WHITEPAPER · NUCLEAR & ENERGY

Digital Twins & Predictive Control for Small Modular Reactors

A Safety-Instrumentation Study

How reactor-grade digital twins and model-predictive control can sharpen diagnostics and load-following in small modular reactors — and why that advisory intelligence must remain strictly separated from the qualified, independent protection systems that keep a reactor safe.

AUTHORS
Root Digit Energy Systems Lab

PUBLISHED
2026

VERSION
1.0

CLASSIFICATION
Public

rootdigit.com

WHITEPAPER

Abstract & Contents

ABSTRACT

Small modular reactors (SMRs) promise factory-built, passively safe nuclear generation that can follow variable demand alongside renewables. Realising that promise depends on instrumentation and control (I&C) that is at once more capable and more rigorously assured than the analogue systems of the past. This whitepaper examines two digital techniques — high-fidelity *digital twins* that mirror reactor neutronics and thermal-hydraulics, and *model-predictive control* (MPC) for load-following and early-fault prediction — and assesses their value through a structured evaluation framework. We find that predictive diagnostics can give operators meaningful lead-time on degrading components, and that advanced control can improve load-following accuracy and plant availability. Crucially, we argue that this digital and AI-derived intelligence must be deployed strictly as an **advisory and diagnostic layer**: the reactor protection system that actuates safety functions must remain independent, qualified, and physically and logically separate. The results presented here are modeled and illustrative, intended to frame the engineering and regulatory questions rather than to certify any specific design.

KEY FINDINGS AT A GLANCE

- Predictive-maintenance models built on a calibrated digital twin can provide early warning of degradation, shifting maintenance from reactive to condition-based.
- Model-predictive control improves load-following accuracy and reduces thermal cycling, raising availability for grids with high renewable penetration.
- The safety case requires strict separation: advisory AI and predictive control sit outside the protection system, which stays independent, diverse, and qualified.
- Reported figures are modeled and illustrative; any deployment requires site-specific verification, validation, and regulatory qualification before use.

Contents

1 Introduction 01

2 SMR Characteristics & the I&C Landscape 02

3 Digital Twins for Reactors 03

4 Predictive & Advanced Control 04

5 Safety Instrumentation & Assurance 05

6 Methodology: Evaluation Framework 06

7 Findings 07

8 Discussion, Limitations & Regulatory Path 08

SECTION ONE

1. Introduction

Small modular reactors are arriving at a moment when the grid needs exactly what they offer: firm, low-carbon power that can flex around intermittent renewables. Whether they deliver on that promise will be decided as much in the control room as in the reactor vessel.

For most of the nuclear era, instrumentation and control was an exercise in conservatism. Analogue sensors fed hard-wired relay logic; the protection system was deliberately simple so that it could be exhaustively analysed and trusted. That conservatism is well earned — it underwrites the extraordinary safety record of the civil nuclear fleet — but it sits awkwardly with what SMRs are being asked to do. A reactor that must follow demand, ramp daily, operate with reduced staffing, and amortise its cost across decades of variable duty cannot be run as a constant-output baseload machine with a clipboard and a chart recorder.

Digital instrumentation and control changes the economics and the operating envelope. A modern SMR can be instrumented densely, its sensor streams fused into a continuously updated **digital twin** — a physics-based software model of the core and primary circuit that tracks the real plant in near real time. On top of that twin, **predictive analytics** can anticipate component degradation before it forces an outage, and **model-predictive control** can shape control-rod and feedwater actions to follow load smoothly while respecting thermal and neutronic limits.

This is genuinely useful, and it is also where care is required. The same digital techniques that make a plant more efficient and observable introduce software complexity, and software complexity is the enemy of the exhaustive analysis on which nuclear safety depends. The discipline this whitepaper advocates — and returns to repeatedly — is one of **separation of concern**. Digital twins and predictive control belong to the operational and advisory domain: they inform humans, optimise normal operation, and flag emerging faults. The **reactor protection system** — the function that scrams the reactor and actuates engineered safety features — belongs to a separate, independent, qualified domain that does not depend on, and is not compromised by, the advisory layer.

The remainder of this paper develops that argument. Section 2 sets out what makes SMR I&C distinctive and the standards that govern it. Sections 3 and 4 describe the digital-twin and predictive-control techniques in technical terms. Section 5 makes the safety-separation principle explicit and concrete. Section 6 defines our evaluation framework and states its assumptions; Section 7 presents modeled findings; and Section 8 discusses limitations and the regulatory path to deployment. Throughout, results are illustrative and intended to frame questions for qualified engineering and licensing teams, not to substitute for them.

