A MULTILAYER NONLINEAR ELIMINATION PRECONDITIONED INEXACT NEWTON METHOD FOR STEADY-STATE INCOMPRESSIBLE FLOW PROBLEMS IN THREE DIMENSIONS*

LI LUO†, XIAO-CHUAN CAI‡, ZHENGZHENG YAN§, LEI XU§, AND DAVID E. KEYES†

Abstract. We develop a multilayer nonlinear elimination preconditioned inexact Newton method for a nonlinear algebraic system of equations, and a target application is the three-dimensional steady-state incompressible Navier–Stokes equations at high Reynolds numbers. Nonlinear steady-state problems are often more difficult to solve than time-dependent problems because the Jacobian matrix is less diagonally dominant, and a good initial guess from the previous time step is not available. For such problems, Newton-like methods may suffer from slow convergence or stagnation even with globalization techniques such as line search. In this paper, we introduce a cascadic multilayer nonlinear elimination approach based on feedback from intermediate solutions to improve the convergence of Newton iteration. Numerical experiments show that the proposed algorithm is superior to the classical inexact Newton method and other single layer nonlinear elimination approaches in terms of the robustness and efficiency. Using the proposed nonlinear preconditioner with a highly parallel domain decomposition framework, we demonstrate that steady solutions of the Navier–Stokes equations with Reynolds numbers as large as 7,500 can be obtained for the lid-driven cavity flow problem in three dimensions without the use of any continuation methods.

Key words. steady-state incompressible flow, lid-driven cavity flow at high Reynolds numbers, inexact Newton, multilayer preconditioning, domain decomposition, parallel processing

AMS subject classifications. 49M15, 65Y05, 65M55, 76D05

DOI. 10.1137/19M1307184

1. Introduction. The Newton method and its variants [17] are frequently used for the numerical solution of large nonlinear systems arising from the discretization of partial differential equations, for example, the incompressible Navier–Stokes equations. Near quadratic convergence can be observed when the nonlinearities in the system are well-balanced and a good initial guess is available. However, if some of the equations have stronger nonlinearities than the others in the system, such methods may suffer from slow convergence in the form of a long stagnation, or not converge at all. These problems usually involve, for example, boundary layers, shock waves, or corner singularities. In these situations, the convergence is rather sensitive to the system parameters, e.g., the Reynolds number, the Mach number, the mesh size, and sometimes the shape of the computational domain. To overcome this difficulty, many continuation techniques have been studied for finding a good initial guess, including parameter continuation [3], mesh sequencing [45], and pseudo-time-stepping [15]. Though they are easy to implement, the drawbacks of continuation are also obvious: it is often difficult to guess the right value of the incremental parameters, or the op-

timal size of an intermediate mesh, or the optimal choice of the time increment for each iteration [34]. Recently, a new class of algorithms, namely, nonlinear preconditioning, draws much attention from the scientific computing community for solving these nonlinearly difficult problems.

Nonlinear preconditioning aims to balance the nonlinearities of the overall system by removing local high nonlinearities that cause the Newton method to take small updates, so that fast convergence can be realized. Similar to a linear preconditioner, a nonlinear preconditioner can be applied on the left or on the right of the nonlinear function. The idea of left preconditioning stems from the additive Schwarz preconditioned inexact Newton method (ASPIN) [13], which changes the original nonlinear system to a more balanced one, and then solve the new system using a Newton-like method. ASPIN and its variants such as the multiplicative Schwarz version (MSPIN) [46, 47, 48] and the restricted additive Schwarz version (RASPEN) [20] have been applied successfully to two-dimensional problems including incompressible flows with high Reynolds numbers [13, 34, 35], transonic compressible flows [36], and multiphase flows in porous media [48]. In contrast, right preconditioning such as nonlinear elimination (NE) [14] does not change the nonlinear function but modifies the unknown variables of the original system, effectively abstracting subsets. The application of NE can be viewed as a subspace correction step to provide a new starting point for the Newton iteration. By using the theory of affine invariance, Gong and Cai [30] provided some insight into why the use of NE can improve the convergence of Newton iterations. In addition, different strategies have been proposed to identify the components of unknown variables that slow down the convergence. Hwang, Lin, and Cai [36] and Hwang, Su, and Cai [37] used a physics-based approach for transonic full potential problems. Huang, Yang, and Cai [33] and Yang and Hwang [62] applied a pointwise approach to eliminate the components associated with certain mesh points that cause the local high nonlinearities for two-dimensional (2D) multicomponent systems. Yang, Sun, and Yang [61] proposed a field-based approach to eliminate the components associated with some field variables, i.e., the saturation field in reservoir simulation. More recently, Luo et al. [49] extended the NE preconditioner to three-dimensional (3D) blood flow problems in a human artery with stenosis.

