Online Quantification of Delayed Neutron Precursor Drift Reactivity in Molten Salt Reactor Flow and Power Transients with a Shadow-Precursor Observer

Jaejin Lee [a*], SeungYop Paek [b]
[a] *Department of Nuclear Engineering, Kyung Hee University, 1732 Deogyeong-daero, Giheung-gu, Yongin-si, Gyeonggi-do 17104*
[b] *Hyundai Engineering & Construction, 75 Yulgok-ro, Jongno-gu, Seoul 03058*
[*] *Corresponding author: jaejin.lee@khu.ac.kr*

***Keywords :** molten salt reactor, delayed neutron precursor drift, effective delayed neutron fraction, inhour equation, Modelica/TRANSFORM*

1. Introduction

In a molten salt reactor (MSR) the fuel salt circulates and carries the delayed neutron precursors (DNP) out of the core. The delayed neutrons emitted in the external loop are lost to the core chain reaction, so the effective delayed neutron fraction β_{eff} falls below the nuclear-data total β_0. Because the precursor field and the core criticality are coupled through the flow, this is a multiphysics phenomenon, and predicting the reactivity that the flow adds or removes is a prerequisite for MSR kinetics and safety analysis.

The modified point-kinetics model transports the precursors through the primary circuit and evaluates β_{eff} from the in-core precursor inventory, capturing the spatially distributed DNP loss[1]. Closed-form analytic solutions of the steady-state precursor distribution and β_{eff} are available for the one-dimensional core–loop channel[2]. On this basis, recent analyses have reproduced the steady-state β_{eff} and the reactivity change of the MSRE zero-power pump tests, in which the measured change reaches 212 pcm[3], across several system codes [4]. In these studies, however, the flow effect is quantified by the inventory reading $\beta_0-\beta_{\text{eff}}(t)$, whose negative is a valid drift reactivity only at constant power. When the power varies freely, the reading conflates the flow effect with the ordinary precursor lag that any reactor exhibits.

TRANSFORM is a Modelica-based dynamic simulation library developed at Oak Ridge National Laboratory for advanced energy systems, and its MSR application has been demonstrated with a full-plant model of the Molten Salt Demonstration Reactor [5]. The modified point kinetics model with precursor advection is implemented natively[1]. TRANSFORM also provides a built-in reactivity sum for monitoring the simulated transient. However, this sum deliberately contains no drift channel, because the flow effect enters the solution through the advected precursor field itself. The drift reactivity is therefore not separately quantified during a general transient.

A companion paper[6] derives a shadow-precursor observer that isolates the drift reactivity online for arbitrary power histories (one auxiliary differential equation per precursor group, non-intrusive and mesh-free), and verifies it in Modelica/TRANSFORM against a three-tier benchmark set whose references are all closed-form. This paper presents the two free-power tiers of that verification: a step reduction of the fuel-salt flow with the power responding freely, and an external-reactivity step at fixed flow.

2. Drift Observer and Closed-Form References

2.1 Shadow-precursor observer

Let $N(t)$ be the neutron population and $C_i^{core}(t)$ the incore inventory of DNP group i with data (β_i , λ_i , Λ), $\beta_0 = \sum_i \beta_i$. In the modified point-kinetics model of the circulating fuel[1], the delayed source of the population equation contains the in-core inventories only,

$$\frac{dN}{dt} = \frac{\rho - \beta_0}{\Lambda} N(t) + \sum_i \lambda_i C_i^{\text{core}}(t), \tag{1}$$

where ρ is the total reactivity input. The inventory diagnostic evaluated by transport codes is $\beta_{\text{eff}}(t) = \sum_i \lambda_i \Lambda C_i^{\text{core}}(t)/N(t)$. Its deficit defines the inventory reading of Section 1,

$$\rho_{\text{loss}}(t) \equiv \beta_0 - \beta_{\text{eff}}(t) \geq 0. \tag{2}$$

This sign convention is kept throughout; the reactivity-equivalent form is always written $-\rho_{\text{loss}}(t)$. The reading contains two effects of different origin: the ordinary precursor lag that accompanies any power change, and the advective (drift) loss. The explicit decomposition into the two parts is exact at all times and is given in [6].

