

## Non-Intrusive Fleet-Level AGC Testbed Built on a Full-Scope Multi-Unit SMR Training Simulator

Jae Min Kim<sup>a\*</sup> and Seo Ryong Koo<sup>a</sup>

<sup>a</sup>Korea Atomic Energy Research Institute, Advanced I&C Research Division, 111 Daedeok-daero 989beon-gil, Yuseong-gu, Daejeon 34057, Republic of Korea

\*Corresponding author: jaemink@kaeri.re.kr

**\*Keywords :** full-scope training simulator, small modular reactor, multi-unit operation, flexible operation, automatic generation control, control rod actuation

### 1. Introduction

Small modular reactors (SMRs) are expected to be deployed as multi-unit sites and to take part in flexible grid operation, from daily load-follow to secondary frequency regulation. Secondary regulation is coordinated by the grid operator through automatic generation control (AGC), which distributes a demand signal to participating plants. In Korea, nuclear units do not take part in AGC at present. Among other reasons, control rod manipulation is reserved for licensed operators, so the output of a nuclear unit is not adjusted remotely by the power exchange. Multi-unit SMR sites are nevertheless expected to share in demand control with the power exchange in the future, and the work reported here is a preparatory research setup for that case. Studies of such operation for SMRs have so far relied on plant models: co-simulation platforms couple plant models to power-system simulators [1], control-rod-driven load-follow is studied with system codes [2], and dispatch analyses of multi-unit SMR sites are carried out at the optimization level [3]. Full-scope training simulators offer a complementary option. They are integrated plant models maintained for operator training, and one is available for a multi-unit SMR design. Being closed commercial products, however, they expose neither model equations nor internal control hooks, which has limited their use in control studies.

This paper describes a testbed built on such a simulator in order to test fleet-level AGC logic, of the kind the power exchange already applies to conventional plants, against a twelve-unit SMR plant of about 50 MWe per unit. The simulator is used as delivered and without access to its model; it is actuated only at the level of operator panel actions. Rod switches are thus operated by software in place of an operator, which is a research expedient and not a proposal for plant practice. The paper reports (i) a measurement campaign that identifies which power-control path is usable from an external client, (ii) an external dispatch layer that closes the loop over all twelve units, and (iii) the initial experiment run on the testbed, which shows how the rod-actuation deadband constrains the allocation of the AGC correction among units. The scope is limited to the AGC timescale of minutes. The external grid is represented only by a minimal boundary

model that generates demand signals, and no grid-integration claim is made.

### 2. Non-Intrusive Testbed

#### 2.1 Control Path Identified by Exclusion

The simulator exposes an operator-client software development kit (SDK), a Windows Communication Foundation service shared with the plant human-machine interface. All candidate write paths for steering electrical output were enumerated and tested against the running simulator. The instructor-station controlled-parameter family (turbine load demand, reactor power demand, manual valve commands) accepts writes, but the plant does not respond. Those variables are consumed only by the instructor-station protocol, and the plant control system overwrites them every cycle. Turbine mode switches and runback commands are likewise inert from the operator client. The only path that responded among those tested is the operator rod-control panel, which requires reproducing the operator interlock: rod-select and bank-select switch pulses, followed by a deadman-style hold of the IN/OUT switch.

The channel was then characterized live. A rod step changes thermal power by about 0.6 to 2.4% of rated depending on rod position, as measured over repeated single-step trials, and the change is reflected within about 30 s. Electrical output follows with a lag of several minutes because the turbine control holds the steam header pressure. The dispatch layer therefore uses thermal power as its feedback variable. The channel is quantized in both amplitude and time. Rod position moves only in integer steps, and switch holds shorter than about 4 s produce no motion at all. One step, up to about 1.2 MWe, is therefore treated as the precision floor of any tracking figure reported below. All maneuvering in this work uses control rods only, and the boron system is left untouched. Boron-free-style load-follow operation is assumed throughout.

