Development of the Molten Salt Fast Reactor Core Analysis Code CA-FMSR and Its Verification against the CNRS Multiphysics Benchmark

Seungsu Yuk*, Tae Young Han, Sang Ji Kim
Korea Atomic Energy Research Institute, 111, Daedeok-daero 989beon-gil, Daejeon, 34057, Korea
**Corresponding author: syuk@kaeri.re.kr*

Keywords : molten salt reactor, reactor analysis code, multiphysics benchmark, delayed neutron precursor advection

1. Introduction

A molten salt reactor (MSR) is a liquid-fuel reactor in which the fuel is dissolved in a molten salt that circulates through the core and an external loop. This gives rise to coupled phenomena that the core analysis framework of a solid-fuel reactor cannot readily handle. Delayed neutron precursors (DNPs) are carried along with the fuel salt, so they change the delayed neutron source distribution. DNPs decay partly outside the core, reducing the effective delayed neutron fraction. Fission heat is transported directly by the flow, and the decrease in salt density with rising temperature provides reactivity feedback through the macroscopic cross sections. The core design of an MSR therefore requires a core analysis code capable of representing the coupled neutron-flux/DNP-advection/thermal-hydraulics behavior.

Accordingly, the present work is developing CA-FMSR (Core Analyzer for Fast-spectrum Molten-Salt Reactors), a reactor core analysis code for molten salt reactors. CA-FMSR performs multigroup neutron diffusion or SP3 calculations for MSR cores with unstructured geometry. It also provides the capabilities for the physics specific to liquid fuel: delayed neutron precursor advection with the external-loop recirculation boundary condition, temperature/density feedback, and transient analysis.

To verify the capabilities for the coupled physics specific to MSRs, the CNRS benchmark [1], a code-to-code numerical benchmark dedicated to MSR analysis codes, was analyzed. Six codes from four institutions participated in the benchmark and provided the reference data (raw data publicly available [2]), and it has been used as a standard verification problem by subsequent MSR codes such as Moltres [3] and Griffin [4]. This paper compares the CA-FMSR results with those of the benchmark participants and presents the verification results and conclusions.

2. CA-FMSR Core Analysis Methods

Fig. 1(a) shows the functions of CA-FMSR and its input/output flow. CA-FMSR consists of a C++ computational core and a Python input/output layer. It takes as input the multigroup constants and equivalence parameters generated by a transport analysis code such as a Monte Carlo code. When accurate thermal-hydraulic phenomena need to be considered, thermal-hydraulic data or the results of a thermal-hydraulic analysis code are used as input. CA-FMSR produces the effective multiplication factor, neutron flux, power distribution, and other core quantities.

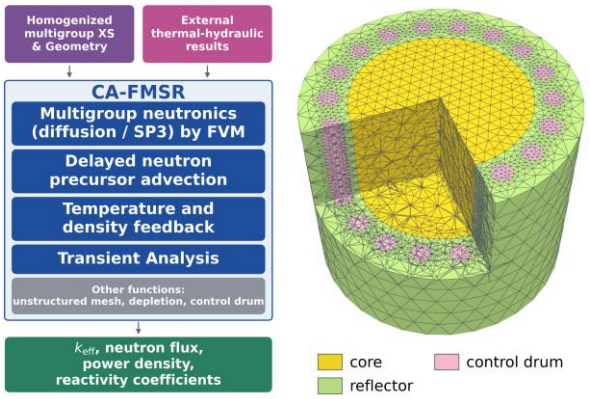


Fig. 1. Features of the CA-FMSR code.

2.1 Multigroup Diffusion/SP3 Finite Volume Solver

CA-FMSR supports the multigroup diffusion and SP3 solvers [5]. For the analysis of unstructured cores, both structured and unstructured meshes (Fig. 1(b)) are supported, and the analysis is performed with the finite volume method (FVM). CA-FMSR also implements core depletion and control drum rotation capabilities. Since the purpose of this paper is to verify the coupled physics specific to MSRs, their verification is planned for future work.