In this paper, we develop a multilayer NE preconditioned inexact Newton method for steady-state incompressible flow problems, specifically, the challenging lid-driven cavity flow problem in three dimensions at high Reynolds numbers. Nonlinear steady-state problems are often more difficult to solve than time-dependent problems because the Jacobian matrix is less diagonally dominant, and a good initial guess from the previous time step is not available. The Newton-like methods may diverge even when they are used with some globalization techniques such as line search, and linear subiterations may be slow to converge. To overcome the difficulty, we introduce an adaptive multilayer NE approach based on the feedback from the intermediate solutions. We show by numerical experiments that the proposed algorithm is superior to the classical inexact Newton method and other single layer NE approaches in terms of the robustness and efficiency. In our implementation, we embed the nonlinear preconditioning step into an overlapping domain decomposition framework [11, 40, 54] so that the overall method can be used for problems defined on large meshes and on machines with a large number of processor cores.

To evaluate the proposed algorithm we focus on the 3D incompressible lid-driven cavity flow problem with highest Reynolds numbers for which the steady-state solution still exists. The problem has attracted lots of attention for decades because of its relevance to many industrial problems and its importance to the fundamental understand-

ing of fluid dynamics. The simplicity of geometry and boundary conditions make it an ideal benchmark case for validating numerical techniques and comparing results obtained from experiments and computations. Since 2D calculations are significantly less costly than the 3D cases, simulations at high Reynolds numbers of the 2D flow have generated a large number of publications, and the main results are summarized in [28, 55]. Recent studies succeeded in computing steady solutions for 2D driven cavity flows with Re $\geq$ 10,000 [7, 10, 22, 23, 32, 57]. Nonlinear preconditioning techniques have been applied to improve the convergence of Newton type solvers for steady driven cavity flows in two dimensions, and the mathematical model varies from lattice Boltzmann equations [33] to incompressible Navier–Stokes equations in the velocity-pressure formulation [34] or the velocity-vorticity formulation [13, 14, 62]. While these calculation are of fundamental interest for benchmarking, they are less relevant for realistic flows.

Recent progress in parallel computing hardware and software makes it possible to investigate 3D flow problems by solving the incompressible Navier–Stokes equations discretized on fine meshes. The flow structure and topology, as well as some benchmarking results were reported in [2, 5, 16, 19, 24, 25, 26, 27, 29, 38, 42, 43, 56]. The most recent results for the driven cubical cavity flow were reviewed in [44]. In Deville, Lê, and Morchoisne [18], a series of benchmark tests for the 3D lid-driven cavity flow were compared for Re = 3,200. The results, however, remained inconclusive, since the numerical solutions obtained by different methods and resolutions differ significantly. An important view is the fact that endwall effects can, to some degree, suppress the intrinsic 3D flow instabilities in the bulk of the cavity [1, 2]. Based on a linear stability approach, Feldman and Gelfgat [24, 25] and Gelfgat [27] assumed that the steady state of time-dependent flows can be computed accurately for Reynolds numbers below 3,000. In particular, using the laser-Doppler anemometer measurements, Koseff and Street [41] and Prasad and Koseff [52] conducted a series of experiments in a rectangular cavity with different spanwise aspect ratios (SARs) between 0.25:1 and 3:1. According to their observations, the flow for Reynolds numbers below 5,000 is essentially laminar and that transition to the turbulence regime takes place in the range of 6,000–8,000. In the numerical experiment section of the present paper, we demonstrate that using the proposed multilayer NE approach, steady solutions with Reynolds numbers as high as 7,500 can be obtained for a rectangular driven cavity flow in three dimensions. As far as we know this is the first time that such solutions are computed directly without using any continuation methods. Our results are validated by a comparison with experimental data and neutral-stability results of previous publications, as well as steady-state solutions obtained from the commercial software CFX [4].

The paper is organized as follows. In section 2, the proposed NE preconditioned inexact Newton method is presented in detail. In section 3, the system of steady-state incompressible Navier–Stokes equations with a finite element discretization is described. In section 4, some numerical experiments for the 3D lid-driven cavity flows are provided, including the validation with experimental results, the comparison with other nonlinear algorithms, and the performance of the parallel solver. Some concluding remarks are given in section 5.

2. The NE preconditioned inexact Newton method. Consider $F: R^n \to R^n$. We aim to find $x^* \in R^n$ such that

$$(2.1) \qquad\qquad\qquad F(x^*) = 0,$$

starting from an initial guess $x^0 \in R^n$, where $F = (F_1, \ldots, F_n)^T$, $F_i = F_i(x)$, and $x = (x_1, \ldots, x_n)^T$. We first recall the inexact Newton algorithm with backtracking

(INB). Assume x^k is the current approximate solution. A new x^{k+1} can be computed via

$$(2.2) \qquad x^{k+1} = x^k + \lambda^k s^k,$$

where the inexact Newton direction s^k satisfies

$$(2.3) \qquad \|F'\left(x^k\right) s^k + F\left(x^k\right)\| \le \eta^k \|F\left(x^k\right)\|.$$

