Transactions of the Korean Nuclear Society Autumn Meeting
Gyeongju, Korea, October 28-30, 2026

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

Kyung Min Kim[a*], Han Gyu Lee[b]

[a]*Korea Atomic Energy Research Institute, 111 Daedeok-daero 989 beon-gil, Yuseong-gu, Daejeon 34057, Korea*
[b]*University of Michigan, 2355 Bonisteel Blvd, Ann Arbor, MI, United States*
**Corresponding author: kmkim3182@kaeri.re.kr*

***Keywords** *: hybrid finite element, saddle-point system, Krylov subspace methods, MINRES, matrix-free solvers*

1. Introduction

The Hybrid Finite Element Method (HFEM) [1], also known as the Variational Nodal Method (VNM) [2], is a widely used spatial discretization for deterministic neutron transport in which volume flux moments are coupled to face unknowns through a variational interface condition. The discrete problem is classically cast in response matrix form and solved by red–black (RB) Gauss–Seidel (GS) sweeps, as in the VARIANT [3] and Griffin [4] codes. Such response matrix sweeps are fast and numerically robust, and they have been further accelerated by partitioned-matrix [5] and p-multigrid techniques [6].

The response matrix formulation, however, requires forming and inverting the element response matrices. The size of these local matrices grows rapidly with angular and spatial approximation order, the number of energy groups, and the mesh granularity, so that the memory required to store and invert them grows multiplicatively in the angular order, spatial order, and energy-group count for high-fidelity problems. This memory growth, rather than the floating-point cost, is the principal obstacle to applying the response matrix sweep at high angular and spatial resolution.

Krylov subspace methods have therefore been proposed as an alternative inner solver for the HFEM system. In one line of work the Krylov iteration is applied to the condensed, non-symmetric response-matrix equations and accelerated by p-multigrid preconditioning [6], so that the response matrices are still assembled and the memory bottleneck noted above is not relieved. Closer to the present approach, Wang et al. [7] solve the HFEM saddle-point system directly with a matrix-free Generalized Minimal Residual method (GMRES) preconditioned by a low-order RB sweep, using only the action of the operator and so avoiding assembly of the full response matrices. That study already observed that the bare iteration converges slowly and that preconditioning is in practice required. However, the full saddle-point system was solved only with GMRES, and Conjugate Gradient (CG), another Krylov option, was applied to the Schur complement system since CG mandates a symmetric positive definite (SPD) operator. Building on it, the present work returns to the unpreconditioned saddle-point system and characterizes systematically how the bare operator scales with spatial order, mesh, and scattering. Particular attention is paid to Minimal Residual Method (MINRES) [8], the Krylov method designed specifically for symmetric indefinite systems and therefore the natural and theoretically most robust solver for this saddle-point, which the prior work did not consider.

The solvers studied here are implemented in Neutronics using Deterministic Finite Element Algorithms (NuDEAL), a deterministic finite-element neutronics code [9]. Its HFEM module implements the general SP_N discretization for arbitrary angular order N, which has been elevated to an independent neutronics solver from the low-order coarse-mesh component coupled to planar Method of Characteristics (MOC) [10]. The higher-order angular modes are already available and the diffusion configuration used in this study is simply its lowest-order instance. The code is developed CPU-first and offloaded to the GPU through a backend-parameterized formulation, so that identical operators execute on either backend and the implementation also serves as a testbed for heterogeneous computing. It was originally built around the k_{eff} eigenvalue problem; for the present work a dedicated fixed-source mode was added so that the convergence and robustness of the inner saddle-point solver can be examined in isolation from the outer power iteration.

This paper reports a preliminary evaluation of *unpreconditioned* GMRES, BiConjugate Gradient Stabilized (BiCGSTAB), and MINRES applied directly to the HFEM saddle-point system, with the classical RB sweep retained as an unaccelerated baseline. To isolate the algebraic behavior of the inner solve from angular effects, the evaluation is carried out in the diffusion limit, which serves as a controlled testbed in which the dependence on spatial order, mesh refinement, and scattering ratio can be separated cleanly. We characterize how the inner iteration count grows with p- and h-refinements, and the scattering ratio, and we contrast the convergence behavior and robustness of the four solvers. The aim of this work is to quantify the unpreconditioned baseline that motivates a dedicated preconditioner for this operator.