The observer separates the two by simulating the precursor inventories C_i^s of a fictitious stationary-fuel reactor driven by the same power history,

$$\frac{dC_i^s}{dt} = \frac{\beta_i}{\Lambda} N(t) - \lambda_i C_i^s(t). \tag{3}$$

The drift reactivity is then defined by an equivalence requirement: it is the reactivity that, when added to the stationary-fuel model driven by the shadow inventories, reproduces the population equation of the circulating model. Requiring

$$\frac{dN}{dt} = \frac{\rho + \rho_{\text{drift}} - \beta_0}{\Lambda} N(t) + \sum_i \lambda_i C_i^s(t) \qquad (4)$$

to coincide with Eq. (1) and eliminating the common terms yields

$$\rho_{\text{drift}}(t) = -\frac{\Lambda}{N(t)} \sum_i \lambda_i \left(C_i^s(t) - C_i^{\text{core}}(t) \right). \qquad (5)$$

Because C_i^s and C_i^{core} both carry the same ordinary precursor-lag deficit, their difference cancels the lag term, and ρ_{drift} therefore measures the flow effect alone, at all times.

Two properties of Eq. (5) fix its physical meaning. First, the residual $\rho_{\text{drift}} + \rho_{\text{loss}}$ equals the lag term at all times. At constant power the lag vanishes, so the observer reduces exactly to $-\rho_{\text{loss}}(t)$. This equality is the principle by which the MSRE benchmark analyses read the drift reactivity from the inventory[4], and the observer generalizes that practice. Second, at zero flow the shadow and core inventories coincide, so $\rho_{\text{drift}} \equiv 0$ for any power history. A stationary-fuel reactor exhibits no drift reactivity.

The observer is implemented as n_g scalar ordinary differential equations attached to the host model as a strict observer: none of its outputs feeds back into the kinetics.

2.2. Exponential-mode closed form at fixed flow

In the one-dimensional uniform-velocity core-loop channel with core and loop transit times τ_C and τ_L, a power history on a single exponential mode $N \propto e^{\omega t}$ turns the precursor transport equation into its steady form with $\lambda_i \to \lambda_i + \omega$. This substitution holds for any axial power profile. For the spatially uniform profile used throughout this paper (Section 3), the in-core fraction of the modal inventory follows in closed form as $F(\lambda_i + \omega)$ with

$$F(\mu) = 1 - \frac{(1 - e^{-\mu\tau_C})(1 - e^{-\mu\tau_L})}{\mu\tau_C (1 - e^{-\mu(\tau_C + \tau_L)})}, \qquad (6)$$

which for $\mu = \lambda_i$ is exactly the steady-state $\beta_{i,\text{eff}}/\beta_i$ factor of the uniform-profile analytic solution[2]. The observer output on the mode follows in closed form,

$$\rho_{\text{drift}}^{\infty}(\omega) = -\sum_i \beta_i \frac{\lambda_i}{\lambda_i + \omega} [1 - F(\lambda_i + \omega)], \qquad (7)$$

with $\rho_{\text{drift}}^{\infty}(0) = -\rho_{\text{loss}}$. The magnitude $|\rho_{\text{drift}}^{\infty}|$ decreases monotonically in ω because on a rising power the in-core inventory is dominated by recently produced precursors whose advective loss has not yet materialized. Even at strictly constant flow, the drift reactivity is therefore a functional of the power history. On the same mode the lag term also closes: $\rho_{\text{drift}} + \rho_{\text{loss}} \to \sum_i \beta_i \, \omega / (\lambda_i + \omega)$.

2.3 Modified inhour relation for the circulating loop

Inserting the modal inventories into the kinetics relates the inserted reactivity, measured from the circulating critical state, to the asymptotic root ω through a delayed-return factor $G_i(\omega)$:

$$\Delta\rho(\omega) = \Lambda\omega + \sum_i \beta_i \left[1 - \frac{\lambda_i}{\lambda_i + \omega} G_i(\omega) \right] - \rho_{\text{loss}}. \qquad (8)$$

Three choices of G_i define three models: $G_i = 1$ is the textbook stationary-fuel inhour relation; $G_i = F(\lambda_i)$ is the common engineering practice of substituting a static β_{eff} into the textbook formula; $G_i = F(\lambda_i + \omega)$ is the exact circulating-loop relation implied by Eq. (6). The difference between the first and the third is precisely the modal drift recovery $\rho_{\text{drift}}^{\infty}(\omega) - \rho_{\text{drift}}^{\infty}(0) > 0$: the circulating loop sustains a given period with a smaller external insertion than the stationary-fuel relation requires.

