PolyEnergetics

Clean, safe, reliable nuclear energy

WHITE PAPER · REACTOR PHYSICS

# Validation of an OpenMC Model of the Molten Salt Reactor Experiment (MSRE)

POLYENERGETICS PRIVATE LIMITED

VERSION 1.0

KEYWORDS

Molten Salt Reactor Experiment, MSRE, OpenMC, ENDF/B-VIII.1, Benchmark Validation, Neutronics, PolyEnergetics

# Abstract

The Molten Salt Reactor Experiment (MSRE) zero-power experiments remain the primary validation basis for liquid-fuelled reactor modelling. This work presents a validation of the open-source Monte Carlo code OpenMC, using CAD-based geometry coupled through DAGMC and ENDF/B-VIII.1 nuclear data, against three experiments from the MSRE zero-power benchmark evaluation [1]. Initial criticality is over-predicted by +1158 pcm, roughly half the +2087 pcm bias of the benchmark's own reference Serpent 2 / ENDF/B-VII.1 calculation. The integral worth of the regulating rod over its full stroke is 2059.6 pcm (C/E = 0.94) and the peak differential worth is 22.64 pcm/cm (C/E = 0.93), with the deficit confined to the withdrawn end of the stroke rather than distributed across it. The $^{235}$U reactivity coefficient is K = 0.2303 ± 0.0040 (C/E = 1.033), agreeing with experiment to within one standard deviation. These results indicate that OpenMC with CAD-based geometry is well suited to molten salt reactor criticality and reactivity-coefficient analysis, while identifying the axial extent of the modelled control-rod absorber as the most likely source of the residual rod-worth deficit. Some of the remaining discrepancy reflects a lack of clarity in material properties and geometry extracted from decades-old ORNL documents, as acknowledged by the benchmark evaluation itself [1].

RESULTS AT A GLANCE

## +1158 pcm

Excess reactivity at initial criticality, against +2087 pcm for the reference Serpent 2 calculation

## 2059.6 pcm

Full-stroke integral worth of the regulating rod, C/E = 0.94

## 22.64 pcm/cm

Peak differential rod worth at 21 in withdrawn, C/E = 0.93

## K = 0.2303

$^{235}$U reactivity coefficient, ± 0.0040, C/E = 1.033 against experiment

READER'S NOTE

A validation compares what a simulation code predicts against what a real reactor did. The MSRE, operated at Oak Ridge National Laboratory in 1965, is the only molten salt reactor for which such measurements exist. This paper reproduces three of those measurements in OpenMC and reports how closely the model matches them. C/E is the ratio of the calculated value to the experimental one, and pcm is a unit of reactivity equal to one hundred-thousandth of a change in the neutron multiplication factor.

CONTENTS

# 1. Introduction

Molten salt reactors (MSRs) were identified by the 2002 Generation IV International Forum roadmap as one of six advanced reactor concepts meriting continued development, and in the years since, renewed interest from both public research programs and private industry has driven multiple MSR designs toward licensing. This has sharpened a long-standing challenge for the technology: the reactor physics codes used to predict MSR safety characteristics such as criticality, control-rod worth, reactivity feedback must be validated against experimental data from a real circulating-fuel system, and very little such data exists. Almost all available integral benchmark data for liquid-fueled reactor physics comes from a single source: the Molten Salt Reactor Experiment (MSRE), designed at Oak Ridge National Laboratory (ORNL) beginning in 1960 to demonstrate the practicality of a high-temperature, fluid-fueled reactor concept [1].

In June 1965, following construction, ORNL conducted a series of zero-power startup experiments on the MSRE: initial criticality with $^{235}$U fuel, control-rod calibration, the reactivity equivalent of $^{235}$U addition, rod-bank worth and shadowing, the reactivity effect of fuel circulation, rod-drop transients, and the isothermal temperature coefficient. For decades this data existed only in ORNL's original design and operations reports, unavailable in a form suitable for direct code validation. Beginning in 2016, the University of California, Berkeley, in collaboration with ORNL, undertook the first formal reactor-physics benchmark evaluation of these experiments for the International Reactor Physics Experiment Evaluation Project (IRPhEP) handbook [1][2]. The first criticality experiment was published as a standalone benchmark in 2019, with a companion Serpent 2 / ENDF/B-VII.1 calculation reporting a multiplication factor of 1.02132 ± 3 pcm. Subsequent evaluations extended the benchmark to the full set of 1965 control-rod calibration and reactivity-coefficient experiments [1,2], providing, for the first time, a complete and independently reviewed reference dataset spanning the seven zero-power experiments.

This benchmark has since motivated independent validation efforts using several codes such as SCALE (Barlow and Clarno, 2020 [13]; Bostelmann et al., 2022 [9]), MCNP (Clarno et al., 2023 [14]), and independently for Serpent (Clarno et al., 2023 [14]). Yilmaz et al. [11] compares the impact of using geometry based on CAD to CSG using OpenMC and Serpent. To date, however, no published validation of the MSRE zero-power benchmark experiments has been carried out using OpenMC v0.15.3 — an open-source Monte Carlo code that has seen rapidly growing adoption in the reactor-physics community — in combination with CAD-based geometry representation and the current ENDF/B-VIII.1 nuclear data library.

This work addresses that gap for three of the benchmark experiments: initial criticality, regulating-rod worth, and the $^{235}$U reactivity coefficient. Results are compared against both the experimental values and the reference Serpent 2 calculations reported in the benchmark evaluation.

# 2. Input Data Used

This section describes the input data used for the three benchmark experiments addressed in this work. All experiments share the fuel-salt composition, materials, geometry, and cross-section library detailed below; per-experiment conditions (rod positions, fuel loading, temperature) are given individually for each experiment.

## 2.1 Fuel salt composition

Fuel-salt composition and properties of molten salt mixture used for analysis in OpenMC code are tabulated below. The value of all the parameters are taken from references as indicated.

TABLE 1. FUEL SALT SPECIFICATION OF THE OPENMC MODEL

| SR. | PARAMETER | VALUE |
|---|---|---|
| 1 | Salt type | LiF–$BeF_2$–$ZrF_4$–$UF_4$ eutectic [1] |
| 2 | Molar composition | LiF 64.88, $BeF_2$ 29.27, $ZrF_4$ 5.06, $UF_4$ 0.79 mol% [1] |
| 3 | Density @ 638 °C | 2.3278 g/cm³ [10] |
| 4 | Volumetric thermal-expansion coefficient | $\beta = +2.16 \times 10^{-4}$ /°C ($+1.20 \times 10^{-4}$ /°F) [12] |

## 2.2 Other materials

Structural and absorber material property values used for the analysis in the OpenMC code are tabulated below. These values are taken from the references as indicated. Density values are at room temperature.