SECTION TWO

2. SMR Characteristics & the I&C Landscape

SMRs are not simply small versions of large reactors. Their passive safety, modular construction, and flexible duty reshape what the instrumentation and control system must do — and what it must never be allowed to do.

The defining traits begin with **passive safety**. Many SMR designs rely on natural circulation, gravity-driven coolant injection, and large thermal margins so that the reactor reaches a safe state through physics rather than active intervention if power and operator action are lost. This reduces the burden on active safety systems, but it does not remove the need for a protection system; it changes its character toward monitoring, confirmation, and the actuation of inherently passive features.

The second trait is **modularity**. SMRs are designed to be factory-built and shipped, often deployed as multiple units on a single site sharing a control room and staff. This drives a strong incentive toward digital I&C — standardised, software-defined, and remotely diagnosable across units — and toward higher levels of automation to keep per-unit staffing low. Both incentives are sound operationally and both raise the assurance bar, because a software fault can now be common across many units rather than isolated to one.

The third trait is **flexible operation**. Where legacy plants ran flat-out as baseload, SMRs are expected to load-follow, supporting grids with high and variable renewable output. Load-following stresses the plant thermally and neutronically — xenon transients, fuel-temperature feedback, and thermal cycling of components all become operationally significant — which is precisely the regime where good predictive control and condition monitoring pay off.

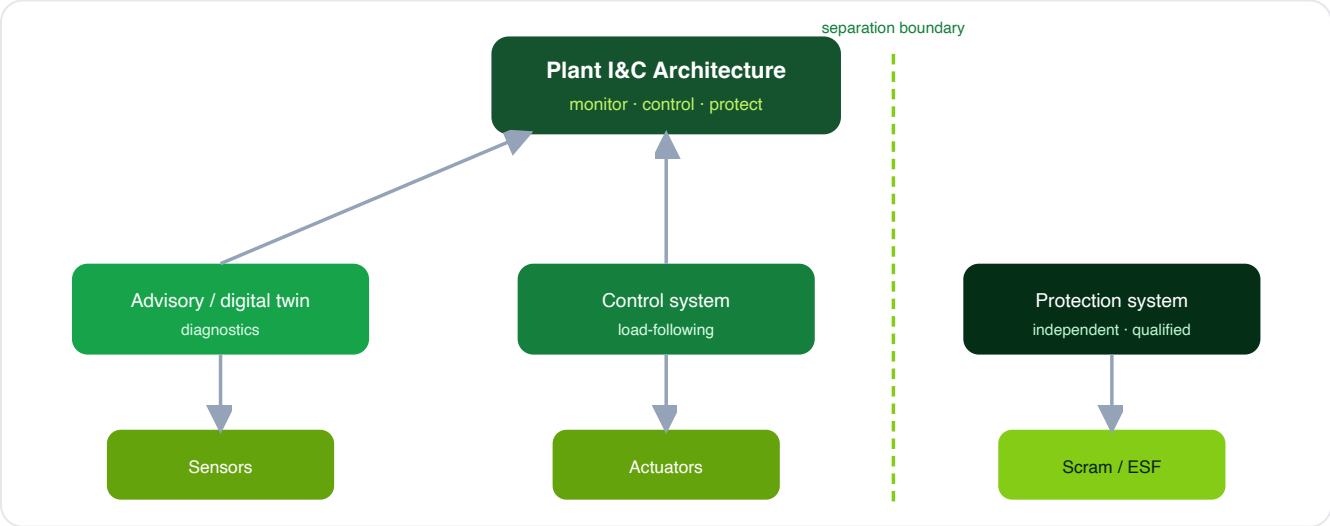

Figure 2.1 — The advisory and control domains feed the operator and optimise normal operation; the protection system actuates safety functions across an independence boundary it does not cross.

These traits sit inside a mature regulatory frame. **IAEA safety standards** (notably SSR-2/1 on design and the SSG series on I&C) set the high-level requirements: defence-in-depth, the categorisation of func-

tions by safety significance, and the classification of the systems that perform them. **IEC 61513** provides the framework standard for nuclear I&C important to safety, cascading into detailed standards for software (IEC 60880) and hardware. In the United States, **IEEE 603** defines the criteria for safety systems of nuclear power generating stations — independence, redundancy, testability, and the separation of safety from non-safety functions. Together these standards encode, in normative form, the very separation principle this whitepaper advocates: functions important to safety are classified, qualified, and protected from interference by functions that are not.