3. Computational Framework

The verification model is a one-dimensional core--loop channel in Modelica/TRANSFORM: a core of $n_C = 10$ first-order upwind finite-volume cells with distributed fission source, an external loop, and a recirculation closure of the DNP concentrations, coupled to the modified point kinetics[1]. The parameters are MSRE-representative: the six-group ^{235}U data of [2], $\beta_0 = 0.006781$, $\Lambda = 2.4 \times 10^{-4}$ s, transit times $\tau_C = 10.3$ s and $\tau_L = 14.9$ s, uniform power profile, and isothermal salt except in the feedback variant, which adds a fuel heat source, a fixed core-inlet temperature, and a fuel-temperature coefficient (FTC) $\alpha = -1 \times 10^{-4}$ K^{-1}. The observer of Eqs. (3) and (5) is attached as six auxiliary equations.

The steady state of this model is verified separately[6]: at $n_C = 10$ the steady-state drift loss is $\rho_{\text{loss}} = 226.4$ pcm. The difference from the analytic value 232.3 pcm of [2] is the spatial discretization error, and mesh refinement confirms its first-order decrease (observed order 1.00). Every transient reference below is therefore evaluated on the same mesh, so that the known mesh bias cancels in the comparisons. Before each transient the model is at the circulating critical state, with the steady-state identity satisfied to $|\rho_{\text{drift}} + \rho_{\text{loss}}| \leq 1.7 \times 10^{-9}$ (dimensionless absolute reactivity). Throughout, ρ_{code} denotes the code's built-in reactivity sum, and $\rho_{\text{tot}} = \rho_{\text{code}} + \rho_{\text{drift}}$ denotes the drift-inclusive total. The asymptotic root ω is fitted from the simulated power history over the final 100 s of each run.

4. Benchmark I: Flow Step with Free Power

The fuel-salt flow is stepped from 100% to 50% at fixed external reactivity, and the power responds freely: Case A without feedback, Case B with the FTC. Fig. 1 shows the reactivity channels of both cases.

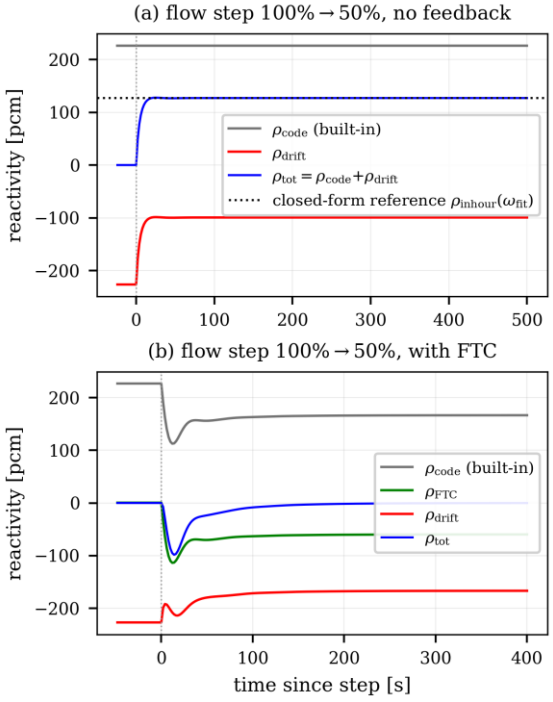


Fig. 1. Flow step with free power: reactivity channels without feedback (a) and with FTC (b)

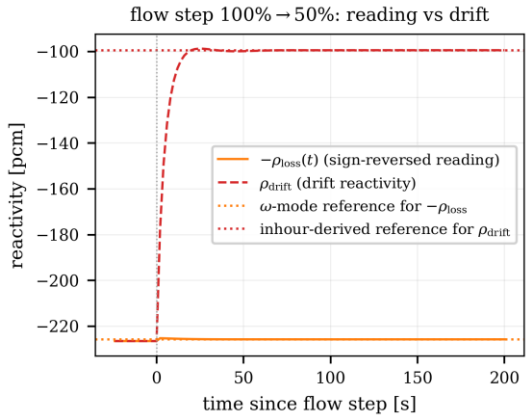


Fig. 2. Flow step with free power, Case A: the sign-reversed reading $-\rho_{\text{loss}}(t)$ and the drift reactivity $\rho_{\text{drift}}(t)$ with their closed-form references

In Case A the reduced flow leaves more precursors in the core, the delayed source rises, and the power climbs 6.9 decades onto a stable period with $\omega = 0.03095$ s^{-1} ($T = 32.3$ s), while ρ_{code} remains constant. On the asymptotic mode the drift-inclusive total must equal the textbook inhour value $\rho_{\text{inhour}}(\omega) = \Lambda\omega + \sum_i \beta_i\, \omega/(\lambda_i + \omega)$, a mesh-free reference. The simulated $\rho_{\text{drift}}(t_{\text{end}})$ matches $\rho_{\text{inhour}}(\omega) - \rho_{\text{ext}}$ to a residual of 5.1×10^{-11}. The inventory reading converges to a different value, $\rho_{\text{loss}} = 225.7$ pcm. This value matches its own reference, the steady-state solution of the code's discrete equations evaluated at $\lambda_i + \omega$, to 1.0×10^{-11}. The gap between the two converged values is the precursor lag. Reading and drift reactivity therefore diverge whenever the power varies. Fig. 2 displays both histories with their closed-form references.