TABLE 2. STRUCTURAL AND ABSORBER MATERIAL SPECIFICATION OF THE OPENMC MODEL

| PARAMETER | VALUE |
|---|---|
| **Graphite – moderator** | |
| Density | ρ = 1.8492 g/cm³ [10] |
| Linear thermal-exp coefficient | α = 2.7 × $10^{-6}$ per °C (1.5 × $10^{-6}$ /°F) [1] |
| Natural B impurity | 0.80 ppm wt total B [1] |
| **INOR-8 (Hastelloy N) – vessel and piping** | |
| Density | ρ = 8.7745 g/cm³ [1] |
| Linear thermal-exp coeff | α = 1.404 × $10^{-5}$/°C (7.8 × $10^{-6}$ /°F) [1] |
| Composition | Full 14-element / 47-isotope specification [1]; nominal 68.5 wt% Ni / 16.5 Mo / 7.0 Cr / 5.0 Fe; plus trace elements, isotope-resolved [10] |
| **Inconel-600 – rod thimble + cable** | |
| Density | ρ = 8.5 g/cm³ [1] |
| Composition | Isotope-resolved weight fractions per [10] |
| **Control-rod bush ($Gd_2O_3$ + $Al_2O_3$)** | |
| Density | ρ = 5.873 g/cm³ [1] |
| Composition (nominal) | 70 wt% $Gd_2O_3$ / 30 wt% $Al_2O_3$ |

## 2.3 Cross-section library & OpenMC configuration

The reference calculation of the benchmark evaluation [1] used Serpent 2 v2.1.30 [8] with ENDF/B-VII.1 [6]; the present work uses OpenMC 0.15.3 [3] with ENDF/B-VIII.1 [5] and differs in geometry, code and nuclear-data library, so comparisons with the reference calculation reflect the combined effect of both. The two library revisions differ in ways that bear directly on this system: ENDF/B-VIII.1 adopts revised $^{233,235,238}$U and $^{19}$F neutron evaluations from the IAEA-coordinated INDEN collaboration, and expands the thermal neutron scattering sublibrary substantially relative to ENDF/B-VIII.0, with many kernels re-evaluated or newly provided [5]. Both the fissile evaluations and the graphite scattering treatment are therefore expected to contribute to the eigenvalue difference reported in Section 4.1.

TABLE 3. COMPUTATIONAL CONFIGURATION OF THE OPENMC MODEL

| SR. | PARAMETER | VALUE |
|---|---|---|
| 1 | Monte Carlo code | OpenMC 0.15.3 [3] |
| 2 | Cross-section library | ENDF/B-VIII.1 [5] |
| 3 | Thermal-scattering data | ENDF/B-VIII.1 graphite S(α,β) [5] |
| 4 | Temperature treatment | Interpolation between tabulated library temperatures (250, 294, 600, 900, 1200, 2500 K) |
| 5 | Convergence settings | 200 batches × 50,000 particles per batch (50 inactive) |

# 3. Geometry

The reactor geometry is represented as a CAD-derived surface mesh (.h5m format) coupled to OpenMC through DAGMC [4]. This approach allows the vessel, core internals, and control-rod assemblies to be modeled with their as-built curvature and detail. The mesh dimensions follow the benchmark specification [1]. We have used the CAD models provided by the OpenMSR project [10] for our experiments.

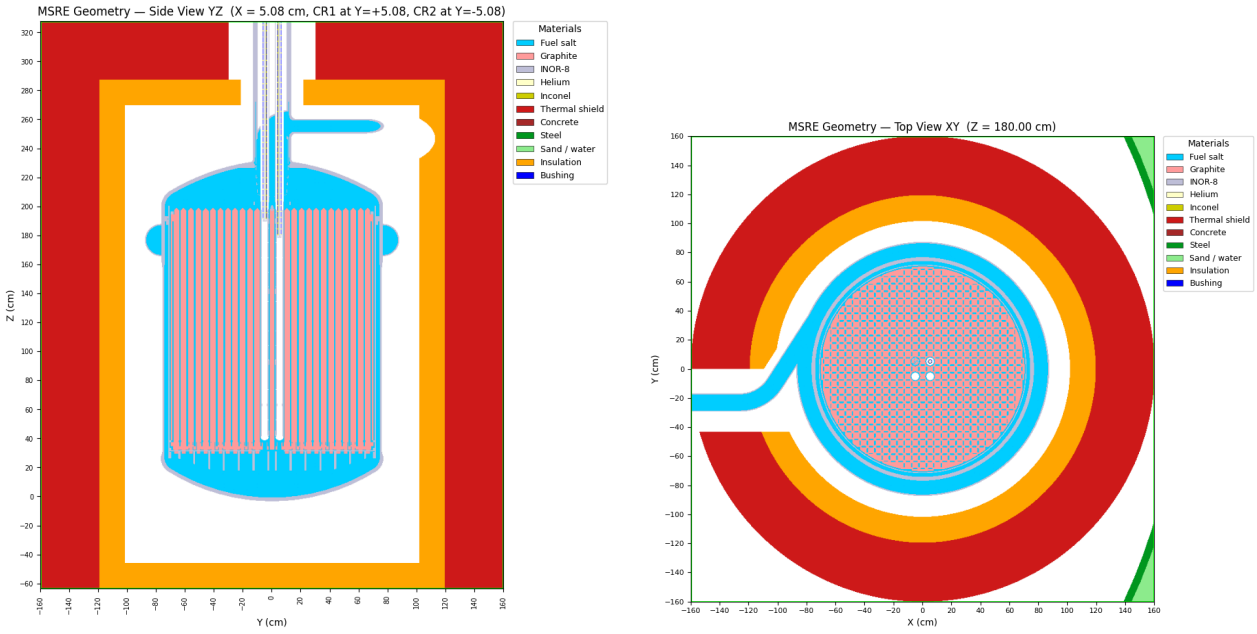


(A) (B)

Figure 1. (a) YZ plot (side view) of the core region. Rod 1 is withdrawn to 46.6 inches, Rod 2 and 3 are fully withdrawn to 51 inches. (b) XY plot (top view) of the core region at z=180 cm. Graphite lattice, control rod thimbles, sample basket, the fuel distributor inlet are visible.