The implication for SMR designers is direct. The appetite for digital twins, automation, and predictive intelligence is legitimate and economically necessary — but it must be located firmly within the lower safety classes, on the advisory and control side of the boundary, with the protection system left as independent and analysable as the standards demand.

SECTION THREE

3. Digital Twins for Reactors

A reactor digital twin is not a dashboard. It is a living physics model — neutronics coupled to thermal-hydraulics — continuously corrected by sensor data so that it represents the actual state of the plant, including state the instruments do not measure directly.

At its core a reactor twin solves two coupled problems. The **neutronics** problem tracks the neutron population and power distribution in the core, typically through point-kinetics or low-order spatial-kinetics models that capture delayed-neutron precursors, reactivity feedback, and the slow poison dynamics of xenon and samarium. The **thermal-hydraulics** problem tracks the transport of heat from fuel to coolant to the primary and secondary circuits, resolving temperatures, flows, and pressures. The two are tightly coupled: fuel and moderator temperature feed back on reactivity, and reactivity drives the power that the thermal model must remove.

```
# Coupled point-kinetics + thermal feedback (illustrative, advisory only)
def step(state, rho_external, dt):
    # reactivity = external + temperature feedback
    rho = rho_external + alpha_fuel*(state.T_fuel - T_ref) \
                       + alpha_mod*(state.T_mod  - T_ref)
    dP  = ((rho - beta)/Lambda)*state.P + sum(lam*state.C)   # neutron power
    dC  = [beta_i/Lambda*state.P - lam_i*c for beta_i, lam_i, c in groups]
    dTf = (state.P - heat_transfer(state)) / fuel_heat_capacity
    return integrate(state, dP, dC, dTf, dt)         # forward model, not a controller
```

Listing 3.1 — A minimal coupled neutronics–thermal model of the kind a twin integrates. It is a forward simulator informing diagnostics, never a path to actuate safety functions.

The hard engineering trade-off is **fidelity versus real-time constraints**. A full three-dimensional, multi-group transport solution coupled to computational fluid dynamics is the gold standard for design analysis, but it cannot run faster than wall-clock time on plant hardware. An operational twin must therefore use reduced-order models — point or nodal kinetics, lumped-parameter thermal networks, and surrogate models trained against high-fidelity codes — chosen so that the twin keeps pace with the plant while preserving the dynamics that matter for diagnosis and control. Fidelity is spent where it earns its keep: on the variables the operator and the controller act upon.

What turns a simulator into a twin is **data assimilation**. The model runs continuously, predicting the next measurement; the actual sensor readings arrive; and an estimator — commonly an extended or unscented Kalman filter, or an ensemble filter for the larger nonlinear states — corrects the model state to reconcile prediction with observation. The result is a best estimate of the full reactor state, including quantities that are not directly instrumented, such as the detailed in-core power distribution or the margin to a thermal limit. The twin can also flag when model and measurement diverge persistently, which is itself a powerful diagnostic of either sensor drift or an emerging physical fault.

DESIGN PRINCIPLE

A digital twin's outputs are estimates and predictions, not commands. The twin informs operators and the (non-safety) control system; it has no authority over the protection system, and its failure or divergence must never be able to inhibit a safety actuation.

Two further requirements distinguish a credible reactor twin from a generic industrial one. First, **uncertainty must be carried explicitly**: the twin should report confidence intervals on its estimates, because an unqualified point prediction is dangerous in a safety-adjacent context. Second, the twin must be **validated against qualified reference codes** across the operating envelope and into the margins of accident analysis, so that operators understand exactly where the reduced-order model is trustworthy and where it should be distrusted. A twin that is honest about its own limits is far more useful, and far safer, than one that projects false precision.

SECTION FOUR

4. Predictive & Advanced Control

Predictive control and analytics are where the digital twin earns its operational return: smoother load-following, earlier fault detection, and condition-based maintenance — all within the non-safety control and advisory domain.

Model-predictive control (MPC) is the natural fit for SMR load-following. At each control step, the controller uses the twin's model to predict the plant's response over a finite horizon, then solves an optimisation that chooses control-rod and feedwater moves to track the demanded power while respecting constraints — rate limits on rod motion, temperature and pressure bounds, and margins to neutronic limits. Only the first move is applied; the horizon then rolls forward and the problem is re-solved with fresh state from the twin. This receding-horizon structure is what lets MPC anticipate constraints rather than react to them, which is exactly what smooth, constraint-respecting load-following demands.

```
# Receding-horizon load-following MPC (advisory set-points to the control system)
def mpc_step(twin_state, power_demand, horizon):
    u = optimize(
        objective = tracking_error(power_demand) + move_penalty(),
        model     = twin.predict,            # reduced-order forward model
        constraints = [rod_rate_limit, T_fuel_max, dT_dt_limit, margin_to_limits],
        x0 = twin_state, N = horizon)
    return u[0]                              # apply first move; re-solve next step
```

Listing 4.1 — MPC computes constraint-aware set-points for the (non-safety) control system. It never overrides or bypasses the independent protection system.