In Case B the feedback terminates the rise: the power settles at 64.5% with the effective fuel temperature higher by 6.0 K, and the drift reactivity recovers by +60.20 pcm. With the drift channel included, the three-channel reactivity balance residual is at machine precision, max $\left|\rho_{ext} + \rho_{FTC} + \rho_{drift}\right| = 1.4 \times 10^{-13}$, over the final 100 s of the run. The built-in sum instead remains at +166.7 pcm at the settled steady state (Fig. 1b).

5. Benchmark II: Reactivity Step at Fixed Flow

The external reactivity is stepped by $\Delta\rho_{\text{ext}} \in \{+20, +50, +100\}$ pcm at constant flow without feedback. Fig. 3 shows the power rise and the net insertion.

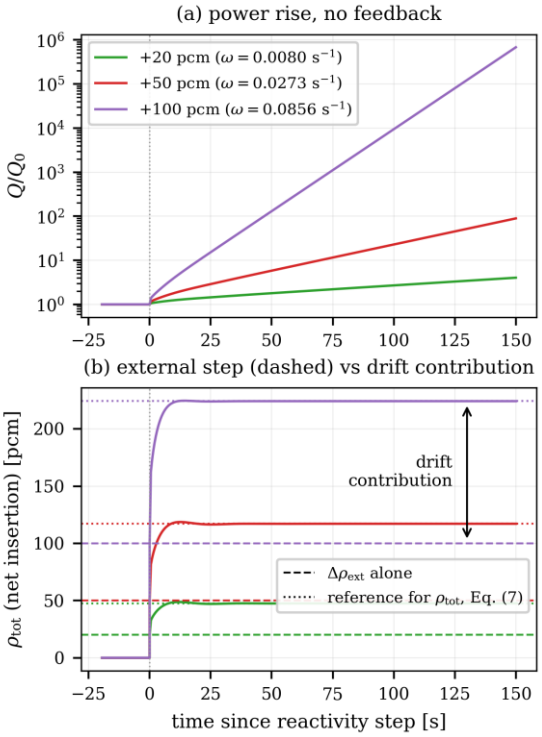


Fig. 3. Reactivity step at fixed flow: power rise (a) and net insertion ρ_{tot}(b) with the exponential-mode reference of Eq. (7).

Although the flow never changes, the drift channel responds to the power rise. The drift reactivity recovers by +67.1 pcm at $\omega = 0.02725$ s^{-1} for the +50 pcm step, and by +124.3 pcm at $\omega = 0.086$ s^{-1} for the +100 pcm step. The observer converges to the exponential-mode closed form Eq. (7) within 4.1×10^{-10} for all three step sizes, and the residual $\rho_{\text{drift}} + \rho_{\text{loss}}$ converges to the modal

lag term within 1.5×10^{-9}. As in the flow-step benchmark, $-\rho_{\text{loss}}(t)$ is not a drift reactivity. Fig. 4 displays the divergence for the +50 pcm case. The two histories separate immediately after the step and converge to different closed-form references. The gap between them, the modal lag term of 116.5 pcm, would be misattributed if the constant-power practice were carried over.

The fitted ω agrees with the numerical root of the exact circulating relation, Eq. (8) with $G_i = F\ (\lambda_i + \omega)$, to a relative difference of 0.8×10^{-5}. Against this reference, both stationary-fuel models underestimate the power rise rate: for a +50 pcm insertion, the textbook inhour relation predicts a period 3.2 times longer (117.6 vs. 36.7 s), and the common static-β_{eff} substitution a period +28% longer (47.1 s). This underestimation is the non-conservative direction for a reactivity insertion assessment (Fig. 5).