The active core occupies a cylindrical region approximately 70.3 cm in radius and 165 cm in height, bounded top and bottom by the graphite-to-plenum transitions visible in Fig. 1 (a). Within this region, the core is built from a matrix of graphite stringers which is the moderating structure of the MSRE and with the fuel salt occupying the channels between adjacent stringers. This stringer-and-channel arrangement is what gives the core its graphite-moderated, salt-cooled/fueled character, distinct from a conventional solid-fuel lattice. Three control-rod thimbles penetrate the core axially, positioned at (5.08, ±5.08) cm from the core centerline (Fig. 1 (b)). Each thimble houses one control rod — rod 1 (the regulating rod) and rods 2 and 3 (the shim/safety rods) — and the Inconel-600 thimble wall separates the rod and its $Gd_2O_3$–$Al_2O_3$ absorber bushing (Table 2) from the surrounding fuel salt. Rod position is represented in the geometry as insertion depth along the thimble axis, spanning the benchmark's 0–51 in (fully inserted–fully withdrawn) stroke; Fig. 1 (a) shows this for rod 1 at its 46.6 in critical position.

The primary vessel and piping enclosing the core are modeled in INOR-8 (Hastelloy N), consistent with the material specification of Table 2. One geometric simplification carries through the model: the vessel's fixed salt-holding volume corresponds to a value below the benchmark's reported total loop inventory, since the benchmark's loop mass includes salt held in piping and the drain tank outside the modeled vessel volume. This mismatch is addressed by anchoring the salt composition rather than the salt mass. In every case the mass fractions of each salt constituent are held to the benchmark specification, and the absolute in-vessel salt and fissile masses then follow from the modelled volume and the salt density at the case temperature. Loading changes are represented by applying the benchmark's per-case composition and density directly (Table 8) rather than by scaling uranium fractions alone. A consequence is that the absolute salt masses quoted in Tables 4, 6 and 10 differ between experiments even where the $^{235}U$ mass fraction is identical; this is expected and does not indicate differing fuel compositions. The choice is discussed where it affects specific results, in Sections 4.1 and 4.3.

# 4. Results

The following subsections present the OpenMC results for three of the MSRE zero-power benchmark experiments. Section 4.1 addresses the initial criticality of the reactor on 1 June 1965. Section 4.2 covers the integral and differential worth of the regulating rod (rod 1) across its full stroke. Section 4.3 presents the $^{235}U$ reactivity coefficient, relating reactivity to fissile loading at fixed rod position. For each experiment, the calculated results are compared against the benchmark's reported experimental values and, where available, its reference Serpent 2 calculation.

## 4.1 Analysis of Initial Criticality

The first criticality of the MSRE was achieved on 1 June 1965. Carrier salt was pre-loaded into the primary loop, and highly enriched LiF–$UF_4$ eutectic (73LiF–27$UF_4$, 93% $^{235}U$) was added incrementally while the neutron source multiplication was monitored after each addition. The 93% figure refers to the enriching salt as added; because the carrier salt already contained depleted uranium, the resulting bulk uranium inventory in the loop was substantially lower in enrichment, at 31.4 wt% $^{235}U$ (Table 4). The just-critical state was reached at a fuel-salt $^{235}U$ mass fraction of 1.408 ± 0.007 wt%, with the core at 638 °C, the fuel salt stationary, control rod 1 withdrawn to 46.6 in, and rods 2 and 3 fully withdrawn (51 in) [1]. The benchmark evaluation assigns a value of k_eff = 0.99978 ± 0.00420, where the deviation from unity and the uncertainty reflect the benchmark model's bias assessment [1][2]. The reference Serpent 2 calculation reports k_eff = 1.02132 ± 0.00003 with ENDF/B-VII.1 [1].

### 4.1.1 Inputs

In the present work, a single eigenvalue calculation was performed at the first-criticality conditions: salt composition as tabulated in Table 1; 1.4103 wt% $^{235}U$ in salt; a uniform temperature of 638 °C; rod 1 at 46.6 in withdrawn, rods 2 and 3 fully withdrawn; thermal expansion applied to salt, graphite, INOR-8 components; and stationary fuel salt. Reactivity is expressed as ρ = (k − 1)/k in pcm relative to the critical condition. This configuration is tabulated in Table 4.

TABLE 4. CALCULATION CONDITIONS FOR THE INITIAL-CRITICALITY CASE

| PARAMETER | VALUE |
|---|---|
| Temperature | 638 °C |
| Rod 1 position | 46.6 in withdrawn |
| Rod 2 / 3 position | 51 in withdrawn (fully out) |
| $^{235}U$ loop loading | 1.4103 wt% at 31.4% enrichment of the loop uranium, corresponding to 61.215 Kg in 4340.43 Kg loop charge |
| Salt density | 2.3278 g/cm³ |
| Thermal expansion | Applied (salt, graphite, INOR-8) |
| Salt motion | Stationary |

Note: The 4340.43 kg figure is the benchmark's reported loop inventory at the time of first criticality. The corresponding benchmark inventories differ for the rod-calibration period (4626.52 kg, Table 6) and the $^{235}U$ loading sweep (4629.80 kg, Table 10), reflecting salt added to the loop between campaigns. These masses are quoted to identify the benchmark condition being reproduced; what is applied in the model is the salt composition in mass fractions, which is held to the benchmark in every case. Absolute masses should therefore not be compared across tables.

### 4.1.2 Comparison of k-effective with Experimental / other Simulation Data

The calculated eigenvalue is k_eff = 1.01172 ± 0.00035, an excess reactivity of +1158 pcm relative to the critical experiment. For comparison, the reference Serpent 2/ENDF-B-VII.1 calculation [1] overpredicts criticality by +2087 pcm.

First-criticality eigenvalue compared with experiment and the reference calculations. $\Delta\rho = (k - 1)/k$.

TABLE 5. FIRST-CRITICALITY EIGENVALUE: PRESENT WORK COMPARED WITH EXPERIMENT AND THE REFERENCE CALCULATIONS

| SR. | SOURCE | NUCLEAR DATA | K_EFF | ΔP VS. CRITICAL (PCM) |
|---|---|---|---|---|
| 1 | MSRE (1 June 1965) | - | 1.00000 (by definition) | 0 |
| 2 | Benchmark Model Expectation [2] | - | 0.99978 ± 0.00420 | −22 |
| 3 | Serpent 2 (reference calc.) [1] | ENDF/B-VII.1 [6] | 1.02132 ± 0.00003 | +2087 |
| 4 | **OpenMC (this work)** | ENDF/B-VIII.1 [5] | 1.01172 ± 0.00035 | +1158 |