2.2 Delayed Neutron Precursor Advection Model

In a solid-fuel reactor the DNPs decay where they are produced, but in a liquid fuel they move with the fuel salt, so an advection term is added to the family-wise DNP balance. At steady state,

$$\nabla \cdot (\vec{u} C_i) + \lambda_i C_i = \frac{\beta_i}{k_{\text{eff}}} \nu\Sigma_f \phi \qquad (1)$$

where C_i, λ_i and β_i are the concentration, decay constant and production fraction of the i-th precursor family, $\vec{u}$ is the fuel salt velocity, ϕ is the neutron flux, $\nu\Sigma_f\phi$ is the fission neutron production rate, and k_{eff} is the multiplication factor. DNP calculation is performed via the FVM in the same framework as the neutron flux calculation. The advection term is discretized with a first-

order upwind scheme [6] so that cell-wise precursor conservation holds, and a zero normal flux condition is imposed at the walls. A recirculation boundary condition is supported, in which precursors leaving the core outlet re-enter the inlet after decaying as $e^{-\lambda_i \tau}$ over the external loop transit time τ; the closed-cavity problem of the CNRS benchmark corresponds to the special case $\tau = 0$. In the eigenvalue calculation the fission neutron emission is assembled as a prompt component $(1-\beta)\chi_p \nu\Sigma_f \phi/k_{\text{eff}}$ and a delayed component $\sum_i \lambda_i C_i \chi_{d,i}$, so that the effect of delayed neutron emission is consistently reflected in the steady-state solution as well (χ_p and $\chi_{d,i}$ are the prompt and delayed neutron emission energy spectra, $\beta = \sum_i \beta_i$).

2.3 Temperature/Density Feedback and Transient Analysis

The temperature feedback of a liquid fuel appears mainly through the density. As the salt temperature rises, thermal expansion reduces the density, and since the atom number density is proportional to the density, the macroscopic cross sections decrease by the same ratio. When the density follows a linear law $\rho(T) = \rho_{\text{ref}}[1 - \alpha(T - T_{\text{ref}})]$ as in the CNRS benchmark (α: volumetric expansion coefficient), the macroscopic cross sections follow

$$\Sigma(T) = \Sigma_{\text{ref}}[1 - \alpha(T - T_{\text{ref}})] \qquad (2)$$

The CA-FMSR macroscopic cross section library reflects this temperature dependence and is interpolated and reassembled cell by cell during the analysis, so that this feedback is accounted for automatically.

In transient analysis, the system of equations is advanced with the time derivative $(1/v)\partial\phi/\partial t$ added to the flux equation and $\partial C_i/\partial t$ to the precursor balance. The flux is discretized with backward Euler, while the precursor equations are integrated analytically with the fission source held fixed within the time step, so that the decay term is treated exactly. When the cross sections vary in time, for instance through temperature feedback, the system matrix is reassembled with the new cross sections at every time step.

3. CNRS Benchmark Analysis

3.1 Benchmark Problem Description

The problem domain is a two-dimensional square cavity 2 m on a side, filled with a LiF-BeF_2-UF_4 molten salt containing U-235. The top lid slides at 0.5 m/s, producing laminar forced convection; buoyancy adds natural convection; and the heat is removed by a volumetric heat sink ($\gamma = 10^6$ W/m³K). The neutron boundaries are vacuum on all four sides, and the precursors are in a closed cavity with zero normal flux at the walls. The benchmark has a ladder structure that switches on the physics one at a time: single-physics checks (Step 0.1 flow, 0.2 static neutronics, 0.3 temperature), added coupling (1.1 precursor advection, 1.2 temperature/density feedback, 1.3 buoyancy, 1.4 fully coupled, 30 operating conditions), and the transient coupling analysis (2.1). The group constants are the six groups specified by the benchmark with eight delayed neutron families ($\beta_{\text{tot}} = 688.25$ pcm), and the temperature dependence of the density is $\rho(T) = 2000[1 - 2\times10^{-4}(T - 900)]$ kg/m³. The problem domain is shown in Fig. 2. A detailed description of the benchmark problem can be found in [1].