2. HFEM-SP$_N$ Saddle-Point Formulation

2.1 Multigroup SP$_N$ equations

This work considers the steady-state multigroup SP$_N$ approximation of the neutron transport equation. For energy group g the even-parity moments of the angular flux satisfy a coupled set of second-order, diffusion-like equations whose within-group form may be written compactly as

$$-\nabla\cdot\left(\mathbf{D}_g\,\nabla\phi_g\right)+\Sigma_{r,g}\,\phi_g \;=\; \mathbf{q}_g, \tag{1}$$

where the vector ϕ_g collects the $n_a=(N+1)/2$ even-parity flux moments at each spatial point and $\mathbf{q}_g$ is the corresponding group source comprising fission, in-scatter, and down/up-scatter contributions. Here $\mathbf{D}_g$ and $\Sigma_{r,g}$ are $n_a\times n_a$ coefficient matrices assembled from the odd- and even-order cross sections. Off-diagonal entries of $\mathbf{D}_g$ couple the angular moments, and this inter-moment coupling is precisely what distinguishes SP$_N$ from a single diffusion equation. For $N=1$ ($n_a=1$) the matrices reduce to scalars and (1) becomes the diffusion equation, whereas $N>1$ retains higher-order angular detail through the coupling. While the solver and its operator are formulated for general SP$_N$, the convergence study reported here is conducted entirely in the SP$_1$ configuration, so that the angular dimension is removed and the behavior of the inner solve can be attributed to spatial order, mesh, and scattering alone. The fission source and the eigenvalue k_{eff} are updated by an outer fission source iteration, within which each group is a fixed source problem of the form (1).

2.2 Hybrid finite element discretization

The HFEM discretizes (1) on a mesh of cells with discontinuous, cell-local polynomial expansions of order p_{int} for the volume moments ψ, while continuity of the relevant trace across cell interfaces is imposed weakly through face multipliers χ expanded to order p_{surf}. Enforcing the interface condition with χ acting as a Lagrange multiplier yields, for each group, a symmetric indefinite saddle-point system

$$\mathbf{K}\begin{bmatrix}\psi\\ \chi\end{bmatrix}=\begin{bmatrix}q\\ 0\end{bmatrix},\qquad \mathbf{K}=\begin{bmatrix}A & B\\ B^{\mathsf{T}} & 0\end{bmatrix}, \tag{2}$$

where the group index has been suppressed. The leading block A combines the streaming operator and the collision operator; it is symmetric positive definite and block diagonal over cells, so that A and A^{-1} act through independent element-local dense solves. The rectangular block B is the primal trace operator coupling the volume moments to the face multipliers, and the zero $(2,2)$ block reflects the absence of a face-local self term. Eliminating $\psi=A^{-1}(q-B\chi)$ gives the surface Schur complement

$$S\chi=B^{\mathsf{T}}A^{-1}q,\qquad S=B^{\mathsf{T}}A^{-1}B, \tag{3}$$

which is SPD under the usual HFEM assumptions.

2.3 Matrix-free operator application

The Krylov solvers require only the saddle-point matrix–vector product $y=\mathbf{K}x$, which for $x=[\psi;\chi]$ is

$$y_\psi=A\,\psi+B\,\chi,\qquad y_\chi=B^{\mathsf{T}}\psi. \tag{4}$$

The decisive practical advantage of this form is that neither the global operator $\mathbf{K}$ nor the dense Schur complement S is ever assembled or inverted, which removes the rapidly growing response matrix storage that limits the classical sweep. Because A is block diagonal, $A\psi$ is a batch of small, independent cell-local products, while B and B^{T} reduce to face-local trace couplings. Because no global matrix is formed, the inner solve carries a small, fixed memory footprint and is amenable to execution on the GPU backend of NuDEAL.

3. Unpreconditioned Iterative Solvers

The present study targets the fixed-source form of (2) so as to isolate the inner saddle-point solve; in an eigenvalue calculation the same fixed-source solve recurs within each step of the outer fission source iteration. We compare the three Krylov methods most commonly applied to such systems, run each without preconditioning so as to isolate the intrinsic convergence behavior of the bare saddle-point operator, and contrast them with the classical RB sweep. For all solvers the iteration is monitored by the relative residual $\|r_k\|/\|r_0\|$, where $r_k=b-\mathbf{K}x_k$, with $b=[q;0]$ and $x_k=[\psi_k;\chi_k]$, and r_0 is the residual of the common initial guess. Crucially, this is the *true* residual, recomputed explicitly at every step rather than the recurrence estimate that each method updates internally. Evaluating it identically for all four solvers provides a single, solver-independent stopping criterion, and the iteration is terminated at a fixed reduction tolerance or a maximum iteration count.