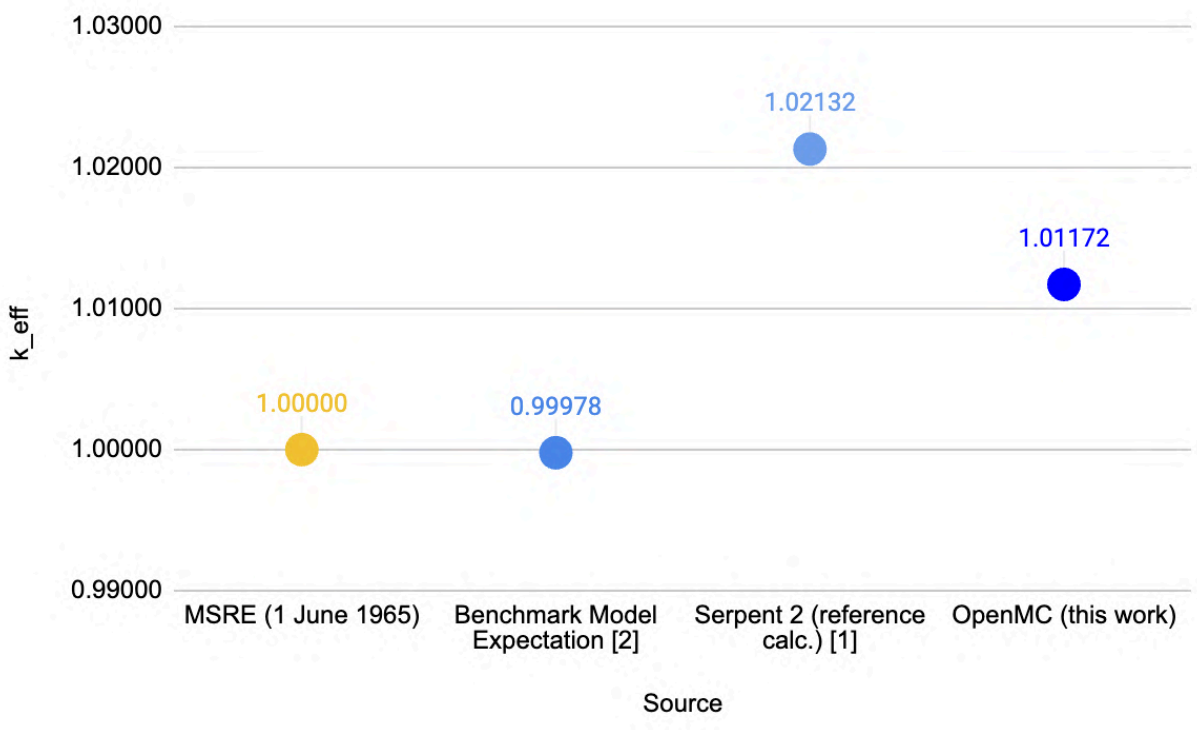


Figure 2 - k-effective comparison

## 4.2 Analysis of Control rod worth calculation

Control rod worth quantifies the change in core reactivity produced by moving a control rod. In this experiment Rod 1 is moved across its stroke — 51 inches withdrawn (fully withdrawn) to 0 inches withdrawn (fully inserted) — with Rods 2 and 3 held fully withdrawn throughout. A dedicated symmetric pair of positions is simulated around each reporting point, giving 17 independent k-eigenvalue solves at −2, 0, 2, 3, 4, 7, 11, 15, ..., 47, 51 inches withdrawn. (The −2 inch case lies 2 inches below nominal full insertion and exists only as the lower partner of the 0 inch reporting point.) Each solve yields k_eff, hence $\rho = (k-1)/k \times 10^5$ pcm, recorded in a ρ-vs-position table.

- **Integral worth** — the reactivity gained by withdrawing Rod 1 from fully inserted (0 inches withdrawn) out to position z: $W(z) = \rho(z) - \rho(0)$, so $W(0) = 0$ and $W(51)$ is the full-stroke rod worth.
- **Differential worth** — the reactivity difference across a symmetric pair of positions, divided by their separation and ascribed to the mean position: $f(x) = [\rho(x + \Delta x) - \rho(x - \Delta x)] / (2\Delta x)$, $\Delta x = 2$ inches. Reported at 14 positions (0, 2, 5, 9, ..., 45, 49 inches withdrawn) in pcm/cm.

### 4.2.1 Inputs used

The materials are as defined in Table 1 and Table 2, and the specific conditions applied for this experiment are as follows:

TABLE 6. CALCULATION CONDITIONS FOR THE ROD-WORTH SWEEP

| SR. | PARAMETER | VALUE |
|---|---|---|
| 1 | Sweep variable | Rod 1: 17 points from 51 in (fully withdrawn) to -2 in (beyond fully inserted, for differential worth calculation). Rods 2 and 3: Fully withdrawn (51 in) in all cases |
| 2 | Temperature | 649 °C (922.15 K) (rod-calibration period) |
| 3 | $^{235}$U loop loading | 1.4103 wt% at 31.4% enrichment, corresponding to 65.25 Kg in 4626.52 Kg loop salt [1] |
| 4 | Salt motion | Stationary |
| 5 | Thermal expansion | Not applied |

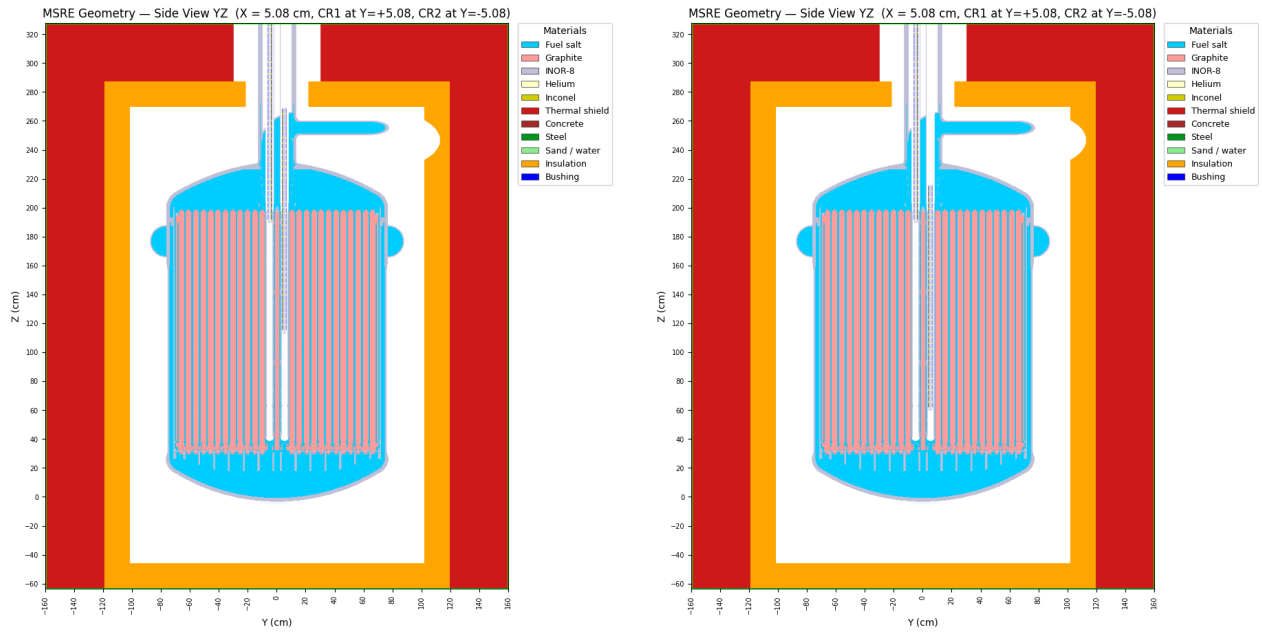


(A) (B)

