Model-Based Scenario Identification for Guided Dynamic Probabilistic Risk Assessment

Nana Yaw Amanyi Angu[a] and Vincent P. Paglioni[b]
[a] Risk, Reliability, & Resiliency Characterization Lab, Department of Systems Engineering, Colorado State University. Fort Collins, CO, USA. Nanayaw.angu@colostate.edu
[b] Risk, Reliability, & Resiliency Characterization Lab, Department of Systems Engineering, Colorado State University. Fort Collins, CO, USA. Vincent.paglioni@colostate.edu

Abstract: Dynamic Probabilistic Risk Assessment (DPRA) provides a framework for analyzing time-dependent accident progression and coupled system behavior in nuclear power plants. However, the need to explore large scenario and uncertainty spaces leads to significant computational demands, limiting the scalability of DPRA. Methods that efficiently identify risk-significant conditions and focus analysis on these could reduce this burden, but developing such approaches remain a key challenge. This paper presents a first step toward addressing this challenge by proposing a model-based methodology for guided DPRA. To support this objective, an executable system architecture model is developed using Model-Based Systems Engineering (MBSE). The model captures system requirements, structure, functions, and time-dependent behavior while embedding qualitative risk information such as failure modes and fault logic. Low-fidelity simulations performed on the resulting model enable the identification of risk-significant accident scenarios and parameter combinations that can guide the application of high-fidelity DPRA simulations. The proposed methodology establishes a foundation for improving the efficiency and scalability of DPRA while maintaining traceability to evolving system definitions and requirements. The methodology is demonstrated through a proof-of-concept case study of a representative safety system, illustrating how the architecture model supports guided DPRA simulation and risk-informed design. Future work will integrate high-fidelity simulations by using the scenario insights generated in this work to guide more detailed DPRA.

1. INTRODUCTION

Nuclear Power Plants (NPPs) require rigorous risk analysis to ensure that accident progression and system failures are adequately understood, managed, and effectively communicated to multiple stakeholders. Probabilistic Risk Assessment (PRA) provides such a method for assessing the risk associated with NPPs by quantifying the risk triplet [1] defined as: (1) what can go wrong, (2) how likely is it, and (3) what are the consequences. Static PRA approaches based on fault trees and event trees have been historically applied in this context but are limited in their capability to capture time-dependent system dynamics and the coupling between deterministic physical behavior and stochastic failure events. Dynamic PRA addresses these limitations by explicitly performing risk assessment within a combined deterministic and stochastic framework, enabling a more complete and less conservative representation of system behavior under uncertainty [2].

Despite its advantages, DPRA faces a fundamental scalability challenge. The method requires propagating uncertainty through time-dependent system models over large numbers of scenarios. Each scenario involves a simulation of accident progression and when high-fidelity thermal-hydraulic codes are used as the simulation engine, the computational cost of a single scenario evaluation can be significant, and the cost of exploring the full uncertainty space becomes prohibitive [3], [4]. This computational challenge is even more relevant for passive systems where the cost and time required to execute high-fidelity codes may outweigh the practical value given their novelty and limited operational data availability [5]. A methodology that enables efficient risk screening of such systems within a structured and traceable framework would support more informed design decisions.

This paper proposes such a methodology in which a low-fidelity, physics-based model constructed from first principles is embedded within a system architecture model developed using Model-Based Systems Engineering (MBSE). The MBSE model captures system requirements, structure, functions, and time-

dependent behavior while embedding qualitative risk information including failure modes and fault logic, capabilities that align with the growing adoption of MBSE in the nuclear industry for design and risk analysis tasks [6], [7]. Recent applications of MBSE to passive heat removal systems have demonstrated its value for structured design requirement analysis [8] and qualitative safety assessment including failure mode identification [5]. This work builds on that foundation by extending the MBSE framework to quantitative probabilistic screening within a DPRA context. The low-fidelity model is derived from governing physical equations, requires no high-fidelity runs to construct, and is available from the outset of the design process. Uncertainty is propagated through the model using Monte Carlo simulation, identifying risk-significant parameter combinations that can guide subsequent high-fidelity assessment. The methodology is demonstrated on the Reactor Cavity Cooling System (RCCS) of the Horizontal Compact High Temperature Gas Reactor (HC-HTGR), a passive natural circulation decay heat removal system whose functional behavior is governed by well-defined physical principles, making it a natural candidate for physics-based low-fidelity screening. This work also provides a concrete example of how quantitative probabilistic screening can be performed directly from a system architecture model.

The remainder of the paper is organized as follows. Section 2 describes the methodology, including the HC-HTGR RCCS system description, the MBSE architecture model, the low-fidelity physics model, and the Monte Carlo uncertainty propagation framework. Section 3 presents the results, including model validation against Argonne National Laboratory's RELAP5-3D reference and the identified failure boundary in the uncertain parameter space. Section 4 summarizes conclusions and outlines directions for future work.

2. METHODOLOGY

2.1. System Description

The HC-HTGR is a prismatic high-temperature gas-cooled reactor in a horizontal layout with a total thermal power of 160-230 MWt, currently under development as part of a DOE-NE Advanced Reactor Concepts program [9]. The reactor pressure vessel (RPV) has an outer diameter of 4.1m and a total length of 12m, housed within a reactor room of 9.2m width and 7.9m height. The RCCS is the primary passive safety system responsible for removing decay heat from the RPV following shutdown, relying entirely on natural forces without active components or external power. Figure 1 displays a simplified version of the design of the HC-HTGR RCCS [9].