3.1 GMRES

GMRES is the most general among the three in a theoretical sense: it minimizes the residual over the full Krylov subspace and therefore exhibits a monotone, nonincreasing residual history for an arbitrary nonsingular operator. Its cost, however, is an Arnoldi basis whose storage and orthogonalization work grow with every iteration, so in practice it is restarted as GMRES(m) after m steps. The restart parameter m trades memory against convergence: a small m caps the memory footprint but can stall the residual, whereas a large m recovers convergence at a correspondingly higher memory load.

3.2 BiCGSTAB

BiCGSTAB uses short recurrences and needs a fixed, small number of work vectors, making it considerably lighter than GMRES in memory. Because it is built on biorthogonality rather than a complete orthogonalization

of the Krylov basis, its residual history is typically non-monotone and can oscillate strongly, and it is susceptible to near-breakdown. It can nonetheless be competitive when it converges smoothly, which makes it a useful point of comparison.

3.3 MINRES

MINRES is tailored to symmetric (indefinite) systems such as (2). It exploits the symmetry of $\mathbf{K}$ through a Lanczos three-term recurrence, giving a constant, minimal memory footprint together with a monotone residual. It is in this sense the natural and most appropriate choice for the HFEM saddle-point operator. Its one structural restriction matters here: MINRES accepts only SPD preconditioning, which rules out the nonsymmetric RB sweeper and instead demands a separately constructed, and more delicate, preconditioner.

3.4 Red–black sweep baseline

For reference the classical response matrix solver is included, in which the volume moments are eliminated cell by cell and the resulting interface system is relaxed by RB GS sweeps. This is the established and numerically robust workhorse of the HFEM, but it is a stationary iteration whose asymptotic convergence factor degrades as the mesh is refined and as the scattering ratio approaches unity. To compare it on the same footing as the Krylov methods, its partial-current iterates are converted to the net-current representation and measured by the identical saddle-point residual.

4. Numerical Results

4.1 Test problem and methodology

The test problem is a homogeneous, one-group, two-dimensional fixed-source box of side 40 cm with a uniform isotropic source and vacuum boundaries, discretized on a base mesh of 10×10 cells, broadly similar to the homogeneous box of Wang et al. [7]. Its single material has $\Sigma_t = 1$ cm^{-1}, so the domain spans 40 mean free paths (mfps) and each base cell is 4 mfps across. Two scattering ratios are used: a moderate operating point $c = 0.9$, at which the solution is smooth and well-conditioned, and a diffusion-limit case $c = 0.9999$, at which the absorption nearly vanishes and the operator approaches its worst conditioning. Because the test problem is homogeneous and solved as a single fixed-source problem, it isolates the dependence of the inner solve on discretization order, mesh, and scattering ratio alone. At the moderate operating point, $c = 0.9$, any growth of the inner iteration count with spatial order therefore reflects the algebraic conditioning of the high-order discretization rather than the classical diffusion-limit hardness, while the $c = 0.9999$ case probes the latter as a stress limit.

Table I: Specification of the test problem.

Property	Value
Dimensionality	2D
Energy groups	1
Materials	homogeneous
Problem type	fixed source
Source profile	spatially flat and isotropic
Domain	40×40 cm^2
Boundary	vacuum
Base mesh	10×10

The three Krylov solvers and the RB sweep are run without preconditioning; all are terminated when the relative residual $\|r_k\|/\|r_0\|$ falls below 10^{-7} or a cap of 5×10^4 iterations is reached, where non-convergence within the cap is itself reported as a valid outcome that quantifies the need for acceleration. GMRES is restarted as GMRES(m) with $m = 100$. Table I summarizes the problem specification.

4.2 Scaling with spatial order and mesh refinement

The central experiment varies the volume/surface order pairs $(p_{\text{int}}, p_{\text{surf}}) = (2,0),(3,1),(4,2),(5,3)$ and refines the mesh through $10^2, 20^2, 40^2, 80^2$ cells, at both the moderate operating point $c = 0.9$ and the diffusion limit $c = 0.9999$. Fig. 1 shows the resulting inner iteration counts for the four solvers, and Table II summarizes the mesh-refinement scaling exponent and the convergence robustness.

