to 1). Severe accidents inherently require knowledge-based responses where human reliability is lowest. Miller’s classic 1956 finding [11] that working memory capacity is limited to 7 ± 2 chunks explains why Three Mile Island’s 100+ simultaneous alarms exceeded operators’ processing capacity.

LIMITATION: *Human factors impose fundamental reliability limits that cannot be overcome through training alone.* Response time limitations constrain human effectiveness—reactor protection systems must respond in milliseconds, 100–1000 times faster than human operators. Cognitive biases systematically distort judgment: confirmation bias, overconfidence, and anchoring bias are inherent features of human cognition, not individual failings [12]. The persistent 70–80% human error contribution despite four decades of improvements demonstrates that these limitations are fundamental rather than remediable.

2.3 HARDENS and Formal Methods

The High Assurance Rigorous Digital Engineering for Nuclear Safety (HARDENS) project, completed by Galois, Inc. for the U.S. Nuclear Regulatory Commission in 2022, represents the most advanced application of formal methods to nuclear reactor control systems to date—and simultaneously reveals the critical gaps that remain.

2.3.1 Rigorous Digital Engineering Demonstrated Feasibility

HARDENS aimed to address the nuclear industry’s fundamental dilemma: existing U.S. nuclear control rooms rely on analog technologies from the 1950s–60s, making construction costs exceed \$500 million and timelines stretch to decades. The NRC contracted Galois to demonstrate that Model-Based Systems Engineering and formal methods could design, verify, and implement a complex protection system meeting regulatory criteria at a fraction of typical cost.

The project delivered far beyond its scope, creating what Galois describes as “the world’s most advanced, high-assurance protection system demonstrator.” Completed in **nine months at a tiny fraction of typical control system costs** [13], the project produced a complete Reactor Trip System (RTS) implementation with full traceability from NRC Request for Proposals and IEEE standards through formal architecture specifications to formally verified binaries and hardware running on FPGA demonstrator boards.

Principal Investigator Joseph Kiniry led the team in applying Galois’s Rigorous Digital Engineering methodology combining model-based engineering, digital twins with measurable fidelity, and applied formal methods. The approach integrates multiple abstraction levels—from semi-formal natural language requirements through formal specifications to verified implementations—all maintained as integrated artifacts rather than separate documentation prone to divergence.

2.3.2 Comprehensive Formal Methods Toolkit Provided Verification

HARDENS employed an impressive array of formal methods tools and techniques across the verification hierarchy. High-level specifications used Lando, SysMLv2, and FRET (NASA JPL’s Formal Requirements Elicitation Tool) to capture stakeholder requirements, domain engineering, certification requirements, and safety requirements. Requirements were formally analyzed for **consistency, completeness, and realizability** using SAT and SMT solvers—verification that current procedure development methods lack.

Executable formal models employed Cryptol to create an executable behavioral model of the entire RTS including all subsystems, components, and formal digital twin models of sensors, actuators, and compute infrastructure. Automatic code synthesis generated formally verifiable C implementations and System Verilog hardware implementations directly from Cryptol models—eliminating the traditional gap between specification and implementation where errors commonly arise.

Formal verification tools included SAW (Software Analysis Workbench) for proving equivalence between models and implementations, Frama-C for C code verification, and Yosys for hardware verification. HARDENS verified both automatically synthesized and hand-written implementations against their models and against each other, providing redundant assurance paths.

This multi-layered verification approach represents a quantum leap beyond current nuclear I&C verification practices, which rely primarily on testing and simulation. HARDENS demonstrated that **complete formal verification from requirements to implementation is technically feasible** for safety-critical nuclear control systems.

2.3.3 Critical Limitation: Discrete Control Logic Only

Despite its impressive accomplishments, HARDENS has a fundamental limitation directly relevant to hybrid control synthesis: **the project addressed only discrete digital control logic without modeling or verifying continuous reactor dynamics**. The Reactor Trip System specification and formal verification covered discrete state transitions (trip/no-trip decisions), digital sensor input processing through discrete logic, and discrete actuation outputs (reactor trip commands). The system correctly implements the digital control logic for reactor protection with mathematical guarantees.

However, the project did not address continuous dynamics of nuclear reactor physics including neutron kinetics, thermal-hydraulics, xenon oscillations, fuel temperature feedback, coolant flow dynamics, and heat transfer—all governed by continuous differential equations. Real reactor safety depends on the interaction between continuous processes (temperature, pressure, neutron flux evolving according to differential equations) and discrete control decisions (trip/no-trip, valve open/close, pump on/off). HARDENS verified the discrete controller in isolation but not the closed-loop hybrid system behavior.

LIMITATION: *HARDENS addressed discrete control logic without continuous dynamics or hybrid system verification.* Hybrid automata, differential dynamic logic, or similar hybrid systems formalisms would be required to specify and verify properties like “the controller maintains core temperature below safety limits under all possible disturbances”—a property that inherently spans continuous and discrete dynamics. Verifying discrete control logic alone provides no guarantee that the closed-loop system exhibits desired continuous behavior such as stability, convergence to setpoints, or maintained safety margins.

2.3.4 Experimental Validation Gap Limits Technology Readiness

The second critical limitation is **absence of experimental validation** in actual nuclear facilities or realistic operational environments. HARDENS produced a demonstrator system at Technology Readiness Level 3–4 (analytical proof of concept with laboratory breadboard validation) rather than a deployment-ready system validated through extended operational testing. The NRC Final Report explicitly notes [13]: “All material is considered in development and not a finalized product” and “The demonstration of its technical soundness was to be at a level consistent with satisfaction of the current regulatory criteria, although with no explicit demonstration of how regulatory requirements are met.”

The project did not include deployment in actual nuclear facilities, testing with real reactor systems under operational conditions, side-by-side validation with operational analog RTS systems, systematic failure mode testing (radiation effects, electromagnetic interference, temperature extremes), actual NRC licensing review, or human factors validation with licensed nuclear operators in realistic control room scenarios.

LIMITATION: *HARDENS achieved TRL 3–4 without experimental validation.* While formal verification provides mathematical correctness guarantees for the implemented discrete logic, the gap between formal verification and actual system deployment involves myriad practical considerations: integration with legacy systems, long-term reliability under harsh environments, human-system interaction in realistic operational contexts, and regulatory acceptance of formal methods as primary assurance evidence.

2.4 Research Imperative: Formal Hybrid Control Synthesis

Three converging lines of evidence establish an urgent research imperative for formal hybrid control synthesis applied to nuclear reactor systems.

Current reactor control practices reveal fundamental gaps in verification. Procedures lack mathematical proofs of completeness or timing adequacy. Mode transitions preserve safety properties only informally. Operator decision-making relies on training rather than verified algorithms. The divide between continuous plant dynamics and discrete control logic has never been bridged with formal methods. Despite extensive regulatory frameworks developed over six decades, **no mathematical guarantees exist** that current control approaches maintain safety under all possible scenarios.

Human factors in nuclear accidents demonstrate that human error contributes to 70–80% of nuclear incidents despite four decades of systematic improvements. The IAEA’s categorical statement that “human error was the root cause of all severe accidents” reveals fundamental cognitive limitations: working memory capacity of 7 ± 2 chunks, response times of seconds to minutes versus milliseconds required, cognitive biases immune to training, stress-induced performance degradation. Human Reliability Analysis methods document error probabilities of 0.001–0.01 under optimal conditions degrading to 0.1–1.0 under realistic accident conditions. These limitations **cannot be overcome through human factors improvements alone**.