Figure 1. Conceptual design of the HC-HTGR RCCS [9].

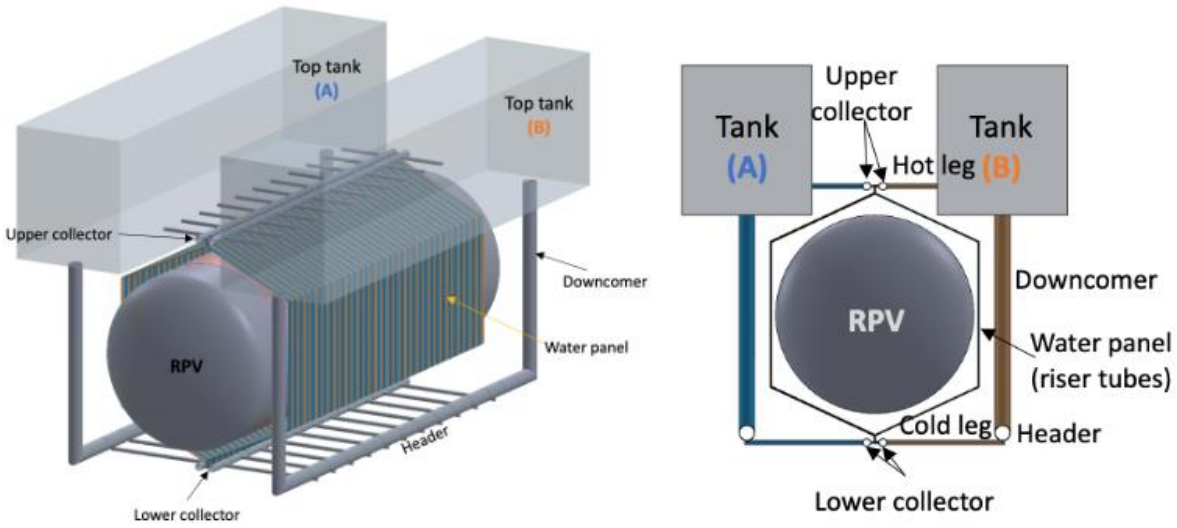

The RCCS is a water-based natural circulation system comprising two independently operating loops, each consisting of a water panel located on one side of the RPV, a top water storage tank, and connecting piping. Heat is transferred from the RPV outer surface to the water panel primarily by thermal radiation, driving natural circulation of the coolant through the loop. Each water panel consists of 66 Schedule 160 one-inch carbon steel riser tubes connected by a carbon steel fin panel, with a reference surface emissivity of 0.8. The system is designed to remove approximately 1 MWt at the limiting design condition while maintaining the RPV temperature below 482°C during a depressurized conduction cooldown event. This condition forms the governing safety requirement for this analysis [9]. The water inventory is sized to sustain passive cooling for a minimum of seven days without replenishment.

2.2. MBSE Architecture Model

A SysML architecture model of the HC-HTGR RCCS was developed to serve as the integrating structure for the guided DPRA methodology. The model was structured around a formal safety requirement that the RPV temperature shall not exceed 482°C during a depressurized conduction cooldown event derived directly from the HC-HTGR RCCS design requirements [9]. Figure 2a shows this requirement and traces it to the maximum vessel temperature, defined as a value property of the analysis. Figure 2b shows the RCCS decomposed into functional blocks representing the key heat transfer interfaces, that is radiation from the RPV to the water panel, convection from water panel to coolant, and natural circulation within the riser tubes. Value properties are also used to define the parameters necessary for each functionality.

Figure 2. (a) SysML requirement traceability to RPV temperature; (b) Functional decomposition of the RCCS.

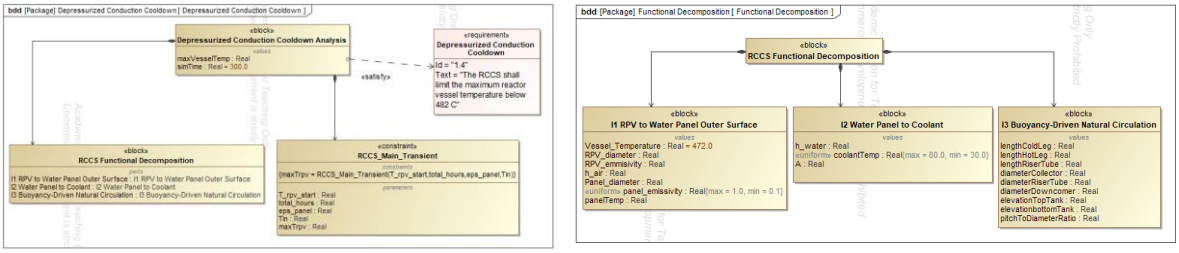

Functional failures were identified for each block and linked to the relevant system interfaces through a dependency relationship, embedding qualitative risk information directly within the system definition. Figure 3a captures this information for the radiation heat transfer interface. The SysML structure diagram of the RCCS, showing the various subsystems and components is captured in Figure 3b. Values in brackets represent the number of such components.

Figure 3. (a) Functional failures for radiation heat transfer interface; (b) SysML structural diagram of the RCCS.