Mesh refinement separates the methods sharply. At the highest order, MINRES grows only as h^{-1}, with a fitted exponent $\alpha \approx 1.0$ at $c = 0.9$ and 1.2 at $c = 0.9999$, consistent with the classical Schur-complement conditioning argument [11]. For the underlying elliptic interface operator the condition number grows as $\kappa = O(h^{-2})$, so a stationary iteration needs $O(\kappa) = O(h^{-2})$ steps while an optimal Krylov process is expected to temper this toward $O(\sqrt{\kappa}) = O(h^{-1})$; the observed MINRES trend follows the latter, even though MINRES is applied to the full indefinite system rather than to the Schur complement directly. The RB sweep instead grows as h^{-2} ($\alpha \approx 1.8$) and is essentially independent of the spatial order, because each sweep eliminates the cell-interior moments exactly and only the cell-interface iteration governs convergence. The restarted GMRES and BiCGSTAB enjoy neither advantage: even where they converge they scale like h^{-2} rather than h^{-1}. This is consistent with the restart of GMRES(100) periodically discarding the Krylov history that would otherwise deliver the h^{-1} rate, whereas BiCGSTAB never attains that rate and converges non-monotonically.

The diffusion limit is decisive for robustness. At $c = 0.9$ all four solvers converge within the cap, but at $c = 0.9999$ restarted GMRES fails on the finer, higher-order configurations and BiCGSTAB fails on the hardest, whereas MINRES and the RB sweep converge in every

case. The order dependence in the lower row of Fig. 1 reinforces the same picture: the RB count is flat in p while every Krylov count rises with p, so that MINRES attains the fewest iterations at every order at this finest mesh.

Table II: Mesh-refinement scaling and robustness on the test problem. α is the fitted exponent at order $p_{\text{int}}/p_{\text{surf}} = 5/3$ in (inner iterations) $\propto$ (intervals per side)$^{\alpha}$; "conv." counts the converged cases out of the 16 order–mesh combinations at each scattering ratio; a dash (—) marks too few converged points to fit a slope.

	$c = 0.9$		$c = 0.9999$	
Solver	α	conv.	α	conv.
MINRES	1.00	16/16	1.22	16/16
RB sweep	1.84	16/16	1.78	16/16
GMRES(100)	1.80	16/16	—	9/16
BiCGSTAB	1.94	16/16	2.04	13/16

4.3 Convergence behavior and robustness

Fig. 2 records the residual history of the four solvers on the hardest configuration, mesh 80^2 at order $5/3$ with $c = 0.9999$. The methods display qualitatively distinct behavior. MINRES decreases smoothly and monotonically, reaching the tolerance in about 7.8×10^3 iterations. The RB sweep also converges, but only after a slowly contracting, strongly oscillatory envelope, taking roughly four times as many iterations. BiCGSTAB is strongly non-monotone, with residual spikes several orders of magnitude above its starting value, and never reaches the tolerance within the cap. Restarted GMRES stalls almost immediately near a relative residual of order unity and likewise fails, consistent with the restart at $m = 100$ repeatedly discarding the subspace information needed to make progress. Together with Table II, this indicates that, without preconditioning, MINRES is the only solver that is at once robust and efficient in iteration count, the RB sweep being robust but asymptotically slower and the two remaining Krylov methods unreliable in the diffusion limit. Even MINRES, however, remains a Krylov iteration whose count still grows as $O(h^{-1})$ under refinement, it is far from fast and would itself benefit substantially from preconditioning, which is the subject of the companion study.

5. Conclusions

This paper presented a preliminary evaluation of unpreconditioned Krylov solvers for the HFEM-SP$_N$ saddle-point system, using the SP$_1$ diffusion limit as a controlled testbed that separates the effects of spatial order, mesh refinement, and scattering ratio. GMRES, BiCGSTAB, and MINRES were compared without preconditioning, with the classical RB sweep retained as an unaccelerated baseline.