#### 2.2 External Dispatch Layer

A dispatch layer written in Python runs on a 10 s tick and is controlled over HTTP. Per-unit actuation is a pulse-width command. The IN or OUT switch is held

for a fraction of each tick proportional to the thermal-power error, with a 0.4% deadband (about 0.2 MWe per unit) and saturation at half a tick. Software interlocks in the dispatch layer forbid insertion below 20% thermal power and withdrawal above 102.5%, and they force switch release on stop or abort. A constant per-unit offset between electrical and indicated thermal power is measured at dispatch start and compensated. The reference fleet controller is an AGC-style scheme. It combines a proportional, droop-like share of the boundary signal with a deliberately slow error integral that is frozen while all units are saturated, which proved necessary against the minutes-scale turbine lag. No primary frequency-regulation capability is implied. A single-bus swing equation converts the generation-demand imbalance into a frequency-like signal. It serves only as a demand-signal generator for experiments, not as a grid model. SDK reads are served from a client-side package cache refreshed every 200 ms, which is negligible against the 10 s tick. Every simulator variable involved is logged at 1 Hz for all runs. Finally, the computation of per-unit targets is pluggable. An external process can take it over through the same HTTP interface, in which case the dispatch layer performs actuation only.

### 3. Initial Results

Results are reported as preliminary. Except for the repeated run noted in Sec. 3.3, each experiment is a single run. Convergence is declared when the fleet tracking error stays within  $\pm 1.5$  MW of demand for 3 min.

#### 3.1 Single-Unit Commissioning

The first closed-loop run on one unit exhibited classic integrator windup. Because the integral accumulated faster than the turbine responded, the unit power target was driven far below demand, and the output oscillated with a period of about 5 min and almost no damping. After slowing the integral tenfold, freezing it under saturation, and clamping it, a 1.8 MW demand step converged in about 7 min and held within  $\pm 0.3$  MW for the following 11 min.

#### 3.2 Twelve Units Simultaneously

Table I summarizes a 56 min staged scenario over all twelve units: steady-state hold, a  $-6$  MW fleet step, a  $-18$  MW ramp, loss of one unit with redispatch over the remaining eleven, and recovery. All phases converged, and the 10 s tick accommodated the switch traffic of all twelve units (up to 24 select pulses per tick) without overruns. During phase B the participation factor of one unit was raised twentyfold at runtime, which concentrated about two-thirds of the power reduction on that unit. When the factor was later restored, the fleet re-balanced automatically. Residual errors of up to 1.42 MW remain in phase B. Their magnitude is plausibly attributable to the per-unit deadband discussed next.

Table I: Twelve-unit staged scenario (fleet demand 600 MW)

| Phase | Disturbance                   | Time to band*      | Residual (MW) | Max $ \Delta f $ (Hz)† |
|-------|-------------------------------|--------------------|---------------|------------------------|
| A     | steady-state hold             | immediate          | +0.00         | 0.009                  |
| B     | $-6$ MW step                  | 8.8 min            | -1.42         | 0.060                  |
| C     | $-18$ MW ramp (10 min)        | 1.5 min after ramp | -1.25         | 0.056                  |
| D     | unit trip, 11-unit redispatch | 6.0 min            | -0.46         | 0.018                  |
| E     | unit return and restore       | 7.2 min            | +0.52         | 0.246                  |

\* fleet error within  $\pm 1.5$  MW of demand held for 3 min. † auxiliary indicator of the boundary model only.