The HARDENS project proved that formal verification is technically feasible and economically viable for nuclear control systems, achieving complete verification from requirements to implementation in nine months at a fraction of typical costs. However, HARDENS addressed only discrete control logic without considering continuous reactor dynamics or hybrid system verification, and the demonstrator achieved only TRL 3–4 without experimental validation in realistic nuclear environments. These limitations directly define the research frontier: **formal synthesis of hybrid controllers that provide mathematical safety guarantees across both continuous plant dynamics and discrete control logic**.

The research opportunity is clear. Nuclear reactors are quintessential hybrid cyber-physical systems where continuous neutron kinetics, thermal-hydraulics, and heat transfer interact with discrete control mode decisions, trip logic, and valve states. Current practice treats these domains separately—reactor physics analyzed with simulation, control logic verified through testing, human operators expected to integrate everything through procedures. **Hybrid control synthesis offers the possibility of unified formal treatment** where controllers are automatically generated from high-level safety specifications with mathematical proofs that guarantee safe operation across all modes, all plant states, and all credible disturbances.

Recent advances in hybrid systems theory—including reachability analysis, barrier certificates, counterexample-guided inductive synthesis, and satisfiability modulo theories for hybrid systems—provide the theoretical foundation. Computational advances enable verification of systems with continuous state spaces that were intractable a decade ago. The confluence of mature formal methods, powerful verification tools demonstrated by HARDENS, urgent safety imperatives documented by persistent human error statistics, and fundamental gaps in current hybrid dynamics treatment creates a compelling and timely research opportunity.

3 Research Approach

This research will overcome the limitations of current practice to build high-assurance hybrid control systems for critical infrastructure. Hybrid systems combine continuous dynamics (flows) with discrete transitions (jumps), which can be formally expressed as:

$$\dot{x}(t) = f(x(t), q(t), u(t)) \quad (1)$$

$$q(k+1) = v(x(k), q(k), u(k)) \quad (2)$$

Here, $f(\cdot)$ defines the continuous dynamics while $v(\cdot)$ governs discrete transitions. The continuous states x , discrete state q , and control input u interact to produce hybrid behavior. The discrete state q defines which continuous dynamics mode is currently active. Our focus centers on continuous autonomous hybrid systems, where continuous states remain unchanged during jumps—a property naturally exhibited by physical systems. For example, a nuclear reactor switching from warm-up to load-following control cannot instantaneously change its temperature or control rod position, but can instantaneously change control laws.

To build these systems with formal correctness guarantees, we must accomplish three main thrusts:

1. Translate operating procedures and requirements into temporal logic formulae
2. Create the discrete half of a hybrid controller using reactive synthesis
3. Develop continuous controllers to operate between modes, and verify their correctness using reachability analysis

The following sections discuss how these thrusts will be accomplished.

3.1 $(Procedures \wedge FRET) \rightarrow TemporalSpecifications$

The motivation behind this work stems from the fact that commercial nuclear power operations remain manually controlled by human operators, despite significant advances in control systems sophistication. The key insight is that procedures performed by human operators are highly prescriptive and well-documented. This suggests that human operators in nuclear power plants may not be entirely necessary given today’s available technology.

Written procedures and requirements in nuclear power are sufficiently detailed that we may be able to translate them into logical formulae with minimal effort. If successful, this approach would enable automation of existing procedures without requiring system reengineering. To formalize these procedures, we will use temporal logic, which captures system behaviors through temporal relations. Linear Temporal Logic (LTL) provides four fundamental operators: next (X), eventually (F), globally (G), and until (U). These operators enable precise specification of time-dependent requirements.

Consider a nuclear reactor SCRAM requirement expressed in natural language: “*If a high temperature alarm triggers, control rods must immediately insert and remain inserted until operator reset.*” This plain language requirement can be translated into a rigorous logical specification:

$$G(HighTemp \rightarrow X(RodsInserted \wedge (\neg RodsWithdrawn U OperatorReset))) \quad (3)$$

This specification precisely captures the temporal relationship between the alarm condition, the required response, and the persistence requirement. The global operator G ensures this property holds throughout system operation, while the next operator X enforces immediate response. The until operator U maintains the state constraint until the reset condition occurs.

The most efficient path to accomplish this translation is through NASA's Formal Requirements Elicitation Tool (FRET). FRET employs a specialized requirements language called FRETish that restricts requirements to easily understood components while eliminating ambiguity. FRETish bridges natural language and mathematical specifications through a structured English-like syntax that is automatically translatable to temporal logic.

FRET enforces this structure by requiring all requirements to contain six components:

1. Scope: *What modes does this requirement apply to?*
2. Condition: *Scope plus additional specificity*
3. Component: *What system element does this requirement affect?*
4. Shall
5. Timing: *When does the response occur?*
6. Response: *What action should be taken?*

FRET provides functionality to check the *realizability* of a system. Realizability analysis determines whether written requirements are complete by examining the six structural components. Complete requirements are those that neither conflict with one another nor leave any behavior undefined. Systems that are not realizable from their procedure definitions and design requirements present problems beyond autonomous control implementation. Such systems contain behavioral inconsistencies that represent the physical equivalent of software bugs. Using FRET during autonomous controller development allows us to identify and resolve these errors systematically.

The second category of realizability issues involves undefined behaviors that are typically left to human judgment during control operations. This ambiguity is undesirable for high-assurance systems, since even well-trained humans remain prone to errors. By addressing these specification gaps in FRET during autonomous controller development, we can deliver controllers free from these vulnerabilities.

FRET provides the capability to export requirements in temporal logic format compatible with reactive synthesis tools. This export functionality enables progression to the next step of our approach: synthesizing discrete mode switching behavior from the formalized requirements.

3.2 $(TemporalLogic \wedge ReactiveSynthesis) \rightarrow DiscreteAutomata$

Reactive synthesis is an active research field in computer science focused on generating discrete controllers from temporal logic specifications. The term “reactive” indicates that the system responds to environmental inputs to produce control outputs. These synthesized systems are finite in size, where each node represents a unique discrete state. The connections between nodes, called *state transitions*, specify the conditions under which the discrete controller moves from state to state. This complete mapping of possible states and transitions constitutes a *discrete automaton*. Discrete automata can be represented graphically as a series of nodes that are discrete states, with traces indicating transitions between states. From the automaton graph, it becomes possible to fully describe the dynamics of the discrete system and develop intuitive understanding of system behavior. Hybrid systems naturally exhibit discrete behavior amenable to formal analysis through these finite state representations.

We will employ state-of-the-art reactive synthesis tools, particularly Strix, which has demonstrated superior performance in the Reactive Synthesis Competition (SYNTCOMP) through efficient parity game solving algorithms. Strix translates linear temporal logic specifications into deterministic automata automatically while maximizing generated automata quality. Once constructed, the automaton can be straightforwardly implemented using standard programming control flow constructs. The graphical representation provided by the automaton enables inspection and facilitates communication with controls programmers who may not have formal methods expertise.