Here $\eta^k \in [0, 1)$ is a forcing term that determines how accurately the Jacobian system needs to be solved. The step length $\lambda^k \in [0, 1]$ determined by line search [17] is the largest value that satisfies

$$(2.4) \qquad f(x^k + \lambda^k s^k) \le f(x^k) + \alpha \lambda^k \nabla f(x^k)^T s^k,$$

where the merit function $f \colon R^n \to R$ is defined as $\|F\|^2/2$, and the parameter α is used to ensure that the reduction of f is sufficient. To enhance the robustness of INB, the forcing term η^k can be computed based on norms that are by-products of the iteration, as suggested by Eisenstat and Walker [21]: start with any $\eta^0 \in [0, 1)$; for $k = 1, 2, \ldots$, we choose

$$(2.5) \qquad \eta^k = \frac{\left|\|F\left(x^k\right)\| - \|F'\left(x^{k-1}\right) s^{k-1} + F\left(x^{k-1}\right))\|\right|}{\|F\left(x^{k-1}\right)\|}.$$

The nonlinear iteration is stopped if

$$(2.6) \qquad \|F\left(x^k\right)\| \le \gamma_r \|F\left(x^0\right)\|,$$

where γ_r is a prescribed relative tolerance for the nonlinear solver.

Remark 2.1. We remark that λ^k is a critically important parameter in INB. The slow convergence of INB happens when the value of λ^k is too small. In practice, the value of λ^k is often determined by a small number of components of F that are more nonlinear than the others.

Remark 2.2. The objective of nonlinear preconditioning is to increase the value of λ^k by balancing the overall nonlinearities of the system so that a single search direction s^k can benefit all components of the residual function F.

Remark 2.3. During the INB process, the relative reduction of the residual $\rho^k = \|F\left(x^k\right)\|/\|F\left(x^{k-1}\right)\|$ can be used to measure the effectiveness of the kth Newton solution x^k. If ρ^k is too large (i.e., $\rho^k > \rho_0$, where $0 < \rho_0 < 1$ is a preselected parameter), a preconditioning step is then introduced to produce a different x^k. In the preconditioning step, the components of F that slow down the convergence are identified and then eliminated approximately using some subspace Newton iterations, that is known as a NE preconditioner.

Remark 2.4. For some really difficult problems, the number of slow components can be large, and the subspace Newton iterations do not converge, thus the nonlinear preconditioning step needs to be further preconditioned. This is the multilayer preconditioning to be discussed below.

2.1. The multilayer NE preconditioner. The key assumption of the single layer NE preconditioning is that the components of F can be decomposed into two groups: a good group of components to be kept, and a bad group to be eliminated

by a subspace INB. The method works well when the number of bad components is relatively small and the subspace INB converges, but for applications that are of interest in this paper, the bad space is too large for the subspace INB to converge.

As suggested in [36] the single layer NE algorithm can be used in a nested fashion, i.e., NE can be used inside another NE when the subspace Newton fails to converge. The idea is simple, but its application is quite difficult. In this paper, we propose a multilayer NE method based on successive sweeping of nested subspace corrections in a multiplicative manner. In contrast to the nested NE method where the elimination restricts to smaller subsets of components as the number of level increases, we start from a small portion of mesh points where the highest nonlinearities are located and then extend this subset to a larger one by including the surrounding mesh points layer by layer. The idea of multilayer is inspired from the *cascadic multigrid* algorithm [9] that performs more iterations on coarser levels in order to use fewer iterations on the finer levels. Denote N_l as the number of layers in NE; a high level description of the multilayer INB-NE method is presented in Algorithm 2.1.

Algorithm 2.1 The multilayer NE preconditioned INB: Given absolute and relative nonlinearity bounds ε and ρ_0.

Step 1 Start from the initial guess x^0 and set $k = 0$, $x^{-1} = x^0$.

Step 2 Check convergence:
- If the global stopping condition $\|F(x^k)\| \leq \gamma_r \|F(x^0)\|$ is satisfied, stop.
- If $\|F(x^k)\| \geq \varepsilon$ and $\|F(x^k)\|/\|F(x^{k-1})\| \geq \rho_0$, go to **Step 3**; otherwise, go to **Step 4**.

Step 3 The NE step: Perform a multilayer subspace correction:
- Set $x^{(0)} = x^k$.
- For $l = 0, \ldots, N_l - 1$:
  - (i) Identify the to-be-eliminated components on the lth layer.
  - (ii) Form and evaluate the nonlinear function $\mathcal{F}^{(l)}(x)$ for the inner Newton iteration.
  - (iii) Approximately solve $\mathcal{F}^{(l)}(x) = 0$ using the classical inexact Newton method with initial guess $x^{(l)}$.
  - (iv) If the inner Newton converges, take the solution as $x^{(l+1)}$; otherwise, set $x^{(l+1)} = x^{(l)}$, break.
  - (v) If $\|F(x^{(l+1)})\|/\|F(x^{k-1})\| < \rho_0$, break.
- Set $x^k = x^{(l+1)}$, go to **Step 4**.

Step 4 The global INB step:
- Inexactly solve $F'(x^k)s^k = -F(x^k)$.
- Compute λ^k using the cubic backtracking line search.
- Update $x^{k+1} = x^k + \lambda^k s^k$.
- Set $k = k + 1$, go to **Step 2**.