Figure 3 - Control Rod positions in the MSRE core. (a) Control Rod 1 at 21 inches withdrawn. (b) Control Rod 1 at 0 inches withdrawn (fully inserted).

### 4.2.2 Comparison of integral and differential rod worth with Benchmark / other simulation data

Raw simulated reactivity ρ (pcm) at each of the 15 swept positions plus the two endpoints (51 in, -2 in) was calculated. A Monotonic decrease in ρ as the rod is inserted is observed.

#### 4.2.2.1 Integral Rod Worth

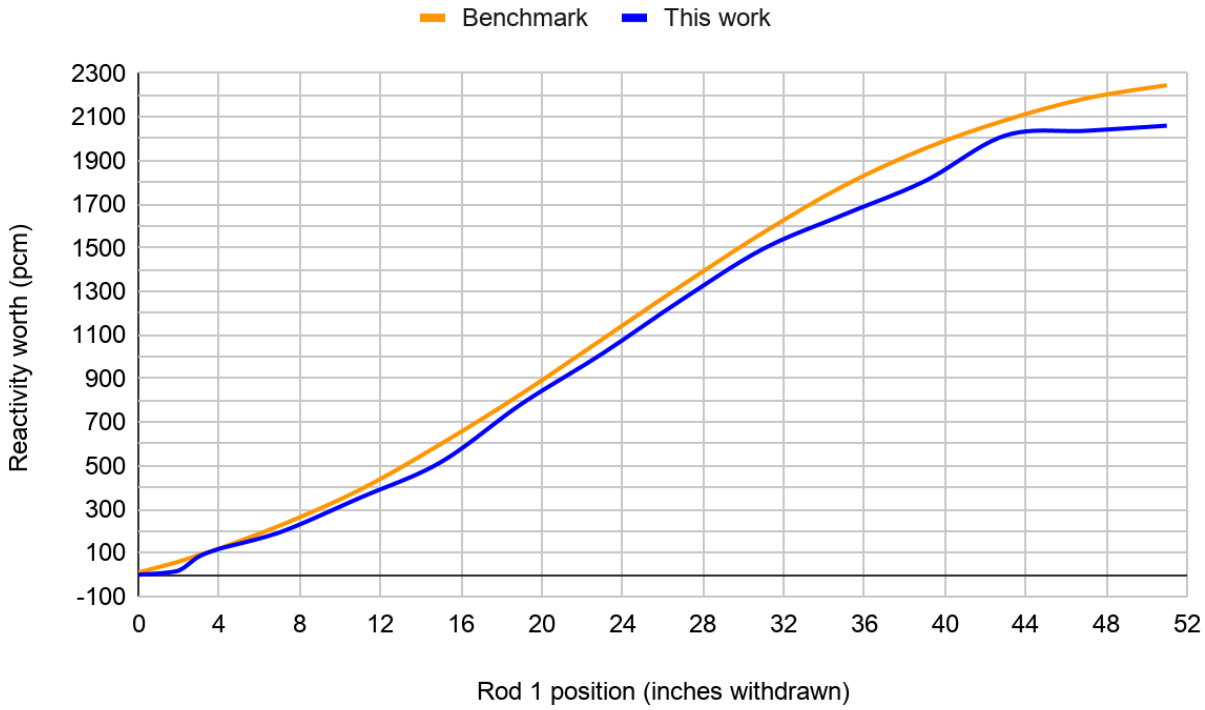


Figure 4 - Integral rod worth comparison

The calculated total Rod-1 worth over the full stroke is 2059.6 pcm, against the benchmark value of approximately 2200 pcm (C/E = 0.94).

The full-stroke deficit relative to the benchmark is 140 pcm, or approximately 6% of the total worth. This deficit is not distributed uniformly along the stroke. As Figure 4 shows, the two integral curves track one another closely from full insertion to roughly 36 inches withdrawn, and the separation opens progressively above that point; the corresponding differential comparison in Figure 5 shows the calculated worth falling below the benchmark most sharply at the extreme withdrawn positions. Essentially the whole of the integral deficit therefore accumulates in the upper portion of the stroke rather than being incurred evenly across it.

#### 4.2.2.2 Differential Rod Worth


Figure 5 - Differential rod worth comparison

The calculated peak differential worth is 22.64 pcm/cm at 21 in withdrawn, which is the peak for benchmark (C/E = 0.93). Running the experiment with a larger number of particles will reduce the sigma to comparable values as the benchmark (0.2 to 0.7 pcm/cm).

## 4.3 Analysis of $^{235}$U Reactivity Coefficient (K)

The $^{235}$U reactivity coefficient K relates the reactivity change to the relative change in fissile loading. Following the benchmark [1, Eq. 3.2],

$$\Delta\rho = K\,(C - C_0)\,/\,C$$

where C is the $^{235}$U mass in the primary loop, $C_0$ = 65.25 kg is the minimum critical loading, and $\Delta\rho = \rho(C) - \rho(C_0)$ with $\rho = (k - 1)/k$. The independent variable is thus the fractional loading change referred to the current mass, $(C - C_0)/C$. K is dimensionless (%Δk/k per %ΔM/M). The benchmark plots this as $\Delta k/(k_1 k_2)$ against Δm/m [1, Fig. 3.11], which is the same pair of quantities.

The benchmark provides two independent determinations. The experimental value, K = 0.223 ± 0.006, derives from the 1965 capsule-addition campaign: ~85 g $^{235}$U capsules of 73LiF–27$UF_4$ eutectic were added through the sampler-enricher, rod 1 was inserted to restore criticality, and the reactivity equivalent of each addition was obtained as the product of the rod displacement and the

calibrated differential worth, $\Delta\rho = \Delta z \cdot dW/dz(\bar{z})$. The experimental K therefore inherits the results of the rod-calibration experiment. The calculated value in the benchmark, K = 0.2228 ± 0.0014 (C/E = 0.999 ± 0.029), was instead obtained from a fixed-rod eigenvalue sweep — k_eff computed at each loading with all rods fully withdrawn, and K extracted as the regression slope of ρ against $\Delta M/M_0$ — and is thus independent of the rod-worth model.