We will use discrete automata to represent the switching behavior of our hybrid system. This approach yields an important theoretical guarantee: because the discrete automaton is synthesized entirely through automated tools from design requirements and operating procedures, we can prove that the automaton—and therefore our hybrid switching behavior—is *correct by construction*. Correctness of the switching controller is paramount to this work. Mode switching represents the primary responsibility of human operators in control rooms today. Human operators possess the advantage of real-time judgment—when mistakes occur, they can correct them dynamically with capabilities that extend beyond written procedures. Autonomous control lacks this adaptive advantage. Instead, we must ensure that autonomous controllers replacing human operators will not make switching errors between continuous modes. By synthesizing controllers from logical specifications with guaranteed correctness, we eliminate the possibility of switching errors.

3.3 $(DiscreteAutomata \wedge ControlTheory \wedge Reachability) \rightarrow ContinuousModes$

While discrete system components will be synthesized with correctness guarantees, they represent only half of the complete system. Autonomous controllers like those we are developing exhibit continuous dynamics within discrete states, as described by $f(\cdot)$ in Equation 1. This section describes how we will develop continuous control modes, verify their correctness, and address the unique verification challenges of hybrid systems.

The approach described for producing discrete automata yields physics-agnostic specifications that represent only half of a complete hybrid autonomous controller. These automata alone cannot define the full behavior of the control systems we aim to construct. The continuous modes will be developed after discrete automaton construction, leveraging the automaton structure and transitions to design multiple smaller, specialized continuous controllers.

The discrete automaton transitions are key to the supervisory behavior of the autonomous controller. These transitions mark decision points for switching between continuous control modes and define their strategic objectives. We will classify three types of high-level continuous controller objectives based on discrete mode transitions:

1. **Stabilizing:** A stabilizing control mode has one primary objective: maintaining the hybrid system within its current discrete mode. This corresponds to steady-state normal operating modes, such as a full-power load-following controller in a nuclear power plant. Stabilizing modes can be identified from discrete automata as nodes with only incoming transitions.
2. **Transitory:** A transitory control mode has the primary goal of transitioning the hybrid system from one discrete state to another. In nuclear applications, this might represent a controlled warm-up procedure. Transitory modes ultimately drive the system toward a stabilizing steady-state mode. These modes may have secondary objectives within a discrete state, such as maintaining specific temperature ramp rates before reaching full-power operation.
3. **Expulsory:** An expulsory mode is a specialized transitory mode with additional safety con-

straints. Expulsory modes ensure the system is directed to a safe stabilizing mode during failure conditions. For example, if a transitory mode fails to achieve its intended transition, the expulsory mode activates to immediately and irreversibly guide the system toward a globally safe state. A reactor SCRAM exemplifies an expulsory continuous mode: when initiated, it must reliably terminate the nuclear reaction and direct the reactor toward stabilizing decay heat removal.

Building continuous modes after constructing discrete automata enables local controller design focused on satisfying discrete transitions. The primary challenge in hybrid system verification is ensuring global stability across transitions. Current techniques struggle with this problem because dynamic discontinuities complicate verification. This work alleviates these problems by designing continuous controllers specifically with transitions in mind. By decomposing continuous modes according to their required behavior at transition points, we avoid solving trajectories through the entire hybrid system. Instead, we can use local behavior information at transition boundaries. To ensure continuous modes satisfy their requirements, we will employ three main techniques: reachability analysis, assume-guarantee contracts, and barrier certificates.

Reachability Analysis: Reachability analysis computes the reachable set of states for a given input set. While trivial for linear continuous systems, recent advances have extended reachability to complex nonlinear systems. We will use reachability to define continuous state ranges at discrete transition boundaries and verify that requirements are satisfied within continuous modes. Recent advances using neural network approximations of Hamilton-Jacobi equations have demonstrated significant speedups while maintaining safety guarantees for high-dimensional systems, expanding the practical applicability of these methods.

Assume-Guarantee Contracts: Assume-guarantee contracts will be employed when continuous state boundaries are not explicitly defined. For any given mode, the input range for reachability analysis is defined by the output ranges of discrete modes that transition to it. This compositional approach ensures each continuous controller is prepared for its possible input range, enabling subsequent reachability analysis without requiring global system analysis.

Barrier Certificates: Finally, we will use barrier certificates to prove that mode transitions are satisfied. Barrier certificates ensure that continuous modes on either side of a transition behave appropriately. Control barrier functions provide a method to certify safety by establishing differential inequality conditions that guarantee forward invariance of safe sets. For example, a barrier certificate can guarantee that a transitory mode transferring control to a stabilizing mode will always move away from the transition boundary, rather than destabilizing the target stabilizing mode.

Combining these three techniques will enable us to prove that continuous components of our hybrid controller satisfy discrete requirements, and thus, complete system behavior. To demonstrate this methodology, we will develop an autonomous startup controller for a Small Modular Advanced High Temperature Reactor (SmAHTR). SmAHTR represents an ideal test case as a liquid-salt cooled reactor design with well-documented startup procedures that must transition through multiple distinct operational modes: initial cold conditions, controlled heating to operating temperature, approach to criticality, low-power physics testing, and power ascension to full operating capacity. We have already developed a high-fidelity SmAHTR model in Simulink that captures the thermal-hydraulic and neutron kinetics behavior essential for verifying continuous controller performance under realistic plant dynamics. The synthesized hybrid controller will be

implemented on an Emerson Ovation control system platform, which is representative of industry-standard control hardware deployed in modern nuclear facilities. The Advanced Reactor Cyber Analysis and Development Environment (ARCADE) suite will serve as the integration layer, managing real-time communication between the Simulink simulation and the Ovation controller. This hardware-in-the-loop configuration enables validation of the controller implementation on actual industrial control equipment interfacing with a realistic reactor simulation, providing assessment of computational performance, real-time execution constraints, and communication latency effects. By demonstrating autonomous startup control on this representative platform, we will establish both the theoretical validity and practical feasibility of the synthesis methodology for deployment in actual small modular reactor systems.

This unified approach addresses a fundamental gap in hybrid system design by bridging formal methods and control theory through a systematic, tool-supported methodology. By translating existing nuclear procedures into temporal logic, synthesizing provably correct discrete switching logic, and developing verified continuous controllers, we create a complete framework for autonomous hybrid control with mathematical guarantees. The result is an autonomous controller that not only replicates human operator decision-making but does so with formal assurance that switching logic is correct by construction and continuous behavior satisfies safety requirements. This methodology transforms nuclear reactor control from a manually intensive operation requiring constant human oversight into a fully autonomous system with higher reliability than human-operated alternatives. More broadly, this approach establishes a replicable framework for developing high-assurance autonomous controllers in any domain where operating procedures are well-documented and safety is paramount.

4 Metrics for Success

This research will be measured by advancement through Technology Readiness Levels, progressing from fundamental concepts to validated prototype demonstration. The work begins at TRL 2-3 and aims to reach TRL 5, where system components operate successfully in a relevant laboratory environment. This section explains why TRL advancement provides the most appropriate success metric and defines the specific criteria required to achieve TRL 5.