We next focus on Step 3 of the algorithm, which is the nonlinear preconditioning step. We assume F arises from the discretization of a PDE on a mesh and each mesh point may have several variables. Note that each component of F involves various operators in the PDE, such as Laplacian and gradient. The strong nonlinearities of the system are often related to some critical features (e.g., boundary layers and local singularities) that appear in certain local regions. In such a situation, F corresponding to a small number of mesh points may contribute to a large percentage of $\|F\|$. We consider a pointwise elimination approach [33, 62], that is, when one variable is selected, all other variables corresponding to that mesh point are also eliminated.
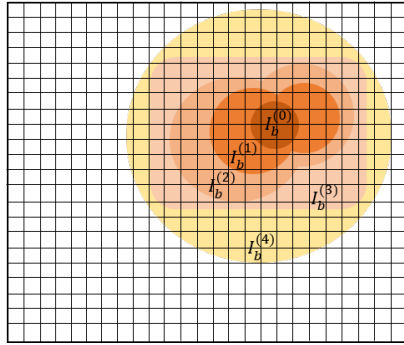
FIG. 1. *An example for illustration of the bad subsets $I_b^{(l)}$ on a 2D mesh, $l = 0, \dots, 4$.*

Specifically, let I be an index set of M mesh points, where each index corresponds to m unknown components x_{i_c} and m nonlinear residual components F_{i_c}, $c = 0, \dots, m-1$. On the lth layer, we decompose I into a "bad" subset $I_b^{(l)}$ with $M_b^{(l)}$ mesh points and a "good" subset $I_g^{(l)} = I \backslash I_b^{(l)}$ with $M - M_b^{(l)}$ mesh points, where $I_b^{(l)}$ and $I_g^{(l)}$ correspond to the variables that have strong and weak nonlinearities, respectively. We denote the bad subset $I_b^{(l)}$ as

$$(2.7) \qquad I_b^{(l)} = \left\{ i \,|\, \text{If } \max_c \{|F_{i_c}(x^{(0)})|\} > \beta_l \|F(x^{(0)})\|_\infty, \; c = 0, \dots, m-1 \right\},$$

where the parameters $\{\beta_l\}$ are a series of descending constants given by the user. As l increases, more mesh points are included in $I_b^{(l)}$, i.e.,

$$(2.8) \qquad I_b^{(0)} \subseteq I_b^{(1)} \subseteq \cdots \subseteq I_b^{(N_l-1)}.$$

An example for illustration of the bad subsets on a 2D mesh is shown in Figure 1.

With the subset $I_b^{(l)}$, we define two subspaces

$$(2.9) \qquad V_b^{(l)} = \left\{ v \,|\, v = (v_0, \dots, v_{n-1})^T \in R^n, \; v_{i_c} = 0 \text{ if } i \notin I_b^{(l)} \right\},$$

$$(2.10) \qquad V_g^{(l)} = \left\{ v \,|\, v = (v_0, \dots, v_{n-1})^T \in R^n, \; v_{i_c} = 0 \text{ if } i \in I_b^{(l)} \right\}.$$

The corresponding restriction operators are denoted as $R_b^{(l)}$ and $R_g^{(l)}$, that map the vectors from R^n to $V_b^{(l)}$ and $V_g^{(l)}$, respectively. Then, the modified nonlinear function $\mathcal{F}^{(l)}$ is defined as

$$(2.11) \qquad \mathcal{F}^{(l)}(x) = R_g^{(l)}(x - x^{(l)}) + R_b^{(l)}(F(x)).$$

The nonlinear system $\mathcal{F}^{(l)}(x) = 0$ is solved by using the classical INB algorithm with the initial guess $x^{(l)}$. $x_*^{(l)}$ is accepted as the approximate solution if the stopping condition $\|\mathcal{F}^{(l)}(x_*^{(l)})\| \leq \gamma_r^{ne} \|\mathcal{F}^{(l)}(x^{(l)})\|$ is satisfied, where γ_r^{ne} is the relative tolerance for the nonlinear solver.

Remark 2.5. For the implementation of the NE preconditioner, we replace the equations corresponding to the good components by $x_{i_c} - x_{i_c}^{(l)} = 0$ and keep the others unchanged, leading to an identity part of the inner Jacobian. Then, the solve of $\mathcal{F}^{(l)}(x) = 0$ can be performed in the whole space. This approach saves the memory as well as the overhead to create inner solvers (with different sizes on different levels) per outer iteration.

Remark 2.6. In (2.7), the selection of bad components depends on the magnitude of the residual F. Scaling the nonlinear system with respect to some field variables may affect the choice of β and the overall performance. We do not account for such a scaling effect in this paper.

Remark 2.7. The multilayer partition of the residual function is adaptive in the sense that it changes with k. On the other hand, the partition depends also on the prescribed parameters $\{\beta_l\}$. As β_l becomes smaller, the dimension of the space for the bad components on the lth layer increases.

Remark 2.8. If the residual norm $\|F(x^k)\|$ is less than the given value of ε, the intermediate numerical solution is considered to be close to the desired one, therefore, the application of NE can be deactivated in order to save the overhead of the nonlinear preconditioning.