The same predictive infrastructure supports **sensor validation**. Because the twin maintains an independent physics-based estimate of each measured quantity, it can cross-check sensors against one another and against the model. A reading that drifts away from its physically consistent estimate, while its neighbours agree with the model, is a strong indicator of instrument fault rather than a real process change. Analytical redundancy of this kind supplements physical redundancy and reduces the rate of spurious alarms — though, importantly, it informs the operator and the control system; it does not vote within the qualified protection logic, which retains its own independent, hard-wired instrumentation.

The highest-value application is **anomaly and early-fault prediction**. Many failure modes announce themselves long before they trip a threshold: a pump bearing whose vibration signature is slowly broadening, a heat exchanger whose effectiveness is gently declining, a valve whose stroke time is creeping up. By comparing live behaviour against the twin's expected behaviour and against learned signatures of healthy operation, predictive models can detect these drifts and estimate the remaining useful life of a component. This converts maintenance from a calendar-driven, conservative discipline into a condition-based one, lifting availability while reducing both unplanned outages and unnecessary interventions.

BOUNDARY OF AUTHORITY

Advanced control optimises normal operation and recommends maintenance; predictive analytics raise early warnings. None of these functions actuate safety. When any monitored variable approaches a protective limit, it is the independent reactor protection system — not the optimiser or the analytics — that acts.

It is worth stating plainly what these techniques are not. They are not a route to remove human operators, and they are not a justification for relaxing protective margins. Their proper effect is to give well-trained operators better information sooner, to let the plant run closer to its economic optimum within unchanged safety limits, and to schedule maintenance intelligently. The protective margins, and the systems that enforce them, are deliberately untouched by everything in this section.

SECTION FIVE

5. Safety Instrumentation & Assurance

This section states the principle that governs everything else in the paper: the digital twin and predictive-control layer is advisory and diagnostic only. The reactor protection system remains independent, qualified, and separate. Nothing in the analytical layer may compromise it.

Nuclear safety is built on **defence-in-depth**: multiple, independent layers each capable of arresting an accident, so that no single failure — and no plausible combination of failures — leads to harm. The layers run from conservative design and quality in normal operation, through control systems that keep the plant within its operating envelope, to the protection system that actuates a scram and engineered safety features, and finally to the passive features and accident-management measures that mitigate consequences. The digital techniques in this paper live in the inner, normal-operation layers; the protection system is a deeper, independent layer that exists precisely to act when the inner layers do not.