Technology Readiness Levels provide the ideal success metric because they explicitly measure the gap between academic proof-of-concept and practical deployment. This gap is precisely what this work aims to bridge. Academic metrics like papers published or theorems proved cannot capture practical feasibility. Empirical metrics like simulation accuracy or computational speed cannot demonstrate theoretical rigor. TRLs measure both dimensions simultaneously. Advancing from TRL 3 to TRL 5 requires maintaining theoretical rigor while progressively demonstrating practical feasibility. Formal verification must remain valid as the system moves from individual components to integrated hardware testing.

The nuclear industry requires extremely high assurance before deploying new control technologies. Demonstrating theoretical correctness alone is insufficient for adoption. Conversely, showing empirical performance without formal guarantees fails to meet regulatory requirements. TRLs capture this dual requirement naturally. Each level represents both increased practical maturity and sustained theoretical validity. Furthermore, TRL assessment forces explicit identification of remaining barriers to deployment. The nuclear industry already uses TRLs for technology assessment, making this metric directly relevant to potential adopters. Reaching TRL 5 provides a clear answer to industry questions about feasibility and maturity in a way that academic publica-

tions alone cannot.

The work currently exists at TRL 2-3. Formal synthesis and hybrid control verification principles have been established through prior research, placing the fundamental approach at TRL 2. The SmAHTR simulation model and initial procedure analysis place specific components at early TRL 3, where proof of concept has been partially demonstrated for individual elements but not integrated. The target state is TRL 5. Moving from current state to target requires achieving three intermediate levels, each representing a distinct validation milestone:

TRL 3 *Critical Function and Proof of Concept* For this research, TRL 3 means demonstrating that each component of the methodology works in isolation. SmAHTR startup procedures must be translated into temporal logic specifications that pass realizability analysis. A discrete automaton must be synthesized with interpretable structure. At least one continuous controller must be designed with reachability analysis proving that transition requirements are satisfied. Independent review must confirm that specifications match intended procedural behavior. This proves the fundamental approach on a simplified startup sequence.

TRL 4 *Laboratory Testing of Integrated Components* For this research, TRL 4 means demonstrating a complete integrated hybrid controller in simulation. All SmAHTR startup procedures must be formalized with a synthesized automaton covering all operational modes. Continuous controllers must exist for all discrete modes. Verification must be complete for all mode transitions using reachability analysis, barrier certificates, and assume-guarantee contracts. The integrated controller must execute complete startup sequences in software simulation with zero safety violations across multiple consecutive runs. This proves that formal correctness guarantees can be maintained throughout system integration.

TRL 5 *Laboratory Testing in Relevant Environment* For this research, TRL 5 means demonstrating the verified controller on industrial control hardware through hardware-in-the-loop testing. The discrete automaton must be implemented on the Emerson Ovation control system and verified to match synthesized specifications exactly. Continuous controllers must execute at required rates. The ARCADE interface must establish stable real-time communication between Ovation hardware and SmAHTR simulation. Complete autonomous startup sequences must execute via hardware-in-the-loop across the full operational envelope. The controller must handle off-nominal scenarios to validate that expulsive modes function correctly. For example, simulated sensor failures must trigger appropriate fault detection and mode transitions, and loss of cooling scenarios must activate SCRAM procedures as specified. Graded responses to minor disturbances are outside the scope of this work. Formal verification results must remain valid with discrete behavior matching specifications and continuous trajectories remaining within verified bounds. This proves that the methodology produces verified controllers implementable on industrial hardware.

These levels define progressively more demanding demonstrations. TRL 3 proves individual components work. TRL 4 proves they work together in simulation. TRL 5 proves they work on actual hardware in realistic conditions. Each level builds on the previous while adding new validation requirements.

Progress will be assessed quarterly through collection of specific data comparing actual results against TRL advancement criteria. Specification development status indicates progress toward TRL 3. Synthesis results and verification coverage indicate progress toward TRL 4. Simulation performance metrics and hardware integration milestones indicate progress toward TRL 5. The research plan will be revised only when new data invalidates fundamental assumptions. Unrealiz-

able specifications indicate procedure conflicts requiring refinement or alternative reactor selection. Unverifiable dynamics suggest model simplification or alternative verification methods are needed. Unachievable real-time performance requires controller simplification or hardware upgrades. Any revision will document the invalidating data, the failed assumption, and the modified pathway with adjusted scope.

This research succeeds if it achieves TRL 5 by demonstrating a complete autonomous hybrid controller with formal correctness guarantees operating on industrial control hardware through hardware-in-the-loop testing in a relevant laboratory environment. This establishes both theoretical validity and practical feasibility, proving that the methodology produces verified controllers and that implementation is achievable with current technology. It provides a clear pathway for nuclear industry adoption and broader application to safety-critical autonomous systems.

5 Risks and Contingencies

This research relies on several critical assumptions that, if invalidated, would require scope adjustment or methodological revision. The primary risks to successful completion fall into four categories: computational tractability of synthesis and verification, complexity of the discrete-continuous interface, completeness of procedure formalization, and hardware-in-the-loop integration challenges. Each risk has associated indicators for early detection and contingency plans that preserve research value even if core assumptions prove false. The staged project structure ensures that partial success yields publishable results and clear identification of remaining barriers to deployment.

5.1 Computational Tractability of Synthesis

The first major assumption is that formalized startup procedures will yield automata small enough for efficient synthesis and verification. Reactive synthesis scales exponentially with specification complexity, which creates risk that temporal logic specifications derived from complete startup procedures may produce automata with thousands of states. Such large automata would require synthesis times exceeding days or weeks, preventing demonstration of the complete methodology within project timelines. Reachability analysis for continuous modes with high-dimensional state spaces may similarly prove computationally intractable. Either barrier would constitute a fundamental obstacle to achieving the research objectives.

Several indicators would provide early warning of computational tractability problems. Synthesis times exceeding 24 hours for simplified procedure subsets would suggest that complete procedures are intractable. Generated automata containing more than 1,000 discrete states would indicate that the discrete state space is too large for efficient verification. Specifications flagged as unrealizable by FRET or STRIX would reveal fundamental conflicts in the formalized procedures. Reachability analysis failing to converge within reasonable time bounds would show that continuous mode verification cannot be completed with available computational resources.

The contingency plan for computational intractability is to reduce scope to a minimal viable startup sequence. This reduced sequence would cover only cold shutdown to criticality to low power hold, omitting power ascension and other operational phases. The subset would still demonstrate the complete methodology while reducing computational burden. The research contribution would remain valid even with reduced scope, proving that formal hybrid control synthesis is achievable for safety-critical nuclear applications. The limitation to simplified operational sequences would be explicitly documented as a constraint rather than a failure.

Reachability analysis specifically can exploit time-scale separation inherent in reactor dynamics. Fast thermal transients can be treated quasi-steady relative to slower nuclear kinetics, which enables decomposition into smaller subsystems. Temperature dynamics operate on time scales of seconds to minutes, while neutron kinetics respond in milliseconds to seconds for prompt effects and hours for xenon poisoning. These distinct time scales permit separate analysis with conservative coupling assumptions between subsystems, dramatically reducing the dimensionality of reachability computations.

Mitigation strategies exist even before contingency plans become necessary. Access to the University of Pittsburgh Center for Research Computing provides high-performance computing resources if single-workstation computation proves insufficient. Parallel synthesis algorithms and distributed reachability analysis can leverage these resources to extend computational feasibility. Compositional verification approaches using assume-guarantee reasoning can decompose monolithic verification problems into tractable subproblems, each of which can be solved independently before composition.