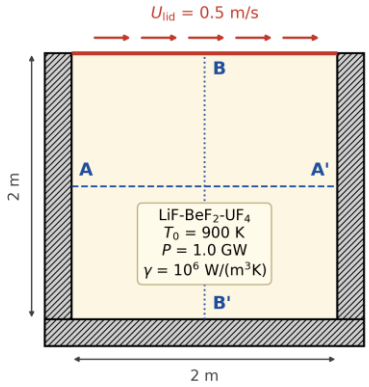


Fig. 2. CNRS benchmark problem domain with the horizontal and vertical centerlines AA' and BB'.

Since CA-FMSR is a three-dimensional code, the two-dimensional problem is set up as a quasi-two-dimensional mesh with a single cell in the z direction and reflective boundaries on the top and bottom faces (vacuum on the four x/y sides); the power is normalized to this 1 cm depth. The same 200×200 mesh (Δ = 1 cm) is used for all steps.

The numerical results for each problem are assessed against the mean curve of the six participating codes and the inter-code distribution band. The agreement of a distributed quantity is quantified by the normalized L_2 metric specified by the benchmark [1] (hereafter the discrepancy), $\varepsilon = [\sum(Q - \overline{Q})^2/\sum\overline{Q}^2]^{1/2}$, where Q is the CA-FMSR value at a sample point and $\overline{Q}$ is the six-code mean at the same point; the sum is taken over the 201 sample points of the centerline profiles [2].

3.2 Generation of the Flow and Temperature Field Inputs

CA-FMSR is designed to treat the flow and temperature fields as external inputs for solving the neutron flux and DNP distributions; accurate analysis of the flow and temperature fields is outside the functional scope of CA-FMSR. Since this benchmark nevertheless requires the full coupling of the three physics, two simple auxiliary modules were written separately to supply the inputs: (i) a two-dimensional laminar flow module based on the stream function–vorticity (ψ–ω) formulation, and (ii) a steady energy equation module based on a first-order upwind FVM. Before being supplied to CA-FMSR, the results of both modules were independently verified

against the benchmark reference data (Step 0.1 velocity discrepancy u_x on BB' 0.26 % and u_y on AA' 0.51 %, and Step 0.3 temperature T on AA' 0.21 % and on BB' 0.16 %, all on a par with the participating codes), and the buoyancy used the standard gravity $g = 9.81$ m/s² without adjustment. The flow field (streamlines) and the temperature field produced by the two modules are shown in Fig. 3. When the externally generated velocity field is supplied to the precursor advection solver, it is corrected by a least-squares projection so as to satisfy the discrete continuity equation of the solver.

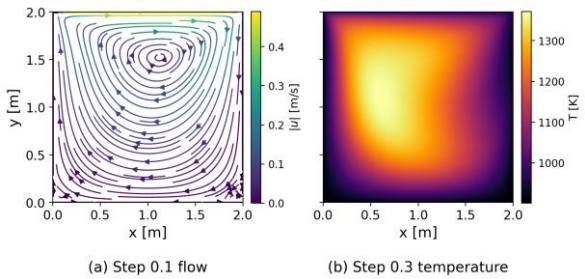


Fig. 3. Output of the flow/temperature-field generation.

3.3 Multiphysics Coupling and Transient Analysis Procedure

There are four multiphysics coupling links: the power feeds the heat source, the temperature enters the cross sections and the buoyancy, and the flow transports both the precursors and the heat. The steady coupled solution is obtained by Picard fixed-point iteration (relaxation factor 0.5) until the reactivity change falls below 5×10^{-7}.