The results show that, in the absence of preconditioning, the choice of method is decisive. The RB sweep converges in all configurations but only at an $O(h^{-2})$ rate and independently of the spatial order, so its cost grows rapidly under mesh refinement. The response matrix memory footprint would still be prohibitive in realistic core problems, even though it is not analyzed in this work since the test problem is homogeneous. Restarted GMRES and BiCGSTAB are competitive on the easier problems but become unreliable in the diffusion limit, stalling or breaking down before reaching the tolerance on the finer, higher-order cases, and even where they converge they degrade to the same $O(h^{-2})$ rate as the stationary sweep. MINRES is the only solver that converges in every case while retaining the optimal $O(h^{-1})$ Krylov rate, confirming it as the natural method for the symmetric indefinite operator. On the worst configuration, MINRES reached the tolerance in about 7.8×10^3 iterations, whereas GMRES and BiCGSTAB did not converge within the 5×10^4 cap.

Even MINRES, however, is not fast in practical applications, since its iteration count still grows under refinement, so that an effective preconditioner remains essential. The structural obstacle is that MINRES admits only an SPD preconditioner and therefore cannot reuse the established nonsymmetric low-order RB sweeper [7] for acceleration. A purpose-built SPD preconditioner of the kind proposed by [12] is thus required, and its construction and evaluation for the HFEM system are quite different from the conventional RB sweeper. Extending the present diffusion-limit characterization to higher SP$_N$ orders and to heterogeneous multigroup eigenvalue cores is left as future work; because the saddle-point structure is invariant in the angular order, the solver ranking established here is expected to carry over.

References

[1] P. Raviart and J. M. Thomas. "Primal Hybrid Finite Element Methods for 2nd Order Elliptic Equations". In: *Mathematics of Computation* 31.138 (1977), pp. 391–413.

[2] I. Dilber and E. E. Lewis. "Variational Nodal Methods for Neutron Transport". In: *Nuclear Science and Engineering* 91.2 (1985), pp. 132–142.

[3] M. A. Smith, E. E. Lewis, and E. R. Shemon. *DIF3D-VARIANT 11.0: A Decade of Updates*. Tech. rep. ANL/NE-14/1. Argonne National Laboratory, 2014.

[4] Y. S. Jung, Y. Wang, and C. Lee. "Development of Transport Solver Based on Hybrid Finite Element Method in Griffin". In: *Annals of Nuclear Energy* 214 (2025), p. 111128.

[5] Y. Li et al. "Partitioned-Matrix Acceleration to the Fission-Source Iteration of the Variational Nodal Method". In: *Progress in Nuclear Energy* 85 (Nov. 2015), pp. 640–647.

[6] E. E. Lewis et al. "Preconditioned Krylov and Gauss-Seidel Solutions of Response Matrix Equations". In: *International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering*. March. 2011.

[7] Y. Wang, C. Rabiti, and G. Palmiotti. "Krylov Solvers Preconditioned with the Low-Order Red-Black Algorithm for the PN Hybrid FEM for the INSTANT Code". In: *International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering*. 2011.

[8] C. C. Paige and M. A. Saunders. "Solution of Sparse Indefinite Systems of Linear Equations". In: *SIAM Journal on Numerical Analysis* 12.4 (1975), pp. 617–629.

[9] K. M. Kim et al. "Verification and Performance Assessment of NuDEAL, a GPU-Accelerated Deterministic Transport Framework on Unstructured Meshes". In: *Nuclear Science and Engineering* (Apr. 2026), pp. 1–25.

[10] K. M. Kim et al. "2D/3D GPU-accelerated Deterministic Neutronics Solver Based on Planar Method of Characteristics and Hybrid Finite Element Method". In: *Joint International Conference on Supercomputing in Nuclear Applications + Monte Carlo*. Vol. 302. 2024, p. 02001.

[11] M. Benzi, G. H. Golub, and J. Liesen. "Numerical Solution of Saddle Point Problems". In: *Acta Numerica* 14 (May 2005), pp. 1–137.

[12] W. Pazner, T. Kolev, and P. S. Vassilevski. "Matrix-Free High-Performance Saddle-Point Solvers for High-Order Problems in H(**div**)". In: *SIAM Journal on Scientific Computing* 46.3 (2024), B179–B204.