5.2 Discrete-Continuous Interface Formalization

The second critical assumption concerns the mapping between boolean guard conditions in temporal logic and continuous state boundaries required for mode transitions. This interface represents the fundamental challenge of hybrid systems: relating discrete switching logic to continuous dynamics. Temporal logic operates on boolean predicates, while continuous control requires reasoning about differential equations and reachable sets. Guard conditions that require complex nonlinear predicates may resist boolean abstraction, making synthesis intractable. Continuous safety regions that cannot be expressed as conjunctions of verifiable constraints would similarly create insurmountable verification challenges. The risk extends beyond static interface definition to dynamic behavior across transitions: barrier certificates may fail to exist for proposed transitions, or continuous modes may be unable to guarantee convergence to discrete transition boundaries.

Early indicators of interface formalization problems would appear during both synthesis and verification phases. Guard conditions requiring complex nonlinear predicates that resist boolean abstraction would suggest fundamental misalignment between discrete specifications and continuous realities. Continuous safety regions that cannot be expressed as conjunctions of half-spaces or polynomial inequalities would indicate that the interface between discrete guards and continuous invariants is too complex. Failure to construct barrier certificates proving safety across mode transitions would reveal that the continuous dynamics cannot be formally related to discrete switching logic. Reachability analysis showing that continuous modes cannot reach intended transition boundaries from all possible initial conditions would demonstrate that the synthesized discrete controller is incompatible with achievable continuous behavior.

The primary contingency for interface complexity is to restrict continuous modes to operate within polytopic invariants. Polytopes are state regions defined as intersections of linear half-spaces, which map directly to boolean predicates through linear inequality checks. This restriction ensures tractable synthesis while maintaining theoretical rigor, though at the cost of limiting expressiveness compared to arbitrary nonlinear regions. The discrete-continuous interface remains well-defined and verifiable with polytopic restrictions, providing a clear fallback position that preserves the core methodology. Conservative over-approximations offer an alternative approach: a nonlinear safe region can be inner-approximated by a polytope, sacrificing operational flexibility to maintain formal guarantees. The three-mode classification already structures the problem to

minimize complex transitions, with critical safety properties concentrated in expulsive modes that can receive additional design attention.

Mitigation strategies focus on designing continuous controllers with discrete transitions as primary objectives from the outset. Rather than designing continuous control laws independently and verifying transitions post-hoc, the approach uses transition requirements as design constraints. Control barrier functions provide a systematic method to synthesize controllers that guarantee forward invariance of safe sets and convergence to transition boundaries. This design-for-verification approach reduces the likelihood that interface complexity becomes insurmountable. Focusing verification effort on expulsive modes—where safety is most critical—allows more complex analysis to be applied selectively rather than uniformly across all modes, concentrating computational resources where they matter most for safety assurance.

5.3 Procedure Formalization Completeness

The third assumption is that existing SmAHTR startup procedures contain sufficient detail and clarity for translation into temporal logic specifications. Nuclear operating procedures, while extensively detailed, were written for human operators who bring contextual understanding and adaptive reasoning to their interpretation. Procedures may contain implicit knowledge, ambiguous directives, or references to operator judgment that resist formalization in current specification languages. Underspecified timing constraints, ambiguous condition definitions, or gaps in operational coverage would cause synthesis to fail or produce incorrect automata. The risk is not merely that formalization is difficult, but that current procedures fundamentally lack the precision required for autonomous control, revealing a gap between human-oriented documentation and machine-executable specifications.

Several indicators would reveal formalization completeness problems early in the project. FRET realizability checks failing due to underspecified behaviors or conflicting requirements would indicate that procedures do not form a complete specification. Multiple valid interpretations of procedural steps with no clear resolution would demonstrate that procedure language is insufficiently precise for automated synthesis. Procedures referencing “operator judgment,” “as appropriate,” or similar discretionary language for critical decisions would explicitly identify points where human reasoning cannot be directly formalized. Domain experts unable to provide crisp answers to specification questions about edge cases would suggest that the procedures themselves do not fully define system behavior, relying instead on operator training and experience to fill gaps.

The contingency plan treats inadequate specification as itself a research contribution rather than a project failure. Documenting specific ambiguities encountered would create a taxonomy of formalization barriers: timing underspecification, missing preconditions, discretionary actions, and undefined failure modes. Each category would be analyzed to understand why current procedure-writing practices produce these gaps and what specification languages would need to address them. Proposed extensions to FRETish or similar specification languages would demonstrate how to bridge the gap between current procedures and the precision needed for autonomous control. The research output would shift from “here is a complete autonomous controller” to “here is what formal autonomous control requires that current procedures do not provide, and here are language extensions to bridge that gap.” This contribution remains valuable to both the nuclear industry and formal methods community, establishing clear requirements for next-generation procedure development and autonomous control specification languages.

Early-stage procedure analysis with domain experts provides the primary mitigation strategy.

Collaboration through the University of Pittsburgh Cyber Energy Center enables identification and resolution of ambiguities before synthesis attempts, rather than discovering them during failed synthesis runs. Iterative refinement with reactor operators and control engineers can clarify procedural intent before formalization begins, reducing the risk of discovering insurmountable specification gaps late in the project. Comparison with procedures from multiple reactor designs—pressurized water reactors, boiling water reactors, and advanced designs—may reveal common patterns and standard ambiguities amenable to systematic resolution. This cross-design analysis would strengthen the generalizability of any proposed specification language extensions, ensuring they address industry-wide practices rather than SmAHTR-specific quirks.

5.4 Hardware-in-the-Loop Integration Complexity

The fourth assumption is that the ARCADE interface can provide stable real-time communication between Simulink simulation and Ovation control hardware at control rates required for reactor dynamics. Hardware-in-the-loop testing introduces timing constraints, communication latency, and platform compatibility challenges that are absent in pure simulation. Control rates for reactor systems typically range from 10-100 Hz for continuous control to millisecond response times for protection system actions. Control loop jitter, communication dropouts, or computational limitations in the Ovation PLC could prevent successful HIL validation even if the synthesized controller is theoretically correct. Real-time operating system constraints, network latency, and hardware execution speed may prove incompatible with verified timing assumptions embedded in the controller design.

Early indicators would identify hardware integration problems before they derail the entire validation effort. Communication dropouts or buffer overruns between ARCADE and Ovation would indicate that the interface cannot maintain stable real-time data exchange. The Ovation PLC proving unable to execute the synthesized automaton at required speed would reveal fundamental computational limitations of the target hardware platform. Timing analysis showing that hardware cannot meet real-time deadlines assumed during verification would demonstrate incompatibility between formal guarantees and physical implementation constraints.

The contingency plan is to demonstrate the controller in software-in-the-loop configuration with detailed timing analysis showing industrial hardware feasibility. Software-in-the-loop testing executes the complete verified controller in a real-time software environment that emulates hardware timing constraints without requiring physical hardware. Combined with worst-case execution time analysis of the synthesized automaton and continuous control algorithms, software-in-the-loop validation can provide strong evidence of implementability even without physical hardware demonstration. This approach maintains TRL 4 rather than TRL 5, but still validates the synthesis methodology and establishes a clear pathway to hardware deployment. The research contribution remains intact: demonstrating that formal hybrid control synthesis produces implementable controllers, with remaining barriers clearly identified as hardware integration challenges rather than fundamental methodological limitations.