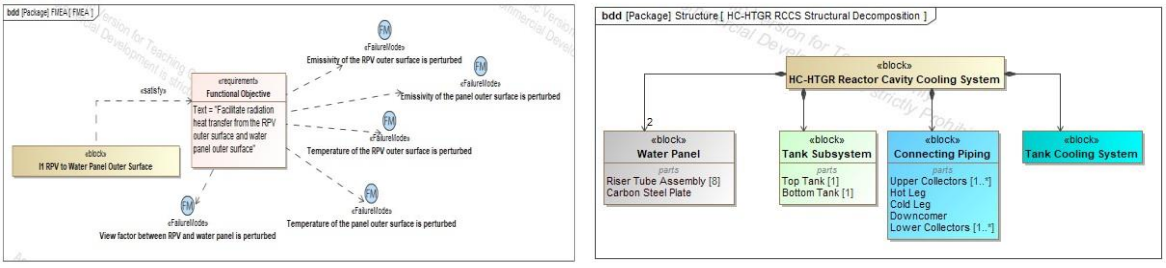

Figure 4 shows the parametric diagram, which connects input value properties defined in Figure 2b to constraint equations implemented in the Python physics model. This interface enables SysML to pass scalar values of the uncertain parameters to Python at each Monte Carlo sample and receive the maximum RPV temperature over the transient in return. At each sample, SysML evaluates whether the returned maximum RPV temperature satisfies the stated safety requirement.

Figure 4. Binding of value properties to low-fidelity model through SysML Parametric Diagram.

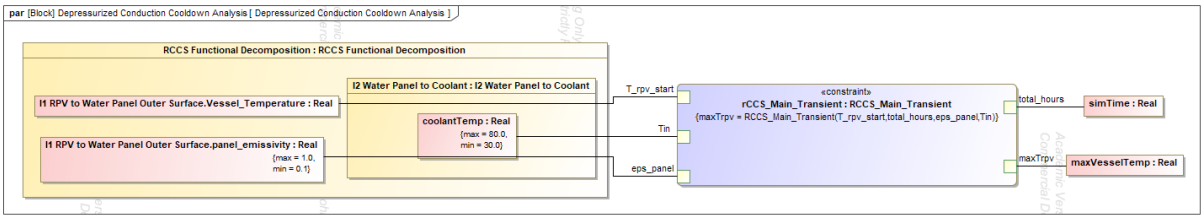

2.3. Low-Fidelity Physics Model

The low-fidelity model follows a system dynamics simulation approach, in which the time-dependent thermal behavior of the RCCS is described through a set of coupled first order differential equations governing conservation, radiation heat transfer, and natural circulation momentum balance. The model is implemented in Python and structured around two coupled subproblems, (1) a thermal resistance network that determines the heat removal rate from the RPV and (2) a natural circulation momentum balance that determines the natural circulation mass flow rate. These two subproblems are mutually independent and are resolved through an iterative outer coupling loop.

2.3.1. Thermal Resistance Network

Heat transfer from the RPV outer surface to the RCCS coolant occurs through two resistances in series. Thermal radiation across the gap between the RPV and the water panel, and natural convection from the panel to the coolant within the riser tubes. Gap air convection is neglected. The panel is treated as a fin with efficiency $\eta = 1.0$, meaning perfect thermal coupling between the panel surface and the riser tube wall is assumed. Radiative heat transfer from the RPV to the water panel is given by Equation 1, where $\sigma = 5.67\times10^{-8}$ W/m²K⁴ is the Stefan-Boltzmann constant, A_{rpv} is the RPV radiating surface area per sector, T_{rpv} is the RPV outer wall temperature, T_{panel} is the panel surface temperature, ε_{rpv} is the RPV surface emissivity, and ε_{panel} is the panel emissivity.

$$Q_{rad} = \frac{\eta \cdot \varepsilon_{RPV} \cdot \sigma \cdot A_{RPV} \cdot \left(T_{RPV}^4 - T_{panel}^4\right)}{\frac{1}{\varepsilon_{RPV}} + \frac{1-\varepsilon_{panel}}{\varepsilon_{panel}} + \frac{D_{RPV}}{D_{panel}}} \quad (1)$$

The natural convection heat transfer from the panel to the coolant is described by the Newton law of cooling given by Equation 2.

$$Q_{conv} = \left(T_{panel} - T_{coolant}\right) * h_{tc} * A_{conv} \quad (2)$$

where $T_{coolant}$ is the coolant inlet temperature of the natural circulation loop, h_{tc} is the convective heat transfer coefficient inside the riser tubes, and A_{conv} is the total convective surface area of the riser tubes. The convective heat transfer coefficient is computed as:

$$h_{tc} = \frac{Nu * k}{D_r} \quad (3)$$

Equations (1) and (2) form a circular dependency, Q_{rad} depends on T_{panel} through equation (1), and T_{panel} depends on Q_{conv} through equation (2). This dependency is resolved by successive substitution. Starting from an initial guess $T_{panel} = (T_{RPV} + T_{coolant})/2$, equations (1) and (2) are alternated until T_{panel} converges, at which radiation in equals convection out. In Equation 3, Nu is determined by a three-regime approach based on the local Reynolds number. $Nu = 4.36$ for Re ≤ 2300 (fully developed

laminar flow). In the transitional regime, the Gnielinski [10] correlation is used for Nu, and the Dittus-Boelter correlation is incorporated for Re ≥ 4000.