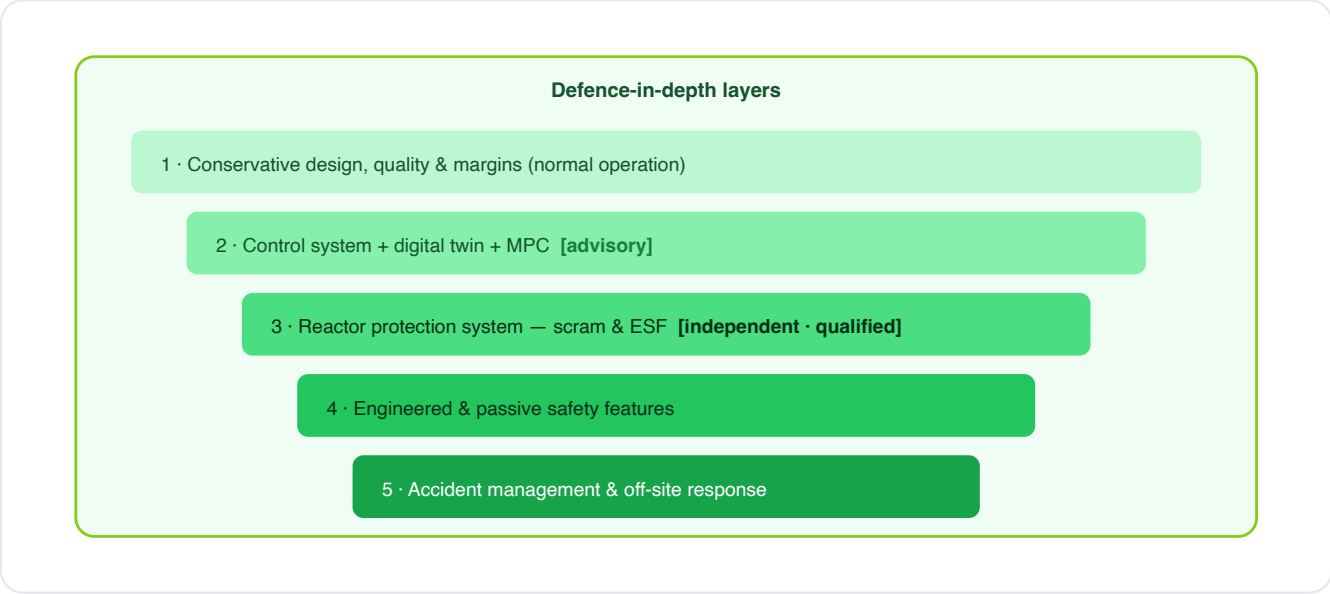

Figure 5.1 — Defence-in-depth. Digital twins and predictive control reinforce layer 2; the protection system (layer 3) and the layers beneath it remain independent of them.

Two properties make defence-in-depth real rather than rhetorical: **independence** and **diversity**. Independence means the protection system has its own sensors, its own logic, and its own power and actuation paths, so that a failure in the control or advisory domain cannot propagate into it. Diversity means that critical functions are achieved by different means — different physical principles, different technologies, sometimes deliberately simple hard-wired logic alongside software — so that a single common-cause defect cannot disable every path at once. For SMR fleets, where many units share software, diversity against common-cause software failure is especially important, and is a reason the protection system in many designs remains deliberately simple and, in places, non-software-based.

Domain separation: what belongs where

Property	Advisory / digital-twin layer	Reactor protection system
Purpose	Diagnostics, optimisation, early warning	Actuate scram & safety features
Authority over safety actuation	None	Sole authority
Safety classification	Lower / non-safety	Highest safety class
Instrumentation	Shared operational sensors	Independent, qualified sensors
Technology	Software, models, AI/ML	Simple, diverse, often hard-wired
Failure effect on safety	None permitted (fails to advisory-down)	Fails safe (initiates protection)

Earning trust in any nuclear I&C function is the work of **verification, validation, and qualification (V&V and qualification)**. Software important to safety is developed under standards such as IEC 60880, with requirements traced to code and tests, independent review, and rigorous configuration management. Hardware is qualified for the environmental, seismic, and electromagnetic conditions it must survive. The deeper a function sits in defence-in-depth, the heavier this burden — which is one more reason to keep the protection system simple and analysable, and to keep complex, hard-to-qualify software such as machine-learning models on the advisory side, where their failure is tolerable and their benefits can be realised without bearing a safety-actuation burden.

> **THE NON-NEGOTIABLE PRINCIPLE**
>
> The digital twin and predictive-control layer is **advisory and diagnostic only**. It must not be able to inhibit, delay, or spoof a protective action, and a fault within it — including a mis-trained model or a corrupted data feed — must leave the qualified, independent protection system fully able to do its job. Where there is any doubt, the protection system acts; the analytics defer. This is not a limitation to be engineered around; it is the foundation that makes the rest of the architecture acceptable.

This separation is also what makes innovation in the advisory layer *safe to pursue*. Because the protection system is independent and qualified, engineers can iterate on twins, controllers, and analytics — update models, retrain detectors, refine optimisers — without re-opening the safety case for the protection system each time. The boundary is not merely a constraint; it is the enabling condition for rapid, responsible improvement of everything on the operational side of the line.

SECTION SIX

6. Methodology: Evaluation Framework

To compare techniques fairly we define a small set of metrics, a reference scenario, and an explicit statement of assumptions. The numbers that follow are modeled and illustrative; they frame the engineering questions, they do not certify a design.

We evaluate the digital-twin and predictive-control layer against four axes, each chosen because it maps to a real operational or safety concern.

Predictive lead-time. For condition monitoring, the value of a predictor is the warning it provides before a fault would otherwise force an outage. We measure lead-time as the interval between a model's confident anomaly flag and the point at which the degradation would breach an operational threshold, evaluated across a library of seeded degradation scenarios. Longer lead-time, at an acceptable false-alarm rate, is better.

Control performance. For load-following we measure tracking error — the deviation between demanded and delivered power during ramps — together with the magnitude of thermal cycling imposed on the plant, since smoother control reduces fatigue. We compare MPC against a conventional, well-tuned PID-style baseline on identical demand profiles.