Remark 2.9. For the purpose of efficiency, it is not necessary to perform elimination on all N_l layers. We consider two exiting mechanisms: (1) if the inner Newton fails to converge on the lth layer, it might fail on the $(l+1)$th layer too. Consequently, one can stop the sweeping and return the intermediate solution $x^{(l)}$ to the global update phase; (2) if the updated residual norm obtained on the current layer is effectively reduced to a desired value, i.e., $\|F(x^{(l+1)})\|/\|F(x^{k-1})\| < \rho_0$, further elimination on the rest layers is not needed and the sweeping can then be stopped. In practice, the necessity of these exiting mechanisms are often problem dependent. Removing one or more of them results in different variants of the algorithm.

Since the intermediate solution $x^{(l+1)}$ is not guaranteed to be smooth on or near the interface points between the good and bad regions, new jumps may be produced in the residual function $F(x^{(l+1)})$ across the interface and this may lead to the relocation of unbalanced nonlinearities. Unfortunately, such interfacial jump pollution cannot be effectively reduced by the global update phase [62]. In [14], Cai and Li used a simple one-dimensional nonlinear two-point boundary value problem to illustrate the issue. The authors suggested that some mechanism must be taken to avoid such jumps. For instance, one can choose carefully the stopping condition for the inner Newton iteration to make sure that the jump is not too large or use a large enough overlap to move the jump away from the interior of the bad region [33, 49, 50, 59]. In [62], Yang and Hwang suggested identifying the particular field variable that dominates the jump pollution and then including its auxiliary linearized subsystem in the elimination process. In practice, this hybrid physical-algebraic approach requires extra analysis from numerical results to determine which field variable is responsible for the dominant part of the residual norm. We claim that these mechanisms are not needed in the proposed multilayer method, since the potential new jumps on the lth layer are smoothed out automatically by the $(l+1)$th subspace correction step, and the strength of the new jumps becomes smaller as l increases, which are evidenced in our numerical experiments.

2.2. The linear solver. A linear solver is required to solve the Jacobian systems arising from both the outer Newton iteration and the NE process. Good candidates include the class of Krylov subspace methods with effective linear preconditioners. In this study, we use a restricted additive Schwarz (RAS) [12] preconditioned generalized minimal residual (GMRES) method [53].

We rewrite the Jacobian system in the following general form:

$$(2.12) \qquad JM^{-1}y = b \quad \text{with} \quad x = M^{-1}y,$$

where J is the Jacobian matrix, M is the preconditioner, x is the solution, and b is the right-hand side. Denoting by np the number of processor cores of the parallel computer, we partition the computational domain (or more precisely the finite element mesh) Ω_h into np nonoverlapping subdomains Ω_p (i.e., $\Omega_i \cap \Omega_j = \emptyset \ \forall i \neq j$) for $p = 1, \ldots, np$, such that $\Omega_h = \Omega_1 \cup \cdots \cup \Omega_{np}$. The subvector associated with Ω_p is denoted as y_p. We then extend Ω_p to overlap with its neighbors by δ layers of mesh elements and denote the overlapping subdomain as Ω_p^δ. On each overlapping subdomain, we define the corresponding subvector y_p^δ and the restriction operator R_p^δ that maps the global vector of unknowns in Ω_h to y_p^δ, i.e.,

$$y_p^\delta = R_p^\delta y = (I \quad 0) \begin{pmatrix} y_p^\delta \\ y \backslash y_p^\delta \end{pmatrix}.$$

We denote by R_p^0 the restriction operator that returns y_p defined on the nonoverlapping subdomain. Then, the RAS preconditioner [12] is defined as

$$(2.13) \qquad M_{\text{RAS}}^{-1} = \sum_{p=1}^{np} \left(R_p^0\right)^T \left(J_p\right)^{-1} R_p^\delta.$$

In (2.13), $\left(J_p\right)^{-1}$ is understood as an approximate inverse of the subdomain Jacobian matrix, its product with a vector is computed by solving a subdomain linear system inexactly using a point-block incomplete LU (ILU) factorization of J_p.

3. The steady-state incompressible Navier–Stokes equations. Let Ω be a bounded domain in R^3. We consider the steady-state incompressible Navier–Stokes equations

$$(3.1) \qquad \begin{cases} \rho(\mathbf{u} \cdot \nabla)\mathbf{u} - \nabla \cdot \boldsymbol{\sigma} = \rho\mathbf{f} & \text{in} \quad \Omega, \\ \nabla \cdot \mathbf{u} = 0 & \text{in} \quad \Omega. \end{cases}$$

Here $\mathbf{u} = (u, v, w)^T$ is the velocity, ρ is the density, $\mathbf{f}$ is a given body force per unit mass, and $\boldsymbol{\sigma} = -p\mathbf{I} + \mu\left(\nabla\mathbf{u} + \nabla\mathbf{u}^T\right)$ is the Cauchy stress tensor, where $\mathbf{I}$ is an identity matrix, p is the pressure, and μ is the viscosity coefficient.