2.3.2. Natural Circulation Momentum Balance

The natural circulation flow rate is determined by requiring that the net pressure driving force balances all pressure losses around the loop. This momentum residual is expressed in Equation 4a, with the components further detailed in Equations 4b – 4e.

$$\Delta P_{buoyancy} - \left(\Delta P_{friction} + \Delta P_{acceleration} + \Delta P_{minor}\right) = 0 \tag{4a}$$

where:

$$\Delta P_{buoyancy} = (\rho_{in} - \rho_{out}) * g * \Delta H \tag{4b}$$

$$\Delta P_{friction} = f \frac{\sum L \rho v^2}{2 D_h} \tag{4c}$$

$$\Delta P_{acceleration} = \dot{m}^2 \left(\frac{1}{\rho_{in}} - \frac{1}{\rho_{out}}\right) \tag{4d}$$

$$\Delta P_{form} = \sum K_j \frac{\rho v^2}{2} \tag{4e}$$

The momentum residual in Equation 4a is a nonlinear function of ṁ (see Equation 4d) because the fluid velocities, densities, and outlet temperature all depend on the flow rate. An analytical solution is therefore not available, and so the equation is solved numerically using the Python root-finding function *scipy.optimize.brentq*, which implements a combination of root bracketing, interval bisection, and inverse quadratic interpolation methods. Given a bracketed interval for $\dot{m}$, brentq iteratively evaluates the momentum residual (Equation 4a) until it equals zero, meaning buoyancy exactly balances all pressure losses and the natural circulation flow rate is established. It is the nested computational structure of this iterative solution, repeated across the full uncertain parameter space, that makes direct evaluation during Monte Carlo simulation computationally expensive and motivates the lookup table approach described in Section 2.3.5.

2.3.3. Outer Cooling Loop

Because the convective heat transfer depends on ṁ through the Reynolds number, and ρ_{out} depends on the natural circulation loop outlet temperature which in turn depends on Q_{rad}, the thermal resistance network and the momentum balance are mutually coupled. This coupling is resolved through an outer iteration loop that alternates between the two subproblems. At each outer iteration, h_{tc} is first updated from the current ṁ, using the three-regime Nu correlation. The heat balance is resolved by iterating Equations 1 and 2 until change in panel temperature is less than the tolerance value. The outlet temperature and density are computed from the converged Q_{rad} and T_{panel}. brentq is then called to find the exact ṁ that reduces the momentum residual to zero. The converged heat removal rate for the full system is then given by Equation 5.

$$Q_{RCCS} = Q_{rad} * 16 \tag{5}$$

The multiplier of 16 in Equation 5 accounts for the eight riser tube assemblies on each of the two sides of the RPV, consistent with the scoping calculations conducted by Argonne National Laboratory [9].

2.3.4. Transient Energy Balance

The time-dependent evolution of the RPV temperature following shutdown is governed by the energy balance in Equation 6.

$$\frac{dT_{RPV}}{dt} = \frac{Q_{decay}(t) - Q_{system}(t)}{M_{RPV} * C_{p,steel} + M_{graphite} * C_{p,graphite}} \quad (6)$$

In Equation 6, M_{RPV} and $M_{graphite}$ are the masses of the RPV steel and graphite core respectively, and $C_{p,steel}$ and $C_{p,graphite}$ are their specific heat capacities. Q_{system} is not itself time-dependent, but is re-evaluated at each timestep from the current T_{RPV} to capture the physical dependency of heat removal on the instantaneous RPV temperature. When Q_{decay} exceeds Q_{system}, the RPV heats up. The RPV cools down when Q_{system} exceeds Q_{decay}. Equation (6) is integrated numerically using a forward Euler method with a timestep of 600 seconds over a 300-hour simulation window. The decay heat is described by the Wigner-Way formula in Equation 7, where P_o= 160 MW is the nominal reactor power and $t_{operated}$ = 1.55×10^7 s is the reactor operating time prior to shutdown.

$$Q_{decay}(t) = 0.0622 * P_o \left[t^{-0.2} - \left(t_{operated} + t\right)^{-0.2}\right] \quad (7)$$

2.3.5. Lookup Table for Computational Efficiency

The nested computational structure of the physics model, the heat balance loop running inside every evaluation, is repeated across the outer coupling iteration. This means that a single call to the low-fidelity model involves a substantial number of operations. Direct evaluation at every timestep of every Monte Carlo (MC) scenario would therefore be computationally prohibitive. To address this, Q_{system} is pre-computed over a three-dimensional grid spanning $T_{RPV} \in [200,510]$℃, $T_{coolant} \in [20,45]$℃, and $\varepsilon_{panel} \in [0.1,1]$, producing a lookup table of 6,400 entries built once before the MC simulation begins. During the simulation, trilinear interpolation using scipy's RegularGridInterpolator replaces the physics solver at each timestep. The accuracy of the lookup table is verified against 20 randomly selected direct physics evaluations prior to the Monte Carlo run, with a maximum interpolation error threshold of 1%. Water properties k, μ, and ρ are evaluated from polynomial fits to IAPWS-IF97 data with maximum errors of 0.7%, 0.3%, and 0.02% respectively over the temperature range of interest [11]. The specific heat Cp is fixed at 4178 J/kgK with variation less than 0.3% over the full temperature range encountered in the model.

3. RESULTS

3.1. Model Validation

The low-fidelity physics model developed in this work was validated against the RELAP5-3D reference result reported by Argonne National Laboratory for the pre-conceptual HC-HTGR RCCS design [9]. At the design condition of T_{RPV} = 472°C, $T_{coolant}$ = 30°C, and ε_{panel} = 0.8, the model predicts a system heat removal rate of 0.566 MW against the ANL RELAP5-3D result of 0.56 MW, representing an error of about 1.0%. The converged natural circulation mass flow rate is 0.41 kg/s with a Reynolds number of 2925, placing the flow in the transitional regime. The 1.0% error therefore reflects that this low-fidelity model provide an accurate representation of the dominant physics at the ANL design point rather than a fitted result and provides confidence that the model correctly captures the radiation-dominated heat transfer mechanism at the condition for which the RCCS was sized.