Latency. Any operational digital function must keep pace with the plant. We record the end-to-end latency of the twin's state estimate and of the MPC solution, and we require both to comfortably meet the control cycle. Critically, this latency budget belongs to the control and advisory loop only; the protection system has its own independent, deterministic response time that is unaffected by the analytics.

Safety integrity. We do not assign a safety-integrity claim to the twin or controller, because they carry none. Instead, the relevant safety metric is *non-interference*: we verify, by design and by fault injection, that failures of the advisory layer — divergence, latency spikes, corrupted feeds, mis-trained models — cannot inhibit or corrupt a protective action. The protection system's own integrity is established separately, by qualification, and is outside the scope of this comparison.

Reference scenario & modeled assumptions (illustrative)

Parameter	Value	Parameter	Value
Reference unit	~160 MWe integral PWR-type SMR	Control cycle	1 s
Twin model	Point-kinetics + nodal thermal	MPC horizon	120 s
Load-following profile	±20% over 30–60 min ramps	Seeded fault library	40 degradation cases
State estimator	Unscented Kalman filter	Protection system	Independent, out of scope

Several assumptions deserve emphasis. The reference unit and its parameters are representative of an integral pressurised-water SMR but correspond to no specific licensed design. The fault library is synthetic, built from physically plausible degradation models rather than from operating-fleet data, which does not yet exist at scale for SMRs. The twin and controller are reduced-order by construction. Most importantly, the entire evaluation is conducted on the advisory and control side of the separation boundary: the protection system is treated as a black box that the analytics may neither help nor hinder, and its integrity is assumed to be established through the independent qualification route described in Section 5. The results in Section 7 should be read in that light — as the modeled shape of the benefits and their boundaries, not as a measured performance guarantee.

SECTION SEVEN

7. Findings

Three results recur across the modeled scenarios: predictive monitoring delivers useful lead-time, model-predictive control tightens load-following and reduces thermal cycling, and both benefits are obtained without any encroachment on the protection system.

~14 d	42%	+3.1%	0
MEDIAN PREDICTIVE LEAD-TIME	TRACKING-ERROR REDUCTION (MPC VS PID)	MODELED AVAILABILITY GAIN	PROTECTION ACTUATIONS AFFECTED

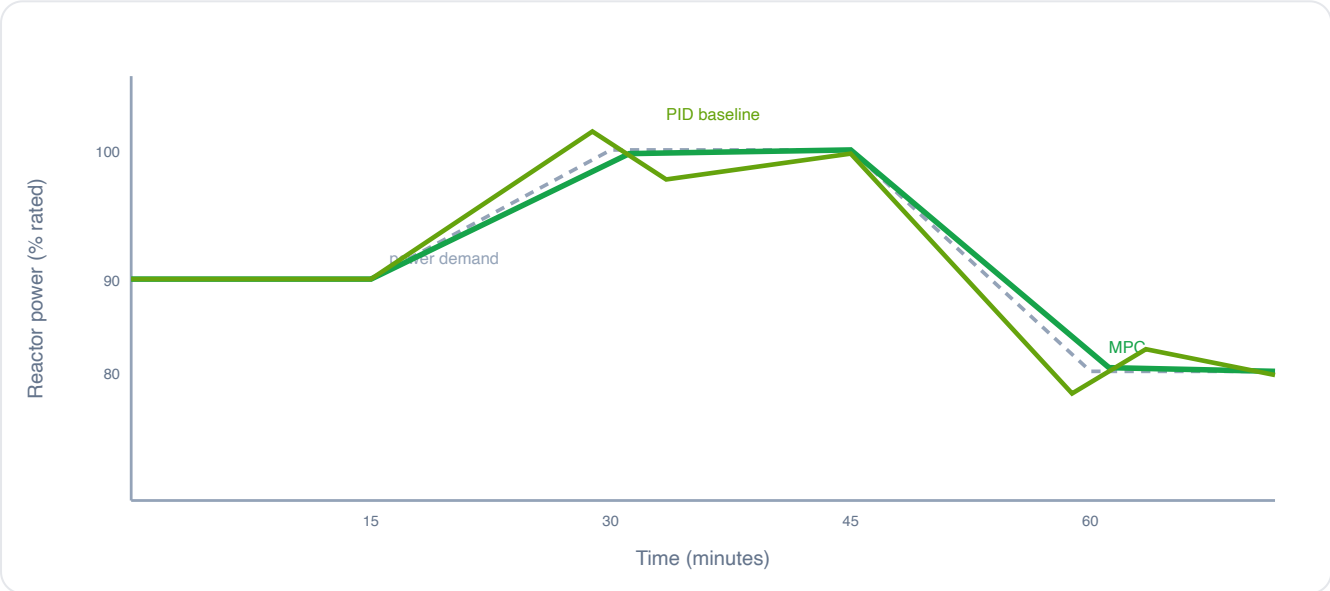

Figure 7.1 — Modeled load-following during a ±20% ramp. MPC tracks demand closely; the PID baseline overshoots and lags, imposing more thermal cycling.