The transient problem (Step 2.1) evaluates the frequency response of the system. When the heat sink coefficient is oscillated as a sine wave of frequency f, $\gamma(t) = \gamma_0[1 + 0.1\sin(2\pi f t)]$, the core power oscillates at the same frequency; the ratio of the relative amplitude of the power oscillation to the relative amplitude of the perturbation (10 %) is the gain, and the time lag of the power oscillation relative to the perturbation, expressed as an angle, is the phase. For seven frequencies (0.0125–0.8 Hz) the gain and phase were obtained with the CA-FMSR transient solver (200 steps per period, temperature-dependent cross sections reassembled at every step). During the transient the neutron flux and the DNPs are solved with the time-dependent system, and the temperature field is advanced together by implicit Euler with the energy equation including thermal inertia (the time-derivative term). The flow field is frozen at that of the Step 1.4 steady coupled solution. This is an approximation that neglects not the buoyancy itself but only its oscillatory variation driven by the temperature oscillation.

4. Numerical Results

4.1 Static Neutronics and Precursor Advection

Table I summarizes the reactivity results for each step. The six-code column is the [minimum, maximum] range of the participating codes. The results of Moltres and Griffin, which solved the same problems later, are listed alongside.

The static reactivity (Step 0.2) of 466.1 pcm, computed with the CA-FMSR diffusion solver, agrees to within 0.5 pcm with Moltres (465.6) and Griffin (465.7). The reactivity loss due to precursor advection (Step 1.1) of −62.73 pcm differs by 0.6 pcm from the six-code average of −62.13 pcm, and Fig. 4 shows that the delayed neutron source distribution skewed upward by the vortex is likewise reproduced (discrepancy AA' 0.75 %, BB' 0.55 %).

Table I: Reactivity and reactivity defect results [pcm].

Step	Item	CA-FMSR		Six codes [1]	Moltres [3]	Griffin [4]
		Diffusion	SP3			
0.2	ρ	466.1	521.6	[353.7, 578.1]	465.6	465.7
1.1	$\rho - \rho_{0.2}$	−62.73	−62.30	[−63.0, −60.7]	−62.7	−61.4
1.2	$\rho - \rho_{1.1}$	−1141.7	−1130.4	[−1161.0, −1122.0]	−1142.2	−1149.3
1.3	$\rho - \rho_{0.2}$	−1206.4	−1194.8	[−1227.0, −1184.4]	−1207.7	−1212.0

In contrast to the reactivity differences, which agree among the codes to within a few pcm, the absolute reactivity of Step 0.2 is spread over 354–578 pcm, about 220 pcm, across the participating codes. This spread appears to originate mostly from the differences in the neutron flux calculation methodology (diffusion, SP3, S_N, and so on). Indeed, changing only the method or the angular approximation within the same code shifts the reactivity (from 411.3 with CNRS SP1 to 353.7 with SP3, and from 482.6 with TU Delft S2 to 578.1 pcm with S6). This offset cancels in the reactivity differences and therefore does not affect the $\rho - \rho_{0.2}$ agreement in Table I.

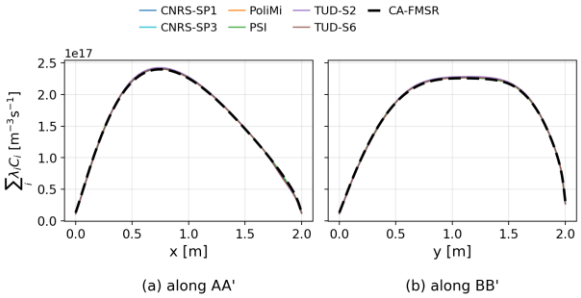


Fig. 4. Step 1.1 delayed neutron source distribution.