Mitigation strategies leverage existing infrastructure and adopt early testing practices. ARCADE has been successfully used for reactor simulation HIL testing at the University of Pittsburgh, establishing feasibility in principle and providing institutional knowledge about common integration challenges. Conducting early integration testing during the synthesis phase, rather than deferring HIL attempts until late in the project, identifies timing constraints and communication requirements that can inform controller design. Early testing ensures that synthesized controllers

are compatible with hardware limitations from the outset, rather than discovering incompatibilities after synthesis is complete. The Ovation platform supports multiple implementation approaches including function blocks, structured text, and ladder logic, which provides flexibility in how synthesized automata are realized and may enable workarounds if one implementation approach proves problematic.

6 Broader Impacts

Nuclear power presents both a compelling application domain and an urgent economic challenge. Recent interest in powering artificial intelligence infrastructure has renewed focus on small modular reactors (SMRs), particularly for hyperscale datacenters requiring hundreds of megawatts of continuous power. Deploying SMRs at datacenter sites would minimize transmission losses and eliminate emissions from hydrocarbon-based alternatives. However, the economics of nuclear power deployment at this scale demand careful attention to operating costs.

According to the U.S. Energy Information Administration’s Annual Energy Outlook 2022, advanced nuclear power entering service in 2027 is projected to cost \$88.24 per megawatt-hour [14]. Datacenter electricity demand is projected to reach 1,050 terawatt-hours annually by 2030 [15]. If this demand were supplied by nuclear power, the total annual cost of power generation would exceed \$92 billion. Within this figure, operations and maintenance represents a substantial component. The EIA estimates that fixed O&M costs alone account for \$16.15 per megawatt-hour, with additional variable O&M costs embedded in fuel and operating expenses [14]. Combined, O&M-related costs represent approximately 23-30% of the total levelized cost of electricity, translating to \$21-28 billion annually for projected datacenter demand.

This research directly addresses the multi-billion dollar O&M cost challenge through implementations of high-assurance autonomous control. Current nuclear operations require full control room staffing for each reactor, whether large conventional units or small modular designs. These staffing requirements drive the high O&M costs that make nuclear power economically challenging, particularly for smaller reactor designs where the same staffing overhead must be spread across lower power output. By synthesizing provably correct hybrid controllers from formal specifications, we can automate routine operational sequences that currently require constant human oversight. This enables a fundamental shift from direct operator control to supervisory monitoring, where operators can oversee multiple autonomous reactors rather than manually controlling individual units.

The correct-by-construction methodology is critical for this transition. Traditional automation approaches cannot provide sufficient safety guarantees for nuclear applications, where regulatory requirements and public safety concerns demand the highest levels of assurance. By formally verifying both the discrete mode-switching logic and the continuous control behavior, this research will produce controllers with mathematical proofs of correctness. These guarantees enable automation to safely handle routine operations—such as startup sequences, power level changes, and normal operational transitions—that currently require human operators to follow written procedures. Operators will remain in supervisory roles to handle off-normal conditions and provide authorization for major operational changes, but the routine cognitive burden of procedure execution shifts to provably correct automated systems that are much cheaper to operate.

SMRs represent an ideal deployment target for this technology. Nuclear Regulatory Commission certification requires extensive documentation of control procedures, operational requirements, and safety analyses written in structured natural language. As described in our approach,

these regulatory documents can be translated into temporal logic specifications using tools like FRET, then synthesized into discrete switching logic using reactive synthesis tools, and finally verified using reachability analysis and barrier certificates for the continuous control modes. The infrastructure of requirements and specifications is already complete as part of the licensing process, creating a direct pathway from existing regulatory documentation to formally verified autonomous controllers.

Beyond reducing operating costs for new reactors, this research will establish a generalizable framework for autonomous control of safety-critical systems. The methodology of translating operational procedures into formal specifications, synthesizing discrete switching logic, and verifying continuous mode behavior applies to any hybrid system with documented operational requirements. Potential applications include chemical process control, aerospace systems, and autonomous transportation, where similar economic and safety considerations favor increased autonomy with provable correctness guarantees. By demonstrating this approach in nuclear power—one of the most regulated and safety-critical domains—this research will establish both the technical feasibility and regulatory pathway for broader adoption across critical infrastructure.

7 Budget and Budget Justification

7.1 Budget Summary

The proposed research will be conducted over three (3) years, corresponding to the expected completion timeline for the PhD dissertation. Table 1 provides a detailed breakdown of costs by category and year.

7.2 Budget Justification

7.2.1 Senior Personnel

Faculty Advisor Funds are requested to support one month of summer salary per year for the faculty advisor (estimated at Associate Professor level, \$96,459/year base salary for 8 academic months = \$12,083/month). A 4% annual salary increase is applied in subsequent years. The advisor will provide: (1) overall project direction and technical guidance; (2) expertise in control systems theory and formal methods; (3) coordination with industry partners at Emerson and through the Cyber Energy Center; (4) review of all technical deliverables; and (5) mentorship in proposal writing, publication preparation, and professional development.

7.2.2 Other Personnel

Graduate Research Assistant (Principal Investigator) Funds are requested to support one full-time graduate research assistant (the PI) for the entire duration of the project at \$38,000 per year in Year 1. This represents a standard graduate research assistantship stipend at the University of Pittsburgh for a PhD student in the Swanson School of Engineering. A 4% annual salary increase is included in Years 2 and 3 to account for cost-of-living adjustments. The GRA will be responsible for all aspects of the research including: (1) formalizing nuclear reactor operating procedures into temporal logic specifications using FRET; (2) synthesizing discrete automata using reactive synthesis tools (Strix); (3) developing and verifying continuous control modes using reachability analysis and barrier certificates; (4) implementing the hybrid controller on industrial control hardware; and (5) conducting hardware-in-the-loop validation experiments.

Table 1: Proposed Budget by Year and Category

Category	Year 1	Year 2	Year 3	Total
Senior Personnel				
Faculty (PI Advisor, 1 mo.)	\$12,083	\$12,566	\$13,069	\$37,718
Other Personnel				
Graduate Research Assistant	\$38,000	\$39,520	\$41,101	\$118,621
Fringe Benefits				
Faculty Fringe Benefits (29.6%)	\$3,577	\$3,720	\$3,868	\$11,165
GRA Fringe Benefits (50%)	\$19,000	\$19,760	\$20,551	\$59,311
<i>Fringe Benefits Subtotal</i>	\$22,577	\$23,480	\$24,419	\$70,476
Equipment				
(No equipment over \$5,000)	—	—	—	—
Travel				
Conference Travel (Domestic)	\$4,000	\$4,000	\$4,000	\$12,000
Industry Collaboration Visits	\$1,500	\$1,500	\$1,500	\$4,500
<i>Travel Subtotal</i>	\$5,500	\$5,500	\$5,500	\$16,500
Participant Support Costs				
(Not applicable)	—	—	—	—
Other Direct Costs				
<i>Materials and Supplies:</i>				
High-Performance Workstation	\$3,500	—	—	\$3,500
Laboratory Materials & Supplies	\$1,500	\$1,000	\$1,000	\$3,500
<i>Publication Costs</i>	\$1,000	\$1,500	\$2,000	\$4,500
<i>Computing/Cloud Services</i>	\$1,500	\$1,500	\$1,500	\$4,500
<i>Other Direct Costs Subtotal</i>	\$7,500	\$4,000	\$4,500	\$16,000
Total Direct Costs	\$85,660	\$85,066	\$88,589	\$259,315
H. Indirect Costs (F&A)				
On-Campus Research (56% MTDC)	\$35,326	\$34,488	\$35,935	\$105,749
TOTAL PROJECT COST	\$120,986	\$119,554	\$124,524	\$365,064

7.2.3 Fringe Benefits

Faculty Fringe Benefits Faculty fringe benefits are calculated at 29.6%, the University of Pittsburgh's approved rate for academic year faculty, covering retirement contributions, health insurance, and other benefits.