In this paper, we consider flows confined in the cavity domains with SARs 1:1 and 0.5:1, respectively, as depicted in Figure 2. The top surface Γ_{lid} moves with velocity U in the positive x-direction. On all walls we impose a no-slip and no-penetration boundary condition $\mathbf{u} = \mathbf{u}_{\text{W}}$ with

$$(3.2) \qquad \mathbf{u}_{\text{W}} = \begin{cases} (U, 0, 0) & \text{on} \quad \Gamma_{\text{lid}}, \\ \mathbf{0} & \text{on} \quad \partial\Omega \setminus \Gamma_{\text{lid}}. \end{cases}$$

The flow pattern inside the cavity consists of several standing vortices with their characteristics dominated by the Reynolds number $\text{Re} = \rho U L / \mu$.
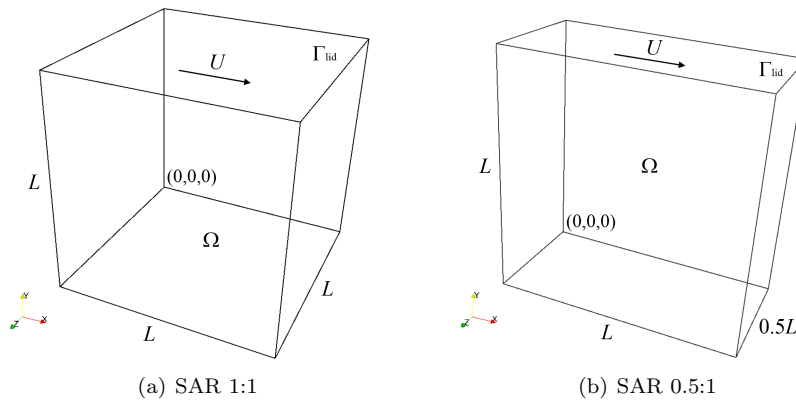
Fig. 2. *Cavity domains with two SARs.*

A P_1-P_1 stabilized finite element method is used for the discretization of (3.1)–(3.2), which results in a nonlinear system $F(x) = 0$ to be solved. The trial function spaces for the velocity and the pressure are defined as

$$\mathbf{V} = \left\{ \mathbf{u} \,|\, \mathbf{u} \in [H^1(\Omega)]^3, \mathbf{u} = \mathbf{u}_W \text{ on } \partial\Omega \right\}, \quad P = \left\{ p \,|\, p \in L^2(\Omega), \int_\Omega p \, \mathrm{d}\Omega = 0 \right\}.$$

The weighting function space for the velocity is defined as

$$\mathbf{V}_0 = \left\{ \mathbf{u} \,|\, \mathbf{u} \in [H^1(\Omega)]^3, \mathbf{u} = \mathbf{0} \text{ on } \partial\Omega \right\}.$$

Following standard notations, the weak form reads: find $(\mathbf{u}, p) \in \mathbf{V} \times P$ such that for $\forall\ (\mathbf{w}, q) \in \mathbf{V}_0 \times P$,

$$(3.3) \qquad\qquad\qquad\qquad B(\mathbf{u}, p; \mathbf{w}, q) = 0$$

with

$$B(\mathbf{u}, p; \mathbf{w}, q) = \int_\Omega \rho \left(\mathbf{u} \cdot \nabla \right) \mathbf{u} \cdot \mathbf{w} \mathrm{d}\Omega + \int_\Omega \mu \left(\nabla \mathbf{u} + \nabla \mathbf{u}^T \right) : \nabla \mathbf{w} \mathrm{d}\Omega$$

$$(3.4) \qquad\qquad - \int_\Omega p \nabla \cdot \mathbf{w} \mathrm{d}\Omega + \int_\Omega (\nabla \cdot \mathbf{u}) \, q \mathrm{d}\Omega - \int_\Omega \rho \, \mathbf{f} \cdot \mathbf{w} \mathrm{d}\Omega.$$

Let $\Omega_h = \{K\}$ be a quasi-uniform conforming tetrahedral mesh of Ω with h the diameter of the element $K \in \Omega_h$; denote by $(\cdot, \cdot)_K$ the L^2-inner product over element K. We define the finite element subspaces $\mathbf{V}^h$, $\mathbf{V}_0^h$, and P^h as the counterparts of their infinite dimensional subspaces. Then, the discretized system of (3.3) is described as follows: find $(\mathbf{u}_h, p_h) \in \mathbf{V}^h \times P^h$ such that for $\forall\ (\mathbf{w}_h, q_h) \in \mathbf{V}_0^h \times P^h$,

$$(3.5) \qquad\qquad\qquad\qquad B_S(\mathbf{u}_h, p_h; \mathbf{w}_h, q_h) = 0$$

with

$$B_S(\mathbf{u}_h, p_h; \mathbf{w}_h, q_h) = B(\mathbf{u}_h, p_h; \mathbf{w}_h, q_h) + \sum_{K \in \Omega_h} \left(\mathcal{S}_h, \tau_m \left((\mathbf{u}_h \cdot \nabla) \mathbf{w}_h + \nabla q_h \right) \right)_K$$

$$(3.6) \qquad\qquad + \sum_{K \in \Omega_h} \left(\nabla \cdot \mathbf{u}_h, \tau_c \nabla \cdot \mathbf{w}_h \right)_K,$$

where $\mathcal{S}_h$, τ_m, and τ_c are defined as

$$\mathcal{S}_h = \rho\left(\mathbf{u}_h \cdot \nabla\right)\mathbf{u}_h + \nabla p_h - \rho\mathbf{f}_h,$$