The present work follows the eigenvalue-sweep method. Eigenvalue calculations were performed at the eight cumulative loadings of the capsule-addition campaign (65.25 to 72.00 kg $^{235}U$ (Table 3.9 of [1]), with all three rods fully withdrawn, and K was obtained by linear regression of ρ on ΔM/M. Because the capsules added LiF–$UF_4$ eutectic rather than pure uranium, each loading step increases not only the uranium content but also the lithium and fluorine inventories and the salt density (2.3223 to 2.3280 g/cm³ at the fixed loop volume of 1.9936 m³). The per-case salt compositions and densities were therefore taken directly from the benchmark's loading table [1, Table 3.9] rather than obtained by scaling the uranium atom fractions alone.

One important mechanism used in the present work is how the salt loadings for each case was simulated. Since there is geometry difference across experiment, benchmark and this work (mostly due to external loop and associated equipment missing in simulations), the absolute masses cannot be simulated as-is. Instead, the wt% of each component of the salt was maintained in the simulation. This is described in Table 8 in the Inputs section below.

### 4.3.1 Inputs

TABLE 7 - EXPERIMENT CONDITIONS

| PARAMETER | VALUE |
|---|---|
| Temperature (K) | 922.15 |
| Rod 1, 2, 3 position | 51" withdrawn |
| Thermal expansion | Not applied |
| Sweep variable | Salt composition - corresponding to each U-235 loading (see Table 8) |

TABLE 8 - SALT COMPOSITION SETTINGS TO SIMULATE CONDITIONS IN WT% INSTEAD OF EXACT KGS SPECIFIED IN BENCHMARK

| COMPONENT | CASE 1 (65.25 KG $^{235}$U) — KG | CASE 1 — WT% APPLIED | CASE 8 (72.00 KG $^{235}$U) — KG | CASE 8 — APPLIED | CHANGE ACROSS SWEEP |
|---|---|---|---|---|---|
| $^{234}$U | 0.67 | 0.0145 | 0.73 | 0.0157 | +8.96% by mass |
| $^{235}$U | 65.25 | 1.4093 | 72.00 | 1.5512 | +10.34% by mass |
| $^{236}$U | 0.27 | 0.0058 | 0.30 | 0.0065 | +11.11% by mass |
| $^{238}$U | 141.91 | 3.0651 | 142.32 | 3.0663 | +0.29% by mass |
| Li | 507.27 | 10.9566 | 507.85 | 10.9415 | +0.11% by mass |
| Be | 293.96 | 6.3493 | 293.96 | 6.3333 | +0.00% by mass |
| Zr | 513.97 | 11.1013 | 513.97 | 11.0734 | +0.00% by mass |
| F | 3103.22 | 67.0271 | 3107.08 | 66.9413 | +0.12% by mass |
| **Total salt** | 4629.80 | 100.0000 | 4641.50 | 100.0000 | +0.25% by mass |
| Salt density (g/cm³) | 2.32233 | — | 2.32820 | — | +0.25% |

Note: Only main nuclides mentioned in this table; impurities are omitted but set correctly in the simulation

### 4.3.2 Results

The resulting coefficient from calculation is K =0.2303 ± 0.0040, close to the benchmark's calculated value 0.2228 ± 0.0014, a C/E of 1.033.

TABLE 10 - INPUTS AND OUTPUTS FOR EVERY CASE

| CASE | U-235 (KG) | LOOP SALT (KG) | SALT DENSITY (G/CM³) | K_EFF | Σ_K | P (PCM) |
|---|---|---|---|---|---|---|
| 1 | 65.25 | 4629.80 | 2.32233 | 1.00994 | 0.00036 | 984.22 |
| 2 | 66.00 | 4631.10 | 2.32298 | 1.01410 | 0.00037 | 1390.40 |
| 3 | 67.00 | 4632.84 | 2.32386 | 1.01674 | 0.00042 | 1646.44 |
| 4 | 68.00 | 4634.57 | 2.32472 | 1.01964 | 0.00035 | 1926.17 |
| 5 | 69.00 | 4636.30 | 2.32559 | 1.02353 | 0.00037 | 2298.91 |
| 6 | 70.00 | 4638.03 | 2.32646 | 1.02645 | 0.00038 | 2576.84 |
| 7 | 71.00 | 4639.77 | 2.32733 | 1.02975 | 0.00037 | 2889.05 |
| 8 | 72.00 | 4641.50 | 2.32820 | 1.03330 | 0.00038 | 3222.68 |

TABLE 11 - CALCULATED RESULTS

| QUANTITY | VALUE |
|---|---|
| K (dimensionless) | 0.2303 ± 0.0040 |
| Fit method | weighted least squares, $w = 1/\sigma_\rho^2$, $\sigma_\rho = \sigma_k/k^2$ |
| C/E | 1.033 vs experiment (0.223). C/C is 1.034 vs benchmark (0.2228) |