The model was further assessed at the normal operation condition (T_{RPV} = 220°C, $T_{coolant}$ = 30°C, ε_{panel} = 0.8), where ANL reports a total heat removal from the RPV of 0.170 MW. The model predicts 0.095 MW at this condition, an underestimate of 44% relative to the ANL total. Both validation results are summarised in Table 1. The underestimate at 220°C is explained by the omission of air gap convection, which is present in the ANL RELAP5-3D model but excluded from the low-fidelity model. According to the ANL study, at 472°C, air convection contributes 27.9% of total RPV heat removal and rises to 44.9% at 220°C. Radiation heat transfer weakens significantly between these two temperature conditions due to the T⁴ dependence, whereas air convection depends on external natural

convection and weakens more gradually. As a result, the air contribution becomes a larger fraction of a smaller total at 220°C, and the model underestimates total heat removal more substantially. This omission is conservative for DPRA purposes, the model underestimates total RCCS heat removal capacity, meaning the true failure threshold lies below $\varepsilon_{\text{panel}} = 0.415$, meaning the system can tolerate greater emissivity degradation than the model predicts.

Table 1. Low-fidelity model validation against ANL RELAP5-3D Model.

Condition	T_{RPV}	$T_{coolant}$	Model [MW]	ANL Total [MW]	Error	Air fraction
Design condition	472°C	30°C	0.566	0.560	1.0%	27.9%
Normal operation	220°C	30°C	0.095	0.170	44%	44.9%

3.2. Failure Boundary and Monte Carlo Screening

The Monte Carlo screening analysis was conducted over N = 10,000 samples drawn from uniform distributions over panel emissivity ∈ [0.10, 0.85] and coolant inlet temperature ∈ [20, 45] °C. Two uncertain parameters were selected for this screening analysis. Panel emissivity was identified as the primary epistemic uncertainty from the ANL parametric study, which demonstrated that thermal radiation is the dominant heat transfer mechanism and that surface condition is the key design sensitivity [9]. Coolant inlet temperature captures the effect of tank thermal conditioning prior to the accident. Other sources of uncertainty including water inventory, decay heat curve, and loop redundancy are acknowledged as important and are identified as extensions for future work. The computational advantage of the low-fidelity model is throughput. The full 10,000-sample Monte Carlo run completes in approximately 180 seconds on a workstation laptop (Intel Core Ultra 7 165H, 22 cores, 32 GB RAM). The low-fidelity model does not replace RELAP5-3D but screens the uncertain parameter space to identify the failure boundary and the high-risk region, which would then be targeted by high-fidelity simulation in a complete guided DPRA workflow. The depressurized conduction cooldown (DCC) event was simulated starting from the conservative design condition of T_{RPV} = 472°C, consistent with the pre-conceptual limiting design condition reported by ANL. The DCC heat load profile applied to the RCCS was reconstructed using guidance from the ANL report, where the load starts at 0.22 MW (equal to the normal-operation parasitic heat loss), decreases to 80% within the first 7 hours following shutdown, rises to a peak of 1.53*parasitic heat loss at approximately 110 hours, then gradually decays. The failure criterion is T_{RPV} > 482°C, corresponding to ANL design requirement 1.6 for DCC events.

Figure 5 presents the Monte Carlo results. The estimated failure probability is = 42.92% across the full sampled parameter space. This figure reflects the fraction of the deliberately broad epistemic uncertainty range that leads to failure, not the operational probability of failure in service. The sampling range $\varepsilon_{\text{panel}} \in [0.10, 0.85]$ is intentionally wide to map the full failure boundary for screening purposes. At the nominal design emissivity of 0.8 the system is comfortably safe, and the result quantifies how far conditions must degrade before failure occurs. The Spearman rank correlation coefficient between the panel emissivity and maximum RPV temperature is –0.919, confirming that panel emissivity is the dominant risk driver. By contrast, the Spearman rank correlation between coolant inlet temperature and maximum RPV temperature is 0.009, confirming that coolant inlet temperature has negligible influence on the failure outcome at this level of analysis. This is consistent with the ANL finding that the thermal performance of the RCCS is primarily radiation-limited, with convection playing a secondary role.

The bisection-based limit surface search identifies a minimum emissivity threshold of 0.415–0.417 across the full temperature range. The nominal panel emissivity of 0.8 provides a safety margin of (0.8 − 0.415)/0.8 = 48% relative to the limit surface, meaning the emissivity must degrade to less than half its nominal value before the system fails to protect the RPV during DCC.

Figure 5. Monte Carlo DPRA screening results. (a) Failure boundary in parameter space with limit surface. (b) Distribution of peak RPV temperature across all samples.