Graduate Research Assistant Fringe Benefits Fringe benefits for the GRA are calculated at 50% of salary, consistent with University of Pittsburgh rates for graduate students on research assistantships. This covers health insurance and other mandatory benefits. Tuition remission is provided separately by the department and is not included in this budget.

7.2.4 Equipment

No equipment purchases over \$5,000 are requested. Computing equipment and software tools under this threshold are budgeted in Materials and Supplies (Section G).

7.2.5 Travel

Conference Travel (\$4,000 per year) Funds are requested for the PI and faculty advisor to attend one major control systems conference annually to disseminate research results. The budget assumes domestic conference attendance with costs including: airfare, hotel, meals and incidentals, ground transportation, and registration for both attendees per conference.

Industry Collaboration Visits (\$1,500 per year) Funds are requested for travel to industry partner sites and potential nuclear facilities to: (1) validate reactor operating procedures with domain experts; (2) present research progress to industry stakeholders; (3) gather feedback on practical implementation considerations; and (4) explore deployment pathways for the developed technology. This includes travel to Emerson facilities or nuclear industry conferences such as the American Nuclear Society Annual Meeting. Budget covers mileage reimbursement at federal rates for regional travel plus hotel and meals for 2-3 one-night trips.

No international travel is planned or budgeted.

7.2.6 Participant Support Costs

Not applicable to this project.

7.2.7 Other Direct Costs

Materials and Supplies *High-Performance Workstation (\$3,500, Year 1):* A dedicated high-performance workstation is required for computationally intensive tasks including: (1) reactive synthesis using Strix, which scales exponentially with specification complexity; (2) reachability analysis for continuous modes, requiring numerical integration of high-dimensional systems; (3) hardware-in-the-loop simulation running at real-time rates with high-fidelity reactor models; and (4) barrier certificate computation using optimization solvers. The workstation specifications include: Intel Core i9 or AMD Ryzen 9 processor (minimum 16 cores), 64 GB RAM, 2 TB NVMe SSD storage, and NVIDIA GPU for potential acceleration of numerical computations. This configuration enables timely completion of verification tasks that would be intractable on standard laboratory computers.

Laboratory Materials and Supplies (\$1,500 Year 1; \$1,000 Years 2–3): Funds are requested for laboratory supplies and materials including: electronic components and sensors for hardware integration, cables and connectors for hardware-in-the-loop setup, and miscellaneous computing

accessories such as external storage devices and backup media. Year 1 costs are higher to account for initial laboratory setup; subsequent years reflect ongoing consumable costs.

Note on Software: All required software tools are either open-source (FRET, Strix, various verification packages) or provided through University of Pittsburgh institutional licenses (MATLAB/Simulink). No software licensing costs are requested. The ARCADE interface software for hardware-in-the-loop testing is also provided through the Cyber Energy Center partnership at no cost to this project.

Note on Equipment: The Emerson Ovation distributed control system hardware required for hardware-in-the-loop validation will be provided through an existing partnership between the University of Pittsburgh Cyber Energy Center and Emerson Process Management. This represents significant in-kind cost sharing and is detailed in Section 7.2.9.

Publication Costs Funds are requested to cover publication fees for disseminating research results in high-quality peer-reviewed venues. Budget includes:

- Year 1 (\$1,000): Conference proceedings fees and one journal submission
- Year 2 (\$1,500): Open-access publication charges for first major journal paper
- Year 3 (\$2,000): Open-access publication charges for dissertation-culminating journal papers

Open-access publication is prioritized to maximize research impact and accessibility, particularly important for work with potential nuclear safety applications. Many high-impact journals (IEEE Transactions on Automatic Control, Automatica) charge \$1,000–\$2,000 for open access.

Computing and Cloud Services Funds are requested for cloud computing resources and online services including: (1) high-performance computing credits for large-scale verification tasks that exceed workstation capabilities; (2) cloud storage and backup services for research data; (3) online collaboration tools and project management platforms; (4) reference management and literature access beyond institutional subscriptions. Cloud computing provides scalable computational resources for particularly demanding verification problems without requiring additional capital equipment purchases.

7.2.8 H. Indirect Costs (Facilities & Administrative)

Indirect costs are calculated at 56% of Modified Total Direct Costs (MTDC), which is the University of Pittsburgh’s federally negotiated rate for on-campus research. MTDC includes all direct costs except equipment purchases over \$5,000, tuition remission, and certain other exclusions. The calculation base includes all personnel costs, travel, and other direct costs as shown in the budget table.

These indirect costs support essential research infrastructure including: laboratory space and utilities in Benedum Hall, administrative support from the Department of Mechanical Engineering and Materials Science, library and information technology services, research compliance and safety oversight, and institutional resources that enable the research enterprise.

7.2.9 Cost Sharing

While no direct cost sharing is proposed or required for this project, significant in-kind contributions will support the research:

Emerson Process Management Partnership Through the University of Pittsburgh Cyber Energy Center, Emerson Process Management will provide access to Ovation distributed control system hardware and ARCADE hardware-in-the-loop interface software. This equipment is essential for TRL 5 validation and represents industry-standard control systems deployed in nuclear facilities. Emerson will also provide technical consultation and domain expertise for practical implementation considerations.

University Infrastructure The University of Pittsburgh provides comprehensive MATLAB/Simulink licenses with all necessary toolboxes through institutional agreements, representing approximately \$3,000–\$4,000 in annual value. Laboratory space in Benedum Hall, high-performance computing resources through the Center for Research Computing, and administrative support are covered through indirect cost recovery and are not included as explicit cost sharing.

Total In-Kind Contributions The estimated value of in-kind contributions over the three-year project duration exceeds \$20,000, primarily from the Emerson partnership. These contributions are essential to achieving project objectives and demonstrate strong industry support for the research.

7.2.10 Budget Inflation and Escalation

Personnel costs include a 4% annual salary increase to account for cost-of-living adjustments, consistent with typical university practices. All other cost categories are held constant in nominal dollars across years, representing a conservative estimate that accommodates minor inflation through efficient resource management.

8 Schedule, Milestones, and Deliverables

This research will be conducted over six trimesters (24 months) of full-time effort following the proposal defense in Spring 2026. The work progresses sequentially through three main research thrusts before culminating in integrated demonstration and validation.