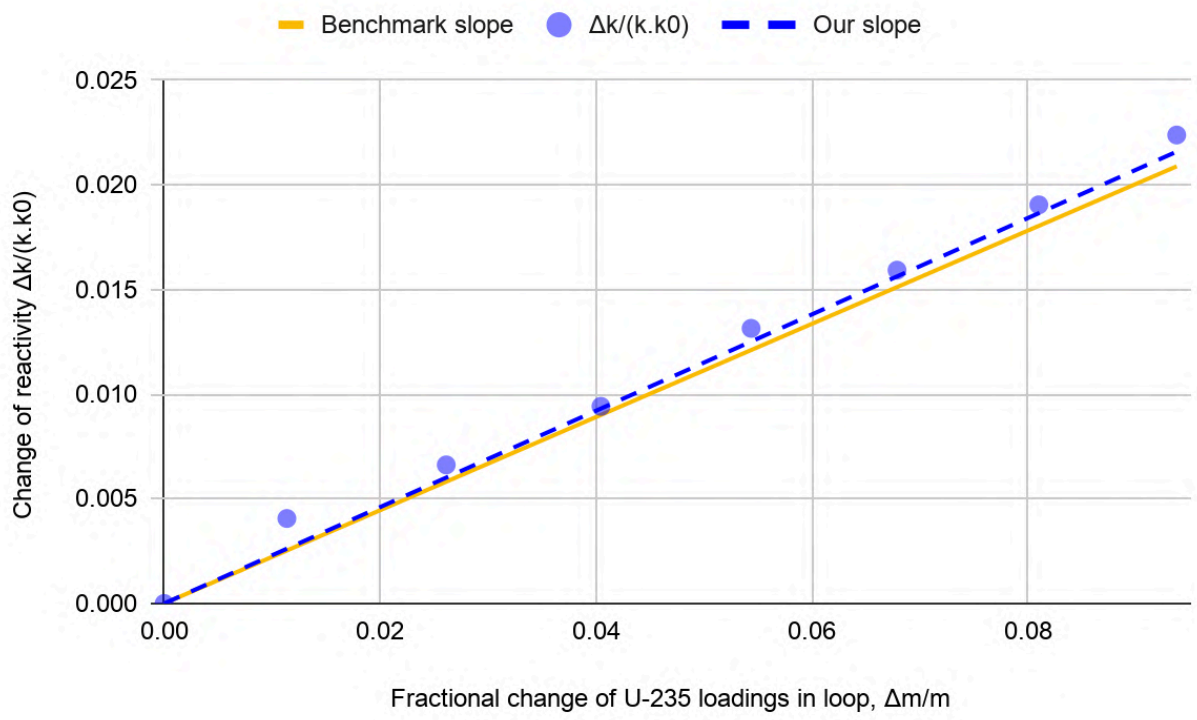


Figure 6 - Effect of $^{235}U$ mass on reactivity

# 5. Discussion

The three experiments analysed in this work probe two largely independent parts of the model. Initial criticality and the $^{235}U$ reactivity coefficient test the fuel-salt composition, density and thermal treatment together with the bulk core geometry; the integral and differential rod-worth sweep additionally tests the control-rod and absorber-bushing representation. Because the $^{235}U$ coefficient was obtained with all three rods fully withdrawn, it is by construction insensitive to the rod model, and the two families of result can therefore be interpreted separately.

## Eigenvalue bias at first criticality

The calculated k_eff of 1.01172 ± 0.00035 corresponds to +1158 pcm of excess reactivity relative to the critical experiment, and to approximately +1180 pcm relative to the benchmark model expectation of 0.99978 ± 0.00420. Since the benchmark's own model uncertainty is ±420 pcm (1σ), the present result sits roughly 2.8σ above the benchmark expectation, whereas the reference Serpent 2 / ENDF/B-VII.1 calculation sits approximately 5.0σ above it. The bias in this work is therefore about 45% smaller than that of the reference calculation, though it remains a systematic over-prediction rather than agreement.

Several mechanisms plausibly contribute, and the present dataset does not separate them. First, the change of nuclear-data library from ENDF/B-VII.1 to ENDF/B-VIII.1 alters both the $^{235}$U thermal fission and capture evaluations and the graphite S(α,β) treatment; in a thermal, graphite-moderated system these revisions are known to move k_eff by hundreds of pcm, and they account for part of the difference between this work and the reference calculation, since code and library were changed together. Second, and specific to the present geometry, the modelled vessel holds less salt than the benchmark's reported loop inventory, because salt resident in the external piping and drain tank is not represented (Section 3). The loading was anchored by preserving salt mass fractions rather than absolute masses, which correctly preserves the fissile-to-moderator ratio inside the active core but removes an out-of-core region that acts as a parasitic absorber and leakage path; omitting it biases k_eff upward. Quantifying this would require a dedicated calculation in which the out-of-core inventory is represented as a homogenised external volume, which was outside the scope of this work. Third, the graphite boron impurity is modelled at the nominal 0.80 ppm of the benchmark specification; because this is a strong thermal absorber distributed throughout the moderator, modest uncertainty in the as-built impurity level maps onto a reactivity effect of the same order as the residual bias. Finally, the temperature treatment interpolates between tabulated library temperatures, and the first-criticality condition of 638 °C (911 K) lies just above the 900 K tabulation point; the associated interpolation error is expected to be small but is not zero.

## Control-rod worth

The rod-worth sweep reproduces the benchmark behaviour well. The full-stroke integral worth of rod 1 is 2059.6 pcm against a benchmark value of approximately 2200 pcm (C/E = 0.94), and the peak differential worth is 22.64 pcm/cm against a benchmark peak of approximately 24.4 pcm/cm (C/E = 0.93). More informative than the endpoint agreement is where the two curves separate. Figures 4 and 5 show that the calculated differential worth tracks the benchmark closely from full insertion to roughly 36 inches withdrawn, including the position of the peak near 21–24 inches, and that essentially the entire 140 pcm integral deficit accumulates in the upper third of the stroke, where the calculated differential worth falls below the benchmark by a factor approaching three at the extreme withdrawn positions.

This distribution is diagnostic. Correct reproduction of the peak position and of the shape over the central stroke indicates that the axial flux weighting along the rod path, and hence the placement of the rod within the core, is modelled correctly. A deficit confined to the withdrawn end is instead consistent with the modelled poison column terminating at a lower axial elevation than the benchmark specification, so that near full withdrawal the modelled absorber has already left the region of appreciable flux while the benchmark rod retains residual worth. This is an observation by the reference paper [1] also, that even in the fully withdrawn position, the rods are few inches inside the graphite stringer region.

Two caveats limit how far this comparison can be pushed. The benchmark rod-worth values used here were digitised from the figures of [1] and carry an unquantified reading error, and it is not stated in [1] whether those values are experimental, calculated, or experimental values adjusted to calculation. A C/E of 0.94 should accordingly be read as agreement at roughly the level at which the reference values themselves are known.

## $^{235}$U reactivity coefficient

The regression of ρ on ΔM/M over the eight cumulative loadings yields K = 0.2303 ± 0.0040, against the experimental value of 0.223 ± 0.006 and the benchmark's calculated value of 0.2228 ± 0.0014. The difference from the experiment is 0.0073 against a combined uncertainty of 0.0072, i.e. agreement at approximately 1σ; the difference from the benchmark calculation is somewhat larger relative to its tighter uncertainty, at approximately 1.8σ. The agreement is meaningful precisely because of how the loading sweep was constructed: since the salt composition was specified in mass fractions rather than absolute masses (Table 8), and K is a ratio of a relative reactivity change to a relative mass change, the coefficient is insensitive to the absolute salt-inventory mismatch that affects the eigenvalue comparison of Section 4.1. The fact that a quantity constructed to be independent of both the rod model and the absolute inventory reproduces experiment to within 1σ supports the conclusion that the fuel-salt composition, density and temperature treatment are sound, and that the residual eigenvalue bias should be sought in the inventory and impurity treatment rather than in the salt model itself.

