Convergence and Robustness of Unpreconditioned Krylov Solvers for the Hybrid Finite Element Saddle-Point Form of Neutron Transport

Kyung Min Kim[1], Han Gyu Lee[2]

Korean Nuclear Society Autumn Meeting · October 29, 2026, 10:00
Room 102, Hwabaek International Convention Center, Gyeongju

[1]Korea Atomic Energy Research Institute [2]University of Michigan

Korea Atomic Energy Research Institute

Core Analysis Across Reactor Types

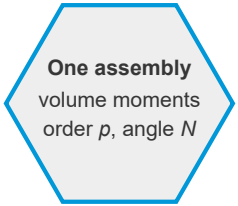
Pre-application design review opens on 20 Nov 2026 (Nuclear Safety Act, Art. 100-2).

Type[a]	Issue	Gap
LWR SMR i-SMR	Rodded depletion, nodal vs. lattice practice: rod-history cross sections	1,300 pcm[b]
SFR Natrium, H-SFR	Sodium void, diffusion vs. transport practice: nodal transport	310–400 pcm, sign flips[c]
MSR seaMSR-100, MARINA	k_{eff}, diffusion vs. Monte Carlo practice: discontinuity factors	+1,100 pcm[d]
HTGR HECTAR, BeSMART	k_{eff}, homogenized diffusion vs. Monte Carlo practice: SPH equivalence	1,700–2,200 pcm[e]
Heat pipe NuSpacePower	Drum worth, deterministic vs. Monte Carlo practice: SPH, drum-aligned meshes	1–20 %[f]

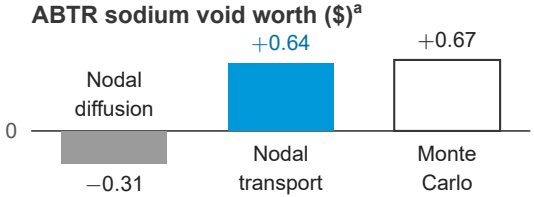
[a]NSSC SMR Roadmap (2026) [b]UNIST, KNS 2024 [c]ANL-AFCI-186 [d]RPHA 2023 [e]INL/EXT-18-51317 [f]ANE 2022; ANL/NEAMS-24/3; ANL/NSE-25/82

High-Order Nodal Transport

- **p-refinement:** raise the angular order (diffusion, SP_N, P_N) and the polynomial order in the same coarse cell, with no transverse-leakage approximation
- **Variational:** the same formulation on Cartesian, hexagonal, and polygonal cells
- **Lineage:** DIF3D-VARIANT (ANL), ERANOS (CEA), INSTANT and Griffin (INL)



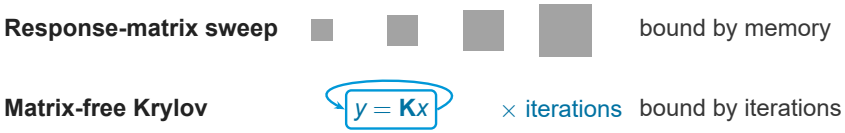
faces: interface multipliers



[a] Same homogenized-assembly model, ENDF/B-V, 33 groups; Yang et al., ANL-AFCI-186 (2007)

Cost of Raising the Order

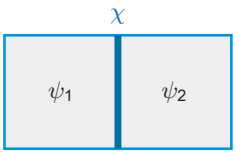
- **Per-cell matrices** grow with angular order × spatial order × energy groups
- **Response-matrix sweep** forms and inverts them: memory is “the remaining inhibitor to using higher order approximations”[a]
- **Matrix-free Krylov** solves the saddle point without response matrices[b], so the cost moves to the iteration count



Question: which Krylov solver suits the HFEM saddle point, and how does its iteration count grow with mesh, order, and scattering?

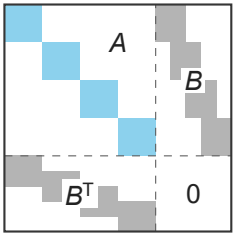
[a] Smith, Lewis, and Shemon, DIF3D-VARIANT 12.0, ANL/NE-14/1 Rev. 1 (2025), p. 40 [b] Wang, Rabiti, and Palmiotti, M&C 2011

HFEM Saddle-Point System



ψ: volume moments, order p_{int}

χ: interface multipliers, order p_{surf}



Eliminate ψ cell by cell

$S\chi = B^{\mathsf{T}}A^{-1}q, \quad S = B^{\mathsf{T}}A^{-1}B$

response matrices, RB sweep (classical)

Keep K, multiply only

$y_\psi = A\psi + B\chi, \quad y_\chi = B^{\mathsf{T}}\psi$

Krylov (this work)

$A\psi$ cell by cell; $B\chi$, $B^{\mathsf{T}}\psi$ face by face

$$\mathbf{K}\begin{bmatrix}\psi\\ \chi\end{bmatrix} = \begin{bmatrix}A & B\\ B^{\mathsf{T}} & 0\end{bmatrix}\begin{bmatrix}\psi\\ \chi\end{bmatrix} = \begin{bmatrix}q\\ 0\end{bmatrix}$$

***A* and *S*:** symmetric positive definite (SPD)

K: symmetric indefinite, with eigenvalues of both signs, so CG does not apply

Three Krylov Solvers

All three need only $y = \mathbf{K}x$, differing in what they assume and what they store.