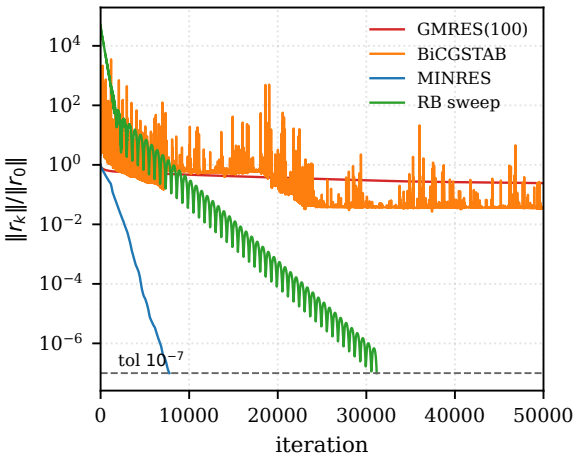


Fig. 2. Relative residual history of the four solvers on the worst configuration (mesh 80^2, order $5/3$, $c = 0.9999$); GMRES and BiCGSTAB do not reach the tolerance within the cap.

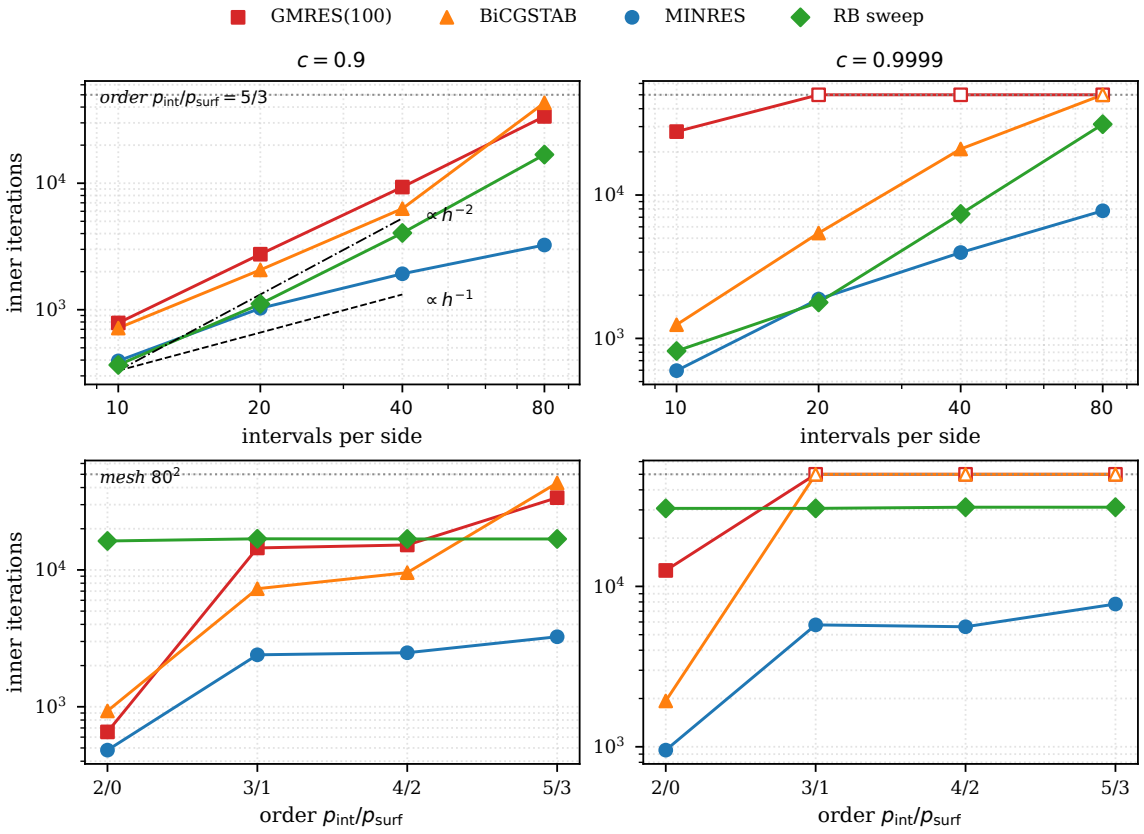


Fig. 1. Inner iteration count on the test problem for the four solvers, at $c = 0.9$ (left) and the diffusion limit $c = 0.9999$ (right). Top: versus mesh refinement at fixed order $p_{\text{int}}/p_{\text{surf}} = 5/3$; bottom: versus spatial order at fixed mesh 80^2. Filled markers denote convergence to $\|r_k\|/\|r_0\| < 10^{-7}$; open markers on the dotted ceiling denote non-convergence within the 5×10^4-iteration cap. Dashed and dash-dotted guides mark the h^{-1} and h^{-2} slopes.