$$\tau_m = \left(\sqrt{\mathbf{u}_h \cdot G \cdot \mathbf{u}_h + 36\mu^2/\rho^2 G : G}\right)^{-1},$$

$$\tau_c = \left(8\tau_m \mathrm{tr}(G)\right)^{-1}.$$

Here $(G)_{ij} = \sum_{l=1}^3 \frac{\partial \xi_l}{\partial x_i}\frac{\partial \xi_l}{\partial x_j}$ $(i,j = 1,2,3)$ is the covariant metric tensor and $\frac{\partial \xi}{\partial x}$ refers to the Jacobian of the mapping between the reference and the physical element. The last two terms in (3.6) are added in order to satify the Ladyzenskaja–Babuska–Brezzi inf-sup condition [8]. We refer to [58] for more details of the discretization scheme.

In accordance with the pointwise NE approach and the point-block ILU subdomain solver, here the components of velocity and pressure are ordered point by point, for example,

$$x = (u_0, v_0, w_0, p_0, u_1, v_1, w_1, p_1, \ldots, u_{M-1}, v_{M-1}, w_{M-1}, p_{M-1})^T.$$

We refer to this ordering as the point-block ordering, in which the unknowns associated with each mesh point are always together in the same 4×4 block. Such an ordering has a significant impact on the convergence properties of the algebraic solver and the parallel scalability of the overall method [31].

4. Numerical experiments. In this section, we first validate the finite element discretization and the proposed algorithm by comparing the velocity profiles of the cavity flows with previous publications. Then, we study the performance of the proposed algorithm and focus on (1) the robustness of the algorithm with respect to the Reynolds numbers, the mesh size, and the preselected parameters; and (2) a comparison of the numerical performance between the new method, the single layer NE, and the classical INB method.

The algorithms are implemented using PETSc [6]. All computations are carried out on the Shaheen2 supercomputer which has two 16-core Intel Haswell CPUs and 128 GB local memory in each of its compute nodes. Quasi-uniform tetrahedral meshes for the cavity domains are generated using CUBIT [51] and the partitions are obtained using ParMETIS [39]. Figure 3 shows a sample partition of the mesh into 16 subdomains. Multilayer nonlinear preconditioning, being algebraic in nature, can also be applied to problems on unstructured adaptive meshes lacking quasi-uniformity. The single layer approach (i.e., $N_l = 1$) is applied on unstructured meshes with local refinement in [49, 50], including the driven cavity flow at low Reynolds numbers where the mesh is locally refined near the moving lid. For the higher Reynolds number considered in this paper, we find that a globally refined mesh better captures the steady-state solution, because there are eddies with diminishing sizes near the bottom corners of the cavity, and flow separation appears in the core region.

In the Navier–Stokes equations, the body force is ignored for all experiments. A zero vector is used as the initial guess, i.e., $x^0 = \mathbf{0}$. The length of the cavity is $L = 1$ m. A fluid with density $\rho = 1$ kg/m^3 is driven by the wall at $y = L$ with a constant velocity $U = 1$ m/s in the positive x-direction. We vary the viscosity μ to test different Reynolds numbers. The Jacobian matrices arising from both the outer Newton iteration and the NE step are computed analytically.

We use the following parameters in our solvers if they are not specifically stated. A sufficiently small relative tolerance $\gamma_r = 10^{-15}$ is used for the global nonlinear
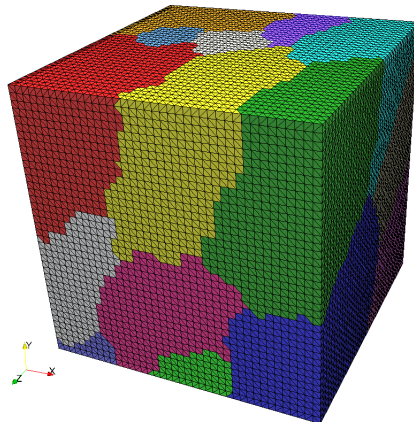
FIG. 3. *A sample partition of the mesh into* 16 *subdomains. Different colors refer to different partitions.*

TABLE 1
The mesh information for the test cases of cavity flows with different Re's.

h	#elements	#nodes	Re	np
		SAR=1:1		
1/100	6,000,000	1,030,301	100–1,000	1,024
1/160	24,576,000	4,173,281	1,500–1,900	2,048
		SAR=0.5:1		
1/40	192,000	35,301	3,200	32
1/80	1,536,000	269,001	3,200	512
1/160	12,288,000	2,099,601	3,200	1,024
1/256	50,331,648	8,520,321	1,000–7,500	4,096

solver to ensure the convergence. The relative tolerance for the nonlinear solver in the subspace correction step is $\gamma_r^{ne} = 10^{-1}$. The application of NE is deactivated if the residual norm is less than $\varepsilon = 10^{-6}$. The restart value of GMRES is fixed at 400. The size of overlap in the linear RAS preconditioner is fixed to $\delta = 2$. A point-block ILU factorization with 3 fill-in levels is used to solve the subdomain linear systems. For the NE preconditioner, we set the preselected parameter ρ_0 to be 0.8 for judging if the residual is reduced slowly.