4.2 Coupled Temperature/Density Feedback

With the temperature/density feedback switched on, the reactivity defects of forced convection (Step 1.2) and buoyant natural convection (Step 1.3) both lie inside the six-code band and are within 0.5 pcm and 1.3 pcm,

respectively, of Moltres, a code using the same diffusion approximation (Table I).

Fig. 5 compares the vertical velocity, temperature and delayed neutron source distributions along AA' of Step 1.3 with the reference data. Each distribution is seen to be very close to those of the participating codes.

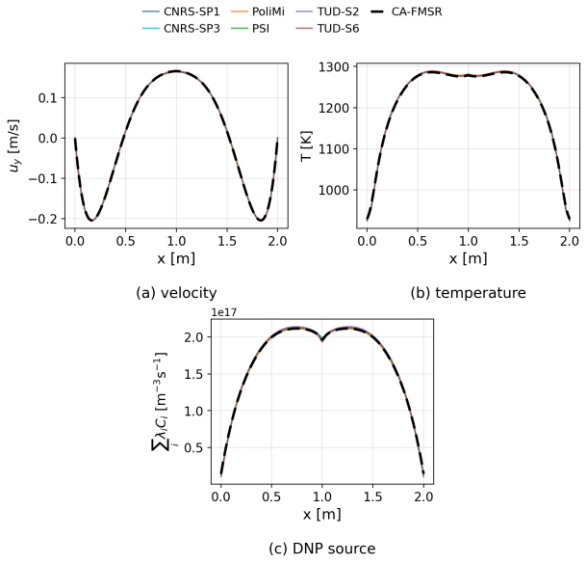


Fig. 5. Comparison of the distributions along AA' of Step 1.3.

The reactivity changes of the 30 operating conditions combining lid velocity and power (Step 1.4) are collected in Table II. All of them lie inside the participant band, and the difference from Moltres is 0.7 pcm RMS (1.3 pcm at most). In Table II the low-power column (P = 0.2 GW) shows the reactivity defect growing with lid velocity, whereas the columns with $P \geq 0.6$ GW move in the opposite direction. At low power an increase in the flow velocity pushes the precursor distribution away from the core center and enlarges the defect, while at higher power the temperature redistribution effect prevails and the trend reverses. This competition between the two mechanisms, including the reversal point, is reproduced in agreement with the reference results.

Table II: Step 1.4 reactivity changes of the 30 operating conditions (CA-FMSR diffusion, vs Step 0.2) [pcm]

U_{lid} [m/s] \ P [GW]	0.2	0.4	0.6	0.8	1.0
0.0	−263.8	−498.2	−730.4	−965.8	−1206.4
0.1	−266.0	−498.6	−730.2	−965.2	−1205.5
0.2	−268.2	−498.3	−729.0	−963.0	−1202.5
0.3	−270.0	−498.4	−727.4	−960.2	−1198.4
0.4	−272.0	−498.2	−726.0	−957.6	−1194.6
0.5	−274.2	−498.3	−725.0	−955.6	−1191.8

The same problems were also analyzed with SP3. The static reactivity (Step 0.2) is +521.6 pcm, inside the full participant range (354–578 pcm), and the reactivity loss of Step 1.1, involving precursor advection only, is essentially the same as diffusion to within 0.4 pcm. SP3, too, stayed inside the participant band for all steps and all operating conditions. In the coupled problems with the temperature/density feedback switched on, however, the reactivity defect was consistently smaller than that of diffusion by about 11 pcm (Step 1.2 −1130.4, Step 1.3 −1194.8 pcm).

4.3 Multiphysics Transient Analysis (Step 2.1)

Fig. 6 compares the gain and phase shift computed at the seven frequencies with the six participating codes. For both diffusion and SP3 the gain/phase curves overlap the bundle of the participating-code curves: they lie inside the participant spread at most frequencies, and even the points that fall outside, in the transition band where the gain rolls off steeply, deviate by less than 1.2 % (less than 0.1° for the phase). The frequency response characteristic is reproduced over the whole band: the power follows the perturbation at low frequency, and at high frequency the response is damped by the precursor/thermal inertia while the phase lag grows. The difference between diffusion and SP3 is also smaller than the participant spread.