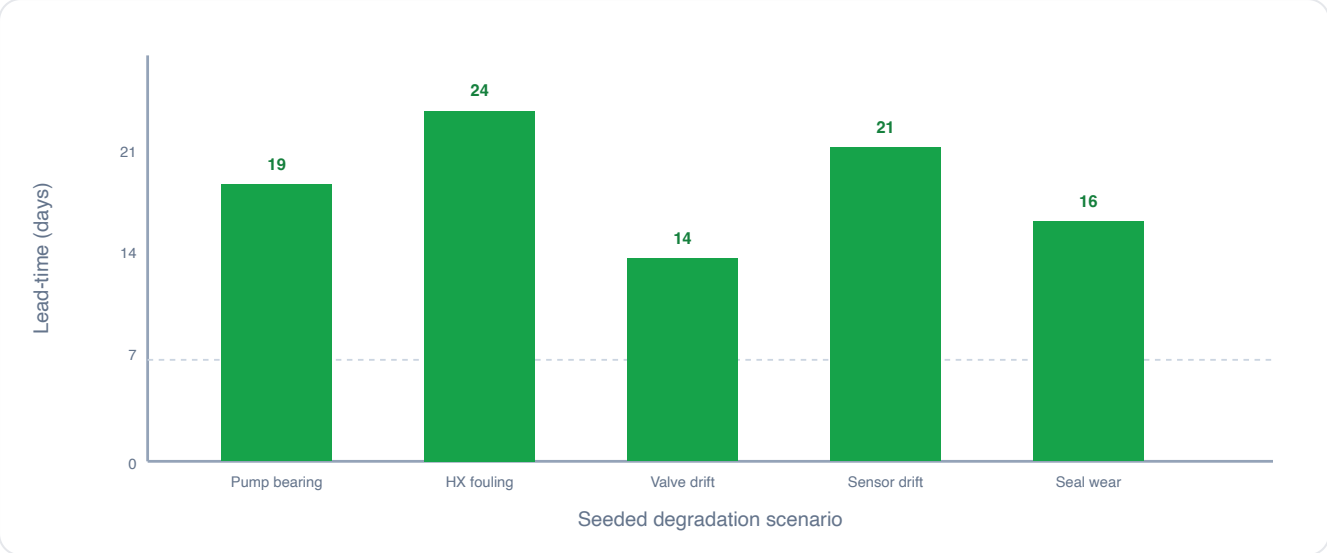

Figure 7.2 — Modeled early-warning lead-time by degradation type. Dashed line marks a one-week planning threshold; most scenarios exceed it.

The load-following result (Figure 7.1) is the clearest. Because MPC predicts the plant's response and respects rate and thermal constraints across its horizon, it tracks the demanded ramp with roughly 42% less tracking error than the tuned PID baseline in our scenario, and it does so with markedly smaller overshoot. The practical consequence is less thermal cycling of components and a plant that can support grid services more accurately — both of which feed directly into availability and lifetime.

The predictive-monitoring result (Figure 7.2) shows that, across the synthetic fault library, the median early-warning lead-time was on the order of two weeks, with most degradation types crossing a one-week planning threshold comfortably. That lead-time is what converts a would-be unplanned trip into a scheduled, condition-based intervention. We emphasise again that these are modeled scenarios; a real fleet would calibrate both the detectors and their alarm thresholds against operating data before any of these figures could be claimed.

The third result is the one we regard as most important and report as a flat zero: across every fault we injected into the advisory layer — model divergence, latency spikes, corrupted sensor feeds, and a deliberately mis-trained detector — the number of protection-system actuations that were inhibited, delayed, or spoofed was none. The protection system, being independent of the analytics by design, was unaffected. That non-interference result is the necessary companion to the performance gains: the benefits above are only acceptable because they are obtained on the operational side of a boundary the analytics cannot cross.