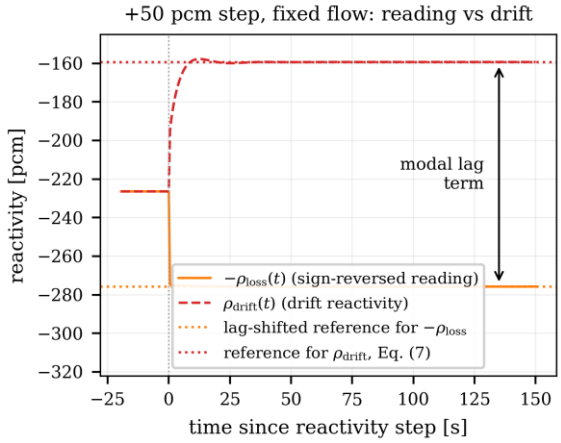


Fig. 4. Reactivity step at fixed flow, +50 pcm: the sign-reversed reading $-\rho_{\text{loss}}(t)$ and the drift reactivity $\rho_{\text{drift}}(t)$ with their closed-form references.

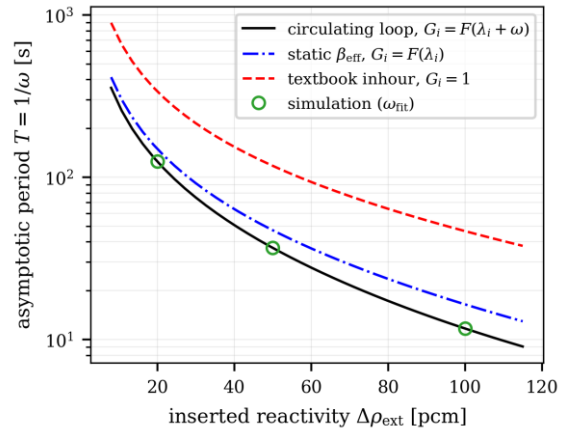


Fig. 5. Asymptotic period versus inserted reactivity: the three delayed-return models of Eq. (8) and the simulated periods

6. Conclusions

A shadow-precursor observer, implemented as six auxiliary differential equations in a non-intrusive and mesh-free form, extends the quantification of MSR drift reactivity to arbitrary power histories. The previous practice was exact only at constant power, where it reduces to the inventory reading $-\rho_{\text{loss}}(t)$[4]. Reported together with the code's built-in sum as $\rho_{\text{tot}} = \rho_{\text{code}} + \rho_{\text{drift}}$, the observer supplies the drift channel that ρ_{code} alone does not carry. In Modelica/TRANSFORM the observer was verified in both free-power transient regimes against closed-form references: the inhour-derived asymptote in a flow step and the exponential-mode closed form in reactivity steps, with residuals below 10^{-3} pcm in every case.

Two practical findings follow. First, the drift reactivity varies with the power history even at strictly constant flow (by +67.1 pcm for a +50 pcm insertion), so tabulating it against flow alone is insufficient. Second, the common static-β_{eff} period practice underestimates the power rise rate of the circulating reactor and is therefore non-conservative for reactivity insertion assessments. The complete derivations, the constant-power tier, and the full verification matrix are given in [6].

ACKNOWLEDGEMENTS

This work was supported by the National Research Foundation of Korea (NRF) funded by the Ministry of Science and ICT of the Republic of Korea (MSIT) (RS-2023-00259713).

REFERENCES

[1] M. S. Greenwood and B. Betzler, Modified Point-Kinetics Model for Neutron Precursors and Fission Product Behavior for Fluid-Fueled Molten Salt Reactors, Nuclear Science and Engineering, Vol. 193, No. 4, pp. 417–430, 2019.

[2] J. J. Jeong and Y. J. Cho, Analytical Derivation of Steady-State Delayed Neutron Precursor Distribution and Effective Delayed Neutron Fraction in a Molten Salt Reactor, Nuclear Engineering and Technology, Vol. 58, p. 104160, 2026.

[3] B. E. Prince et al., Zero-Power Physics Experiments on the Molten-Salt Reactor Experiment, ORNL-4233, Oak Ridge National Laboratory, 1968.

[4] J. J. Jeong et al., Benchmarking the MARS Code for Molten Salt Reactor Applications Using MSRE Transient Experiments, Nuclear Engineering and Technology, Vol. 58, p. 104438, 2026.

[5] M. S. Greenwood et al., Demonstration of the Advanced Dynamic System Modeling Tool TRANSFORM in a Molten Salt Reactor Application via a Model of the Molten Salt Demonstration Reactor, Nuclear Technology, Vol. 206, No. 3, pp. 478–504, 2020.

[6] J. Lee and S. Paek, On-line Quantification of Delayed Neutron Precursor Drift Reactivity in Molten Salt Reactors: a Shadow-Precursor Observer Verified against Closed-Form Benchmarks in Modelica/TRANSFORM, submitted to Nuclear Engineering and Technology, 2026.