The residuals of the fit are not entirely structureless. The lowest-loading point lies visibly above the fitted line in Figure 6, by roughly twice its statistical uncertainty, and the highest-loading point does likewise, suggesting either mild curvature in ρ against ΔM/M over this range or an underestimate of the per-case statistical uncertainty. With per-case σ_k of 3.5–4.2 × $10^{-4}$, the reactivity differences being regressed are only one to two orders of magnitude above the noise floor at the lower loadings. Repeating the sweep with a larger particle count would establish whether the structure is statistical or physical, and would also bring the differential rod-worth uncertainties down to the 0.2–0.7 pcm/cm of the benchmark.

### Limitations

Beyond those already noted, two treatments in the present model are worth recording. Thermal expansion was applied for the first-criticality case but not for the rod-worth sweep or the loading sweep, where densities were taken directly from the benchmark tables; this is defensible but means the three experiments are not modelled on a fully uniform basis. Second, the salt is stationary in all cases, so no circulation effect on delayed-neutron precursor distribution is represented — appropriate for the zero-power static comparisons made here, but a limitation on extending the model to the circulating-fuel experiments of the benchmark suite.

## 6. Conclusions

An OpenMC/DAGMC model of the MSRE built on CAD-derived geometry and ENDF/B-VIII.1 nuclear data has been validated against three of the 1965 zero-power benchmark experiments: initial criticality, the integral and differential worth of the regulating rod, and the $^{235}U$ reactivity coefficient.

The $^{235}U$ reactivity coefficient is reproduced to within one standard deviation of the experimental value ($K = 0.2303 \pm 0.0040$ against $0.223 \pm 0.006$, C/E = 1.033). Rod worth is reproduced to within approximately 6–7% (integral C/E = 0.94, differential peak C/E = 0.93), with the peak position and the shape of the differential curve over the central stroke matched closely and the deficit confined to the withdrawn end of the stroke — a signature consistent with the modelled poison column having insufficient axial extent rather than with an error in absorber cross sections or rod placement. Initial criticality is over-predicted by +1158 pcm, roughly $2.8\sigma$ of the benchmark model uncertainty and about 45% less biased than the benchmark's own reference Serpent 2 / ENDF/B-VII.1 calculation, with the residual attributable to some combination of the library change, the salt inventory held outside the modelled vessel volume, and graphite impurity uncertainty.

Taken together, these results support the suitability of OpenMC with DAGMC CAD geometry for MSR criticality and reactivity-coefficient analysis, while indicating that the axial representation of control-rod absorber regions warrants explicit attention in calculations where rod worth is a licensing-relevant quantity. Three extensions follow directly: a rod-worth calculation using explicit, non-homogenised CSG geometry for the $Gd_2O_3$–$Al_2O_3$ bushing rings, to separate homogenisation error from absorber-extent error; an eigenvalue calculation with the out-of-core salt inventory represented, to quantify its contribution to the criticality bias; and repetition of the sweeps at higher particle counts to bring statistical uncertainties to benchmark level. Extension of the model to the remaining zero-power experiments — rod-bank worth and shadowing, rod-drop reactivity, fuel-circulation experiment, and the isothermal temperature coefficient — would complete the validation against the benchmark suite.

# 7. References


[1] Fratoni, M., Shen, D., Ilas, G., Powers, J. (2020). Molten Salt Reactor Experiment Benchmark Evaluation. Project 16-10240, University of California, Berkeley. OSTI ID 1617123.

[2] Shen, D., Ilas, G., Powers, J. J., & Fratoni, M. (2021). Reactor Physics Benchmark of the First Criticality in the Molten Salt Reactor Experiment. Nuclear Science and Engineering, 195(8), 825–837. https://doi.org/10.1080/00295639.2021.1880850

[3] Romano, P.K., Forget, B. (2013). The OpenMC Monte Carlo Particle Transport Code. Annals of Nuclear Energy, 51, 274-281.

[4] Wilson, P.P.H., Tautges, T.J., Kraftcheck, J.A., Smith, B.M., Henderson, D.L. (2010). Acceleration techniques for the direct use of CAD-based geometry in fusion neutronics analysis. Fusion Engineering and Design.

[5] Nobre, G.P.A., et al. ENDF/B-VIII.1: Updated Nuclear Reaction Data Library for Science and Applications. Nuclear Data Sheets. arXiv:2511.03564. Library available at https://doi.org/10.11578/endf/2571019

[6] Chadwick, M.B., Herman, M., Obložinský, P., et al. (2011). ENDF/B-VII.1 Nuclear Data for Science and Technology. Nuclear Data Sheets, 112, 2887-2996.

[7] Shibata, K., et al. (2011). JENDL-4.0: A New Library for Nuclear Science and Engineering. J. Nucl. Sci. Technol., 48, 1-30.

[8] Leppanen, J., Pusa, M., Viitanen, T., Valtavirta, V., Kaltiaisenaho, T. (2015). The Serpent Monte Carlo Code: Status, Development and Applications in 2013. Annals of Nuclear Energy, 82, 142-150.

[9] Bostelmann, F., Skutnik, S.E., Walker, E.D., Ilas, G., Wieselquist, W.A. (2022). Modelling of the Molten Salt Reactor Experiment with SCALE. Nuclear Technology, 208(4), 603-624.

[10] OpenMSR (Copenhagen Atomics). MSRE: detailed CAD model and OpenMC benchmarks of the Molten Salt Reactor Experiment. GitHub repository. https://github.com/openmsr/msre

[11] Yilmaz S, Romano PK, Chierici L, Knudsen EB and Shriwise PC (2024) CAD and constructive solid geometry modeling of the Molten Salt Reactor Experiment with OpenMC. Front. Nucl. Eng. 3:1385478. doi: 10.3389/fnuen.2024.1385478

[12] Prince, B. E., et al. "ZERO-POWER PHYSICS EXPERIMENTS ON THE MOLTEN-SALT REACTOR EXPERIMENT.." , Jan. 1968. https://doi.org/10.2172/4558029

[13] Barlow, John & Clarno, Kevin. (2020). Validation of SCALE for Molten Salt Reactors using the MSRE Benchmark. 1467-1470. 10.13182/T130-33057.

[14] Kevin T. Clarno, John E. Barlow, Tucker Sawyer, Jonathan Scherr, Jason Hearne, Pavel Tsvetkov. Evaluation of SCALE, Serpent, and MCNP for Molten Salt Reactor applications using the MSRE Benchmark. Annals of Nuclear Energy, Volume 194, 2023, 110092, ISSN 0306-4549, https://doi.org/10.1016/j.anucene.2023.110092

PolyEnergetics Private Limited
www.polyenergetics.com

VERSION 1.0