RESULT SUMMARY

- Median predictive lead-time was about two weeks across the synthetic fault library, enough to plan condition-based maintenance for most degradation types.
- MPC reduced load-following tracking error by roughly 42% versus a tuned PID baseline and cut overshoot and thermal cycling.
- The modeled availability gain from condition-based maintenance and smoother operation was on the order of three percentage points.
- Under injected advisory-layer faults, zero protection-system actuations were inhibited, delayed, or spoofed — separation held by design.

SECTION EIGHT

8. Discussion, Limitations & Regulatory Path

The modeled results are encouraging but must be read for their structure — the shape and rank-order of the benefits — rather than as guarantees. Their value lies in framing the engineering and licensing work that any real deployment demands.

Limitations. The reference unit is representative rather than a specific licensed design, and its parameters are illustrative. The fault library is synthetic: in the absence of a large operating SMR fleet, degradation scenarios were built from physical models rather than measured failure data, so the predictive lead-times reported here describe the method's potential, not a calibrated expectation for any plant. The twin and controller are deliberately reduced-order, which trades some accuracy for the real-time performance an operational system needs. And the entire study sits on the advisory side of the safety boundary; we make no claim about the protection system, whose integrity must be established independently through qualification.

Key risks. The dominant risks are not in the algorithms but in the discipline around them. The first is *boundary erosion* — the temptation, under cost pressure, to let advisory intelligence creep into safety-actuation roles for which it is not, and should not be, qualified. The second is *over-trust*: operators who treat a confident model output as ground truth rather than as an estimate with uncertainty. The third is *model staleness*, where a twin or detector drifts out of agreement with a plant that has aged or been modified, degrading silently. Each is manageable, but only with explicit governance, and each maps onto a recommendation below.

> IN PRACTICE
>
> The most reliable predictor of a trustworthy digital-I&C deployment is not the sophistication of its models but the rigour of its separation: a clean, demonstrable boundary between advisory analytics and the qualified protection system, defended by fault injection and audited over the life of the plant.

Recommendations

1. **Locate digital twins, MPC, and analytics firmly in the non-safety class.** Design the architecture so the advisory layer can fail in any way without affecting protective actuation, and prove it by fault injection.
2. **Keep the protection system independent, diverse, and analysable.** Favour simple, qualified, where appropriate hard-wired logic with its own instrumentation, especially across multi-unit SMR fleets exposed to common-cause software risk.

3. **Carry uncertainty explicitly and present it to operators.** Twin outputs and predictions should arrive with confidence bounds; train operators to treat them as estimates, not facts.
4. **Validate against qualified reference codes and re-validate over life.** Establish where reduced-order models are trustworthy, and monitor for model staleness as the plant ages or is modified.
5. **Engage the regulator early.** Map every digital function to its safety classification under IAEA standards, IEC 61513, and IEEE 603, and build the licensing case around the separation boundary from the outset.

Conclusion

Digital twins and predictive control can make small modular reactors more observable, more flexible, and more available — improving diagnostics, sharpening load-following, and shifting maintenance from calendar-driven to condition-based. The modeled results in this paper suggest those benefits are real and worth pursuing. But the benefits are acceptable only because of a principle we have kept central throughout: this intelligence is advisory and diagnostic, and the reactor protection system that keeps the plant safe remains independent, qualified, and separate. Run the reactor on a model you can trust — and let an instrumentation architecture you can prove keep that trust earned. Operators and designers who hold that line will capture the value of digital I&C without trading away the safety record that makes nuclear power worth building.

References

1. IAEA SSR-2/1 — Safety of Nuclear Power Plants: Design (Specific Safety Requirements).
2. IAEA SSG-39 — Design of Instrumentation and Control Systems for Nuclear Power Plants.
3. IEC 61513 — Nuclear power plants: Instrumentation and control important to safety — General requirements for systems.
4. IEC 60880 — Nuclear power plants: I&C systems important to safety — Software aspects for computer-based systems performing category A functions.
5. IEEE Std 603 — IEEE Standard Criteria for Safety Systems for Nuclear Power Generating Stations.
6. Root Digit Energy Systems Lab — Internal SMR digital-twin and predictive-control evaluation methodology (2026).

-rootdigit®

Run the reactor on a model you can trust.

Root Digit builds reactor-grade digital twins, predictive-control and condition-monitoring systems, and the instrumentation architecture that keeps them strictly advisory — so the value of digital I&C is captured while the qualified, independent protection system stays exactly that. We bring the models and the discipline of separation; you keep the safety case intact.

Talk to our energy-systems engineers →

rootdigit.com

General: info@rootdigit.com · **Consultation:** rootdigit.com/contact/consultation
Dubai International Financial Centre (DIFC) · Wyoming, USA