#### 3.3 Fleet-Level Quantization Floor

The plugin interface was exercised with a deliberately naive external client that distributes demand equally over the units and applies proportional-integral trim on the fleet error. Without integral action the client stalls about 2.2 MW short of demand. The per-unit share of the residual falls below the rod deadband, and every rod stops. With a slow integral, a  $-6$  MW step produces the equal-share trajectory of Fig. 1. The error drops quickly to about 3 MW and then oscillates for some 11 min in a sawtooth between 3 and 4.7 MW, near twelve times the 0.2 MWe per-unit deadband. It reaches only about 2 MW by the end of the 15 min run and never holds the convergence band. Equal sharing parks all units inside their deadbands at once, so the integral must re-accumulate past the deadband to move the whole fleet; this produces the sawtooth and makes equal sharing the worst-case residual allocation for this actuator. To confirm the mechanism, the same step was repeated with one change: the whole corrective term was concentrated on a single unit. The error then falls through the floor and enters the convergence band 6.7 min after the step, settling at  $-0.4$  MW (Fig. 1). Three repeats from the same stored initial condition entered the band in 6.3 to 6.7 min with final errors of  $-0.4$  to  $-0.6$  MW. Two interface requirements follow for any external model driving a rod-actuated fleet of this type. Integral action is mandatory, and residual errors should be concentrated on few units so that per-unit corrections exceed the deadband.




Fig. 1. Fleet tracking error after a  $-6$  MW step: equal sharing (black) versus correction concentrated on one unit (blue).

### 3.4 Hour-Timescale Core Behavior

To probe whether core physics beyond prompt feedback is modeled, one unit was reduced by four rod steps and then left with rods stationary for 4.2 h, which was verified by rod-position and switch-lamp readback. An unperturbed unit was logged as a control. The reduced unit drifted downward by a further 1.18% of rated at a decaying rate ( $-0.46$  to  $-0.13\%$  per hour), whereas the control unit stayed within  $\pm 0.01\%$  (Fig. 2). This unit-specific, decelerating drift after a power reduction indicates that hour-timescale core dynamics consistent with xenon buildup are present in the simulator. Multi-hour load-follow scenarios therefore cannot treat the plant as time-invariant. The electrical-to-thermal indication offset was also found to be power-dependent,  $+0.5$  percentage points at full power versus  $-1.9$  at 96%. The constant compensation of Sec. 2.2 is therefore valid only near its calibration point.




Fig. 2. Thermal power over 4.2 h with rods stationary: unit reduced by four rod steps versus unperturbed control unit.

### 4. Conclusions

A closed full-scope multi-unit SMR training simulator was operated closed-loop, without modification, at the level of operator panel actions over all twelve units. The testbed accepts a fleet-level demand signal and any external allocation logic through a plugin interface, so AGC logic in use elsewhere can be tested against it. Three limitations delimit the claims. The simulator is certified for training fidelity and is not validated for engineering transients. Hour-timescale core drift consistent with xenon buildup is present (Sec. 3.4), but axial offset and burnup effects remain unexamined. Most results are single runs. The main finding of the initial experiment is practical. On a rod-actuated fleet with a per-unit deadband, an AGC correction spread equally over all units stalls at a floor of roughly the number of units times the deadband, whereas concentrating the correction on a single unit brings the fleet into the tracking band. Dispatch logic tested on this plant should therefore include integral action and assign the residual correction to few units rather than to all of them. Characterizing rod worth versus power level, longer xenon transients, and power-dependent offset compensation is left to follow-up work.

### ACKNOWLEDGEMENT

This research was supported by the National Research Council of Science & Technology (NST) grant by the Korea government (MSIT) (No. GTL24033-000).

### REFERENCES

- [1] P. Astero, P. Laakso, S. Hänninen, R. J. Millar, and M. Lehtonen, Electrical and Thermomechanical Co-Simulation Platform for NPP, *Energies*, Vol. 14, No. 4, 939, 2021.
- [2] K. Park, T. Park, S. K. Zee, and B. S. Koo, Control Rod Driven Long-term Load-follow Operation in Small Modular Reactor, *Annals of Nuclear Energy*, Vol. 182, 109620, 2023.
- [3] S. Alhadhrami, G. J. Soto, and B. Lindley, Dispatch Analysis of Flexible Power Operation with Multi-unit Small Modular Reactors, *Energy*, Vol. 280, 128107, 2023.