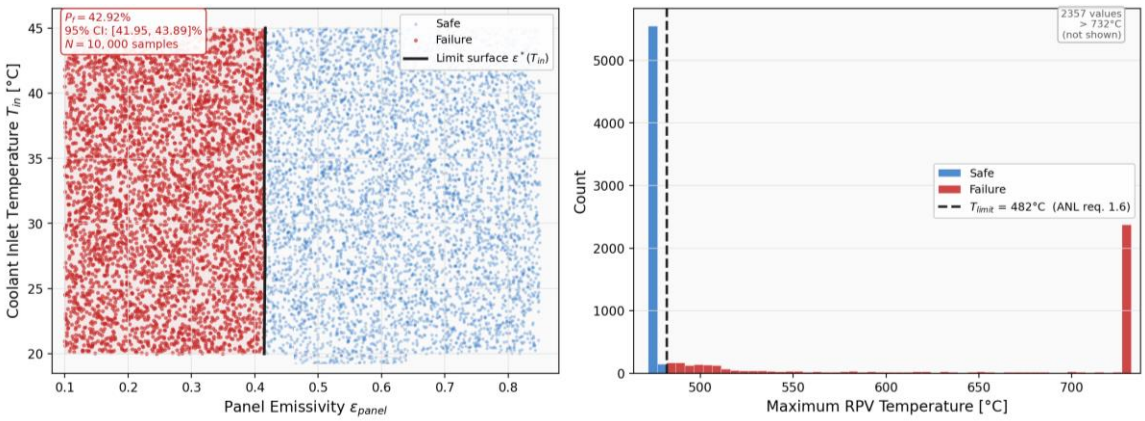

3.3. Natural Circulation Physics Maps

Figure 6 shows the natural circulation flow regime and mass flow rate map across the parameter space at the design RPV temperature of 472°C. At low emissivity and high inlet temperatures, the Reynolds number falls below 2,300 and the flow transitions to the laminar regime. This transition has a direct bearing on the convective heat transfer coefficient. The laminar Nusselt number of 4.36 (fully developed, constant heat flux) is significantly lower than the turbulent Dittus-Boelter value, reducing heat removal by the RCCS beyond what the emissivity reduction alone implies. The failure boundary (red line) lies entirely within the laminar-to-transitional regime, indicating that flows in the failure region are governed by laminar heat transfer correlations. This is consistent with the low natural circulation flow rates, where reduced radiation limits the buoyancy driving force.

Figure 6. Natural circulation flow regime and mass flow rate map at $T_{RPV} = 472$ °C. Filled regions indicate laminar (blue), transitional (orange), and turbulent (green) regimes. Contour lines show mass flow rate in kg/s. The red line is the failure limit surface.

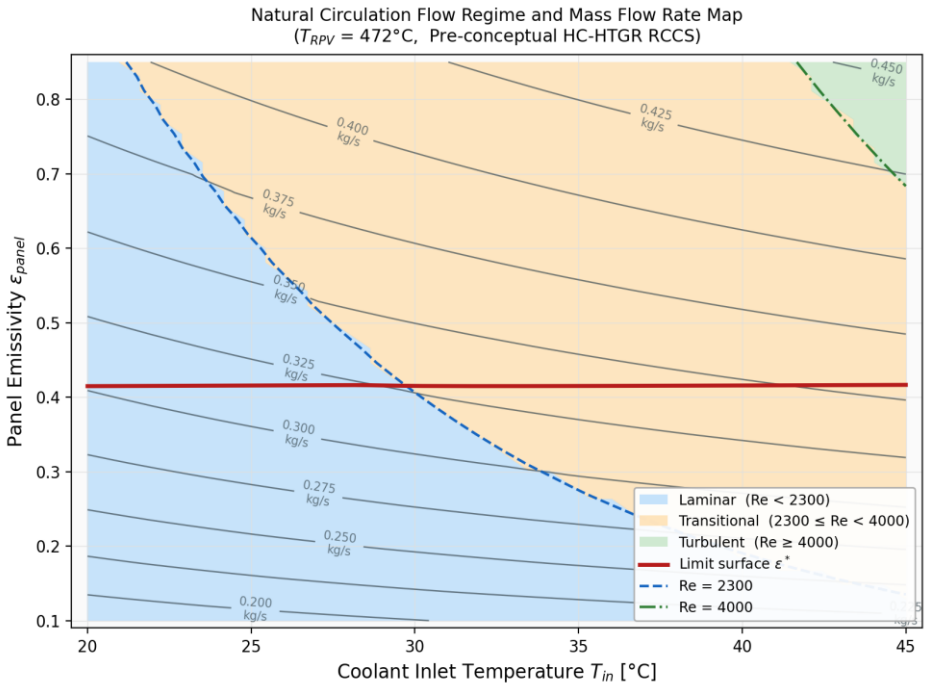

Figure 7 presents the coolant outlet temperature and RCCS heat removal margin maps. Plot (a) shows that the outlet temperature increases with both T_{RPV} and $\varepsilon_{\text{panel}}$ reaching a maximum of approximately 66°C at the highest emissivity and inlet temperature in the sampled space. The coolant remains well within the single-phase regime across the entire parameter space at the design condition, confirming that the single-phase flow assumption used in the model is valid throughout. Plot (b) shows the difference between heat removal by the RCCS and the peak DCC load of 336.6 kW across the parameter space. Green regions (positive margin) indicate parameter combinations where the RCCS can remove more heat than the DCC peak load at the design condition, while red regions (negative margin) indicate insufficient capacity. The most significant feature in plot (b) is the gap between the dashed black line and the red failure boundary. This gap arises because the DCC heat demand does not start at its peak. At t = 0 the demand is only 0.22 MW, well within the removal capacity even at low emissivity and does not reach its peak of 0.337 MW until after 110. Scenarios in the gap $0.415 < \varepsilon_{\text{panel}} < 0.47$ therefore spend the first 110 hours cooling down, building thermal margin against the eventual peak. By the time the peak arrives, the RPV temperature has already fallen below 472°C and the subsequent temperature rise will be below the 482°C limit. A static analysis, which evaluates only whether heat removal by the RCCS can match the peak load, misses this time-dependent recovery entirely and incorrectly classifies these scenarios as failures. This gap is the quantitative demonstration of the value of DPRA over static analysis. At nominal emissivity of 0.8 the heat removal margin is approximately +230 kW.