The first semester (Spring 2026) focuses on Thrust 1, translating SmAHTR startup procedures into formal temporal logic specifications using FRET. This establishes the foundation for automated synthesis by converting natural language procedures into machine-readable requirements. The second semester (Summer 2026) addresses Thrust 2, using Strix to synthesize the discrete automaton that defines mode-switching behavior. With the discrete structure established, the third semester (Fall 2026) develops the continuous controllers for each operational mode through Thrust 3, employing reachability analysis and barrier certificates to verify that each mode satisfies its transition requirements.

Integration and validation occupy the remaining three semesters. Spring 2027 integrates the verified hybrid controller into high-fidelity Simulink simulation, achieving TRL 4 through extensive testing across the operational envelope. Summer 2027 implements the controller on Emerson Ovation hardware with ARCADE providing the hardware-in-the-loop interface, demonstrating TRL 5 with actual industrial control equipment. The final semester (Fall 2027) completes validation testing, addresses any remaining issues discovered during hardware testing, and finalizes the dissertation.

Figure 1 shows the complete project schedule including research thrusts, major milestones, and planned publications.




Figure 1: Project schedule showing major research thrusts, milestones (orange row), and publications (green row). Red diamonds indicate completion points. Overlapping bars indicate parallel work where appropriate.

8.1 Milestones and Deliverables

Six major milestones mark critical validation points throughout the research. M1 (Month 4) confirms that SmAHTR startup procedures have been successfully translated to temporal logic using FRET with realizability analysis demonstrating consistent and complete specifications. Domain expert review validates that formalized specifications accurately capture procedural intent. This milestone delivers an internal technical report documenting the translation methodology and complete specification set.

M2 (Month 8) validates computational tractability by demonstrating that Strix can synthesize a complete discrete automaton from the formalized specifications. Success requires an automaton with fewer than 1,000 states generated within 48 hours of computation time, with all operational modes present and interpretable structure. This milestone delivers a conference paper submission to NPIC&HMIT documenting the procedure-to-specification translation methodology.

M3 (Month 12) achieves TRL 3 by proving that continuous controllers can be designed and verified to satisfy discrete transition requirements. At least one complete operational sequence must have continuous controllers for all required modes, with reachability analysis confirming satisfaction of transition requirements and barrier certificates proving safety constraints. This milestone delivers an internal technical report demonstrating component-level verification.

M4 (Month 16) achieves TRL 4 through integrated simulation demonstrating that component-level correctness composes to system-level correctness. The complete hybrid controller must execute autonomous startup sequences in high-fidelity Simulink simulation with zero safety violations across 100 consecutive runs with varied initial conditions. Discrete behavior must match synthesized specifications exactly while continuous trajectories remain within verified reachable sets. This milestone delivers a journal paper submission to IEEE Transactions on Automatic Control presenting the complete hybrid synthesis methodology.

M5 (Month 20) achieves TRL 5 by demonstrating practical implementability on industrial

hardware. The discrete automaton must be implemented on Emerson Ovation distributed control system with ARCADE providing stable real-time communication at required control rates. Complete autonomous startup sequences must execute successfully via hardware-in-the-loop, including correct responses to off-nominal scenarios such as sensor failures and loss of cooling events. This milestone delivers a conference paper submission to NPIC&HMIT or CDC documenting hardware implementation and experimental validation.

M6 (Month 24) completes the dissertation documenting the entire methodology, experimental results, and research contributions. All verification artifacts including specifications, automata, reachable sets, and barrier functions are publicly released along with implementation code for both simulation and hardware integration. This enables independent replication and provides the nuclear industry with practical tools for developing autonomous control systems.

The publication strategy targets both the nuclear instrumentation community and the broader control systems research community. The first conference paper (Summer 2026) presents the procedure formalization methodology to NPIC&HMIT, building on a previous best student paper award received at that conference. This paper addresses the practical challenge of translating natural language procedures into formal specifications, presenting a taxonomy of procedural ambiguities and systematic resolution strategies that will be valuable to the nuclear I&C community.

The journal paper (Spring 2027) targets IEEE Transactions on Automatic Control to establish theoretical rigor and reach the broader control systems community. This paper presents the complete hybrid synthesis methodology as a contribution to control theory, demonstrating how reactive synthesis and continuous verification can be unified to address continuous autonomous hybrid systems. The three-mode classification and compositional verification framework have applications beyond nuclear systems to any safety-critical domain with documented operational procedures.

The final conference paper (Summer 2027) documents hardware implementation and experimental validation, demonstrating that the methodology produces controllers compatible with existing industrial control infrastructure. Venue selection between NPIC&HMIT and IEEE Conference on Decision and Control depends on whether emphasis is placed on nuclear industry adoption or experimental validation of theoretical control methods. This decision will be made based on results from preliminary hardware testing and feedback from the TRL 4 milestone review. Additional dissemination includes release of all software tools and verification artifacts.

References

- [1] U.S. Nuclear Regulatory Commission. Operators' licenses. 10 CFR Part 55. Code of Federal Regulations.
- [2] John G. Kemeny et al. Report of the president's commission on the accident at three mile island. Technical report, President's Commission on the Accident at Three Mile Island, October 1979.
- [3] U.S. Nuclear Regulatory Commission. Guidelines for the preparation of emergency operating procedures. Technical Report NUREG-0899, U.S. Nuclear Regulatory Commission, 1982.
- [4] U.S. Department of Energy. Human performance handbook. Handbook DOE-HDBK-1028-2009, U.S. Department of Energy, 2009.
- [5] World Nuclear Association. Safety of nuclear power reactors. <https://www.world-nuclear.org/information-library/safety-and-security/safety-of-plants/safety-of-nuclear-power-reactors.aspx>, 2020.
- [6] International Atomic Energy Agency. Human error as root cause in severe nuclear accidents. IAEA Safety Report. Analysis of TMI, Chernobyl, and Fukushima accidents.
- [7] Y. Wang et al. Analysis of human error in nuclear power plant operations: A systematic review of events from 2007–2020. *Journal of Nuclear Safety*, 2025. Analysis of 190 events at Chinese nuclear power plants.
- [8] D. Gertman et al. The spar-h human reliability analysis method. Technical Report NUREG/CR-6883, U.S. Nuclear Regulatory Commission, 2005.
- [9] U.S. Nuclear Regulatory Commission. Cognitive basis for human reliability analysis. Technical Report NUREG-2114, U.S. Nuclear Regulatory Commission, 2016.
- [10] J. Rasmussen. Skills, rules, and knowledge; signals, signs, and symbols, and other distinctions in human performance models. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-13(3):257–266, 1983.
- [11] George A. Miller. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63(2):81–97, 1956.
- [12] James Reason. *Human Error*. Cambridge University Press, 1990.
- [13] Joseph Kiniry, Alexander Bakst, Michal Podhradsky, Simon Hansen, and Andrew Bivin. High assurance rigorous digital engineering for nuclear safety (hardens) final technical report. Technical Report ML22326A307, Galois, Inc. / U.S. Nuclear Regulatory Commission, 2022. NRC Contract 31310021C0014.
- [14] U.S. Energy Information Administration. Levelized costs of new generation resources in the annual energy outlook 2022. Report, U.S. Energy Information Administration, March 2022. See Table 1b, page 9.

[15] Environmental and Energy Study Institute. Data center energy needs are upending power grids and threatening the climate. Web article, 2024. Accessed: 2025-09-29.