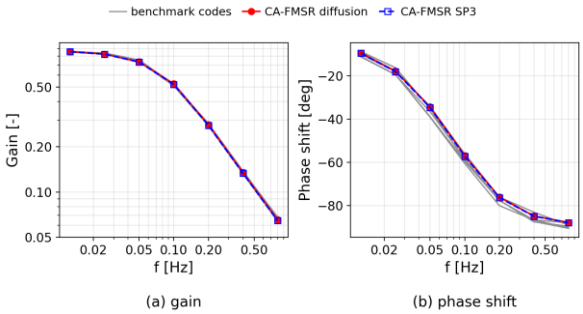


Fig. 6. Step 2.1 multiphysics transient analysis.

5. Conclusions

The fast-spectrum molten salt reactor core analysis code CA-FMSR was developed and verified against the full set of CNRS multiphysics benchmark problems. The CA-FMSR reactivity results for the steady-state and multiphysics-coupled problems were all inside the participating-code bands, and in the multiphysics transient analysis the gain and phase shift of the frequency response were likewise nearly identical to those of the benchmark participants for both diffusion and SP3. The analysis capability of CA-FMSR for the coupled physics specific to MSRs (delayed neutron precursor advection, temperature/density feedback, and the transient behavior) is thereby confirmed against the CNRS benchmark.

The limitations of this verification are as follows. First, the causes of the reactivity and reactivity-defect differences between SP3 and diffusion have not been precisely identified. Second, the flow/temperature input modules are a simple configuration prepared for this verification and are outside the functional scope of CA-FMSR. In future work the accuracy will be improved

through coupled calculations with an external code capable of accurately analyzing the flow and temperature of a molten salt reactor (for example GAMMA+ [7]). In addition, code-to-code verification against a Monte Carlo code will be carried out on a realistic core geometry (three-dimensional, with control drums).

ACKNOWLEDGEMENTS

This work was supported by Korea Research Institute for defense Technology planning and advancement (KRIT) grant funded by the Korea government (DAPA (Defense Acquisition Program Administration)) (KRIT-CT-22-017, Next Generation Multi-Purpose High Power Generation Technology (Liquid Fueled Heat Supply Module Design Technology), 2022).

REFERENCES

[1] M. Tiberga et al., Results from a multi-physics numerical benchmark for codes dedicated to molten salt fast reactors, Annals of Nuclear Energy, Vol. 142, 107428, 2020.
[2] M. Tiberga et al., CNRS benchmark raw results, Zenodo, doi:10.5281/zenodo.3648639, 2020.
[3] S. M. Park and M. Munk, Verification of moltres for multiphysics simulations of fast-spectrum molten salt reactors, Annals of Nuclear Energy, Vol. 173, 109111, 2022.
[4] M. K. Jaradat, N. Choi, and A. Abou-Jaoude, Verification of Griffin-Pronghorn-Coupled Multiphysics Code System Against CNRS Molten Salt Reactor Benchmark, Nuclear Science and Engineering, Vol. 198, pp. 2403-2436, 2024.
[5] P. S. Brantley and E. W. Larsen, The Simplified P_3 Approximation, Nuclear Science and Engineering, Vol. 134, pp. 1-21, 2000.
[6] Y. S. Jung, J. Fang, D. Shaver, and B. Feng, Multi-Physics Simulations of Molten Chloride Fast Reactor using Nek5000 and PROTEUS-NODAL, ANL/NSE-20/40, Argonne National Laboratory, 2020.
[7] H. S. Lim, GAMMA+ 2.0 Volume II: Theory Manual, KAERI/TR-8662/2021, Korea Atomic Energy Research Institute, 2021.