Figure 7. Thermal performance maps at $T_{RPV} = 472$ °C. (a) Coolant outlet temperature. (b) RCCS heat removal margin.

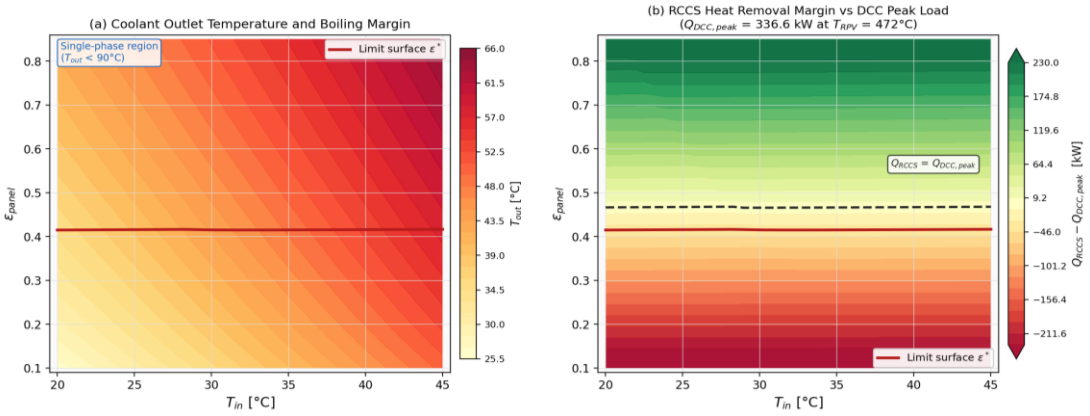

Figure 8. Transient RPV temperature for three scenarios.

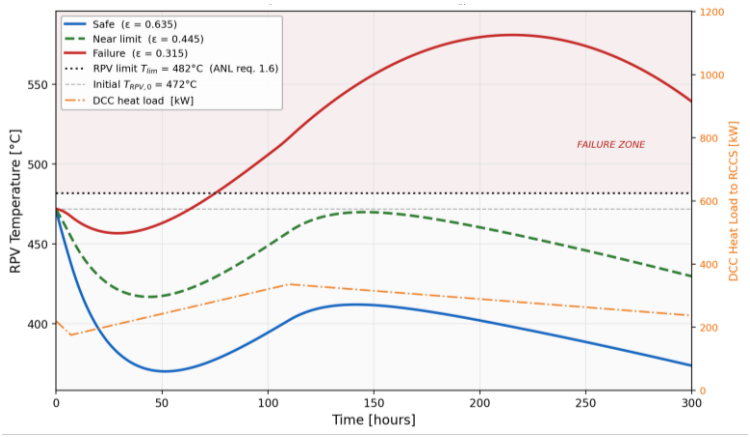

3.4. Transient Scenario Analysis

Figure 8 shows how the RPV temperature evolves over 300 hours for three scenarios at $T_{coolant}$ = 35°C. The three emissivity values were chosen as deliberate offsets from the limit surface ε = 0.415, with one well above it ($\varepsilon_{\text{panel}}$ = 0.635), one just above it ($\varepsilon_{\text{panel}}$ = 0.445), and one below it ($\varepsilon_{\text{panel}}$ = 0.315). The DCC heat demand on the secondary axis provides the time-varying load context starting at 0.22 MW at shutdown, dipping briefly at 7 hours, peaking at 0.337 MW around 110 hours, and gradually decaying thereafter.

The safe and failure scenarios tell a straightforward story. At $\varepsilon_{panel} = 0.635$ the RPV cools steadily throughout the event, staying well below the 482°C limit with a large margin. At $\varepsilon_{panel} = 0.315$, the RPV temperature rises continuously after a small dip, crossing the 482°C limit at approximately 95 hours and reaching 535.4°C by the end of the simulation. Both outcomes are clear and quite robust as the margins in both directions are large enough that no plausible modelling uncertainty would reverse them. The two regions are effectively resolved by the low-fidelity screening approach. The near-limit temperature evolution at $\varepsilon_{panel} = 0.445$ is where the analysis becomes interesting. This scenario sits inside the gap $0.415 < \varepsilon_{panel} < 0.47$ that Figure 7b identified between the DPRA failure boundary and the static balance line. The transient shows what that gap looks like in practice. The RPV actually cools during the first 110 hours because the early DCC demand is well within the RCCS capacity, then warms slightly as the peak demand arrives, and settles at a maximum of 471.6°C just 10.4°C below the failure limit. The low-fidelity model says this scenario is safe, but 10.4°C margin is not enough confidence given the modelling uncertainties. This presents an interesting region where high-fidelity simulation runs can be used to probe further and support risk-informed decision-making. The low-fidelity model screened everything below $\varepsilon_{panel} = 0.415$ as a failure case, everything above $\varepsilon_{panel} = 0.47$ as a safe case, and the band in between these values is where high-fidelity resources must be leveraged to gain further understanding.