Solver	Built for	Memory	Residual	Preconditioner
GMRES(m)	any nonsingular	grows to m, restarts	monotone	any
BiCGSTAB	any nonsingular	fixed, small	non-monotone	any
MINRES	symmetric indefinite	fixed, small	monotone	**SPD only**

Baseline: the response-matrix RB sweep, a stationary solve on the face partial currents

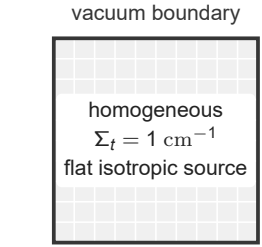
Earlier work[a]: GMRES on **K** with a low-order RB-sweep preconditioner; CG on S

Note for MINRES: the nonsymmetric RB sweep cannot serve as its preconditioner

[a] Wang, Rabiti, and Palmiotti, M&C 2011

A Controlled Test Problem

One-group homogeneous box, SP_1 (diffusion), fixed source, no preconditioning



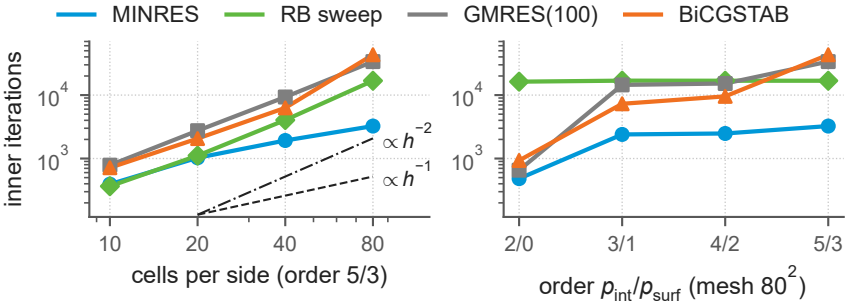
Varied	Values
Order $p_{\mathrm{int}}/p_{\mathrm{surf}}$	2/0, 3/1, 4/2, 5/3
Cells per side	10, 20, 40, 80
Scattering ratio c	0.9 and 0.9999 (diffusion limit)
Solver	MINRES, RB sweep, GMRES(100), BiCGSTAB

128 runs, each stopped at true residual $\|r_k\|/\|r_0\| < 10^{-7}$ or 5×10^4 iterations

Same residual for all four: RB-sweep partial currents are converted to net currents

Mesh and Order Dependence

$c = 0.9$: all four solvers converge in all 16 order–mesh combinations



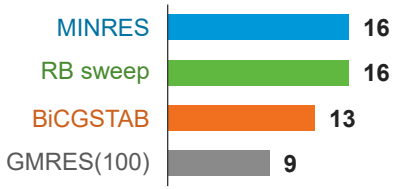
Mesh: MINRES grows as h^{-1} (slope 1.00); the others as h^{-2} (1.80 to 1.94)

Order: Krylov counts rise with the order; the RB sweep stays flat, paying in memory

Robustness in the Diffusion Limit

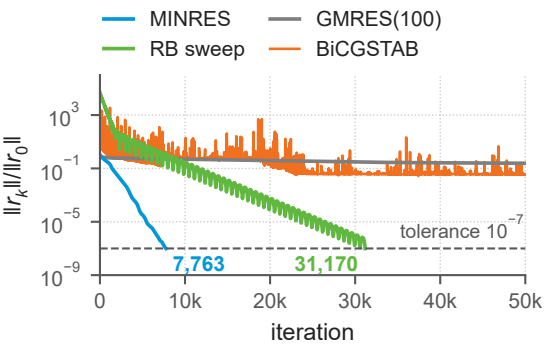
$c = 0.9999$: absorption nearly vanishes, and the conditioning is at its worst

Converged runs (of 16)



mesh slope: MINRES 1.22, RB 1.78

Hardest case: mesh 80^2, order 5/3



Only MINRES and RB converge in all 16; MINRES takes a quarter of the RB iterations

Summary and Next Steps

1. **MINRES** converged in all 32 runs, growing only as h^{-1}
2. **RB sweep** converged in all 32 but grew as h^{-2}; flat in order, at the cost of memory
3. **GMRES(100) and BiCGSTAB** failed in the diffusion limit

One high-order nodal transport for all five reactor types needs three things:

	Obstacle	Next step
Memory	preconditioners need the cell inverses again, growing with angle × space × groups	factored cell inverses (in preparation)
Iterations today	MINRES still grows as h^{-1}; the RB sweep is not SPD	SPD preconditioner for the saddle point
Angle	SP_N does not converge to transport in multiple dimensions	P_N in the same saddle-point form

Thank You

A1. Cost per Iteration

Unpreconditioned, per iteration; hardest case: mesh 80^2, order 5/3, $c = 0.9999$

Solver	Kx	Dot	axpy	Mflop / iter.	Iterations	Total
MINRES	1	2	5	16.8 (1)	7,763	130 Gflop
BiCGSTAB	2	4	6	32 (1.9)	$> 5 \times 10^4$	$>$ 1.6 Tflop
GMRES(100)	1	$j+1$	$j+1$	52 (3.1)	$> 5 \times 10^4$	$>$ 2.6 Tflop
RB sweep[a]	–	–	–	3.3 (0.2)	31,170	102 Gflop

K$x = 2N_c(n_v^2 + 2n_v n_f) = 14.2$, dot or axpy $= 2n = 0.37$, RB sweep $= 2N_c n_f^2 = 3.3$ Mflop

$N_c = 6{,}400$, $n_v = 21$, $n_f = 16$, $n = 186{,}240$. GMRES: step j of a cycle, averaged. RB stores $8N_c n_f^2$ bytes of response matrices

Note: totals are comparable at 80^2

[a] In place of **K**x, dot, and axpy: one response-matrix product per cell