4. CONCLUSIONS

This paper presented a model-based methodology for guided Dynamic Probabilistic Risk Assessment applied to the Reactor Cavity Cooling System of the Horizontal Compact High Temperature Gas Reactor. A SysML architecture model was developed to formally trace the governing safety requirement, decompose the system into functional blocks with embedded failure mode information, and interface with a low-fidelity Python physics model through a parametric diagram. The physics model couples a thermal resistance network for radiative and convective heat transfer with a natural circulation momentum balance, solved iteratively, and integrates a transient energy balance over a 300-hour depressurized conduction cooldown simulation window.

Monte Carlo screening over N = 10,000 samples identified a failure probability of 42.92% across the full uncertain parameter space, with a limit surface at emissivity value ≈ 0.415–0.417 that is nearly independent of coolant inlet temperature. The nominal panel emissivity of 0.8 provides a 48% margin above the failure threshold, quantifying the tolerance of the system to emissivity degradation through possible corrosion and fouling. The transient analysis demonstrated that the DPRA methodology captures physically meaningful distinctions between safe, near-limit, and failure scenarios that are inaccessible to static analysis, including the critical near-limit case in which a scenario that would appear to fail a static peak-load check is in fact safe when the time-dependence of the DCC heat load is accounted for.

The low-fidelity model required no high-fidelity simulation runs to construct, and the full 10,000-sample Monte Carlo analysis completed in approximately 90 seconds, demonstrating the computational efficiency benefit of the proposed screening approach. The results carry a direct design implication, the panel emissivity must remain above 0.415 to ensure RPV protection during DCC, providing a quantitative basis for surface treatment and inspection requirements. The 48% margin between the nominal emissivity and the failure threshold quantifies the system tolerance to corrosion and fouling

degradation over the 40-year design life. This paper constitutes the first stage of a two-stage guided DPRA workflow. This first attempt focuses on establishing the methodology and demonstrating the screening capability. Future work will extend the analysis to additional uncertain parameters including water inventory, decay heat curve uncertainty, and loop redundancy failure modes, and will integrate high-fidelity simulations targeted at the identified risk region to complete the guided DPRA workflow.

ACKNOWLEDGEMENTS

The authors acknowledge the use of the publicly available conceptual design documentation for the HC-HTGR RCCS (ANL/NSE-23/89) as the reference basis for all system parameters and design requirements used in this work.

REFERENCES

[1] S. Kaplan and B. J. Garrick, "On The Quantitative Definition of Risk," *Risk Analysis*, vol. 1, no. 1, pp. 11–27, Mar. 1981, doi: 10.1111/j.1539-6924.1981.tb01350.x.

[2] D. Mandelli, C. Parisi, N. Anderson, Z. Ma, and H. Zhang, "Dynamic PRA Methods to Evaluate the Impact on Accident Progression of Accident Tolerant Fuels," *Nuclear Technology*, vol. 207, no. 3, pp. 389–405, Mar. 2021, doi: 10.1080/00295450.2020.1794234.

[3] D. R. Karanki, S. Rahman, V. N. Dang, and O. Zerkak, "Epistemic and aleatory uncertainties in integrated deterministic and probabilistic safety assessment: Tradeoff between accuracy and accident simulations," *Reliability Engineering & System Safety*, vol. 162, pp. 91–102, Jun. 2017, doi: 10.1016/j.ress.2017.01.015.

[4] D. Mandelli *et al.*, "Linking classical PRA models to a dynamic PRA," *Annals of Nuclear Energy*, vol. 149, p. 107746, Dec. 2020, doi: 10.1016/j.anucene.2020.107746.

[5] I. Ibrahim *et al.*, "Safety Assessment of a Passive Decay Heat Removal System Using a Model-Based Systems Engineering Approach," *Nuclear Science and Engineering*, pp. 1–22, Mar. 2026, doi: 10.1080/00295639.2026.2631317.

[6] S. C. Pair and M. C. Jones, "Model-Based Systems Engineering and Simulation of a Molten Salt Reactor Power Plant for Requirements Analysis".

[7] E. Roumili, J. Bossu, V. Chapurlat, N. Daclin, J. Tixier, and R. Plana, "Contribution to Nuclear Safety Demonstration Through System Modelling and Artificial Intelligence," *INSIGHT*, vol. 24, no. 4, pp. 31–33, Dec. 2021, doi: 10.1002/inst.12360.

[8] M. Tang, J. Yang, P. Zhao, and K. Wang, "Research on design requirements for passive residual heat removal system of lead cooled fast reactor via model-based system engineering," *Nuclear Engineering and Technology*, vol. 56, no. 8, pp. 3286–3297, Aug. 2024, doi: 10.1016/j.net.2024.03.029.

[9] Y. Jeong, D. Lisowski, Q. Lv, and R. Hu, "A Conceptual Design of the Reactor Cavity Cooling System for the Horizontal Compact High Temperature Gas Reactor (HC-HTGR)," ANL/NSE--23/89, 2427308, 186605, Sep. 2023. doi: 10.2172/2427308.

[10] J. P. Abraham, E. M. Sparrow, and W. J. Minkowycz, "Internal-flow Nusselt numbers for the low-Reynolds-number end of the laminar-to-turbulent transition regime," *International Journal of Heat and Mass Transfer*, vol. 54, no. 1–3, pp. 584–588, Jan. 2011, doi: 10.1016/j.ijheatmasstransfer.2010.09.012.

[11] F. P. Incropera and F. P. Incropera, Eds., *Fundamentals of heat and mass transfer*, 6th ed. Hoboken, NJ: John Wiley, 2007.