In the rest of this paper, "np" denotes the number of processor cores used for the test; "NI_g" denotes the number of global Newton iterations, "LI_g" denotes the averaged number of GMRES iterations per global Newton iteration, "N_{ne}" is the number of applications of the NE preconditioner, "NI_{sub}" refers to the averaged number of Newton iterations per subspace correction step, "LI_{sub}" is the averaged number of GMRES iterations per Newton in the subspace correction, "$\mathrm{T}_{ne}(\mathrm{s})$" is the compute time in seconds for all NE applications, and "$\mathrm{T}_t(\mathrm{s})$" is the total compute time in seconds for the overall nonlinear solver.

4.1. Validation of the proposed numerical method. We first validate the finite element discretization and the proposed solver by comparing our results for benchmark problems with experimental and numerical data in previous literature. A sequence of refined meshes ranging from 192,000 to 50,331,648 elements are used for the tests, as summarized in Table 1.
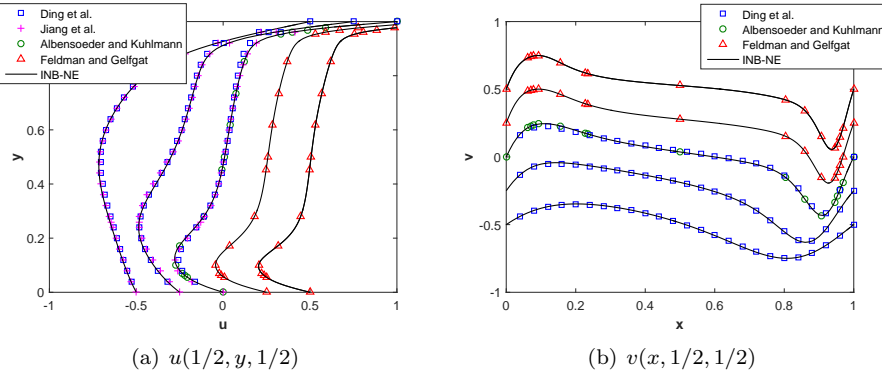
(a) $u(1/2, y, 1/2)$ (b) $v(x, 1/2, 1/2)$

FIG. 4. *Velocity profiles of the cavity flow with SAR* $1 : 1$ *at different Re on the midplane* $z = 1/2$. *Note that the profiles are shifted for visual comparison.* (a) u, *from left to right:* $Re = 100$, 400, $1,000$, $1,500$, *and* $1,900$; (b) v, *from bottom to top:* $Re = 100$, 400, $1,000$, $1,500$, *and* $1,900$.

TABLE 2
Case configuration for the cavity flow with SAR $0.5 : 1$. Δt *is the time step size in LBM.*

Method	Model	Mesh	h	Δt	Re	Stopping condition
INB-NE	Steady-state	Tetrahedral	1/256		1,000–7,500	$\|F(x^k)\|/\|F(x^0)\| \leq 10^{-15}$
CFX	Steady-state	Cartesian	1/256		1,000–3,200	$\sqrt{\sum_{i=1}^{n}(x_i^k - x_i^{k-1})^2/n} \leq 10^{-8}$
LBM	Time-dependent	Cartesian	1/200 1/400 1/800	1/200 1/400 1/800	3,200 5,000 7,500	$\|x^n - x^{n-1}\|_\infty \leq 10^{-7}$

Figure 4 shows the two velocity components u and v along the vertical and horizontal centerlines on the midplane of the cavity with SAR 1:1. The results obtained using INB-NE for cases Re = 100, 400, 1,000, 1,500, and 1,900 show excellent agreement with the results of Albensoeder and Kuhlmann [2], Ding et al. [19], and Feldman and Gelfgat [26]. Note that their results are obtained by solving time-dependent equations until the steady states are reached. Based on computational studies, Gelfgat concluded in the recent publication [27] that the critical Reynolds number of the steady-oscillatory transition for the case SAR 1:1 is around 2,000, in other words, our proposed solver resolves the problem well by solving the steady-state equation directly.

For the case with SAR 0.5:1, the drag effect of the end walls is expected to be stronger than the case SAR 1:1, causing the flow structure to become more organized. Due to the lack of numerical results for this case to date, we compare the results of INB-NE with steady-state solutions obtained using the commercial software ANSYS-CFX (CFX), neutral-stability solutions obtained using an in-house code for solving the time-dependent lattice Boltzmann equations (LBM), as well as the mean velocities measured by laboratory experiments in Prasad and Koseff [52]. CFX uses an element-based finite volume method with a *high resolution* scheme [4] in space for the steady-state simulation. LBM uses a second-order finite different method for spatial discretization and an explicit Euler scheme for temporal discretization. The Smagorinsky's algebraic eddy viscosity approach is incorporated into LBM for large-eddy simulation of high Reynolds number flows [63]. Details of the case configuration are listed in Table 2. Figure 5 shows the two velocity components u and v along the vertical and horizontal centerlines on the midplane of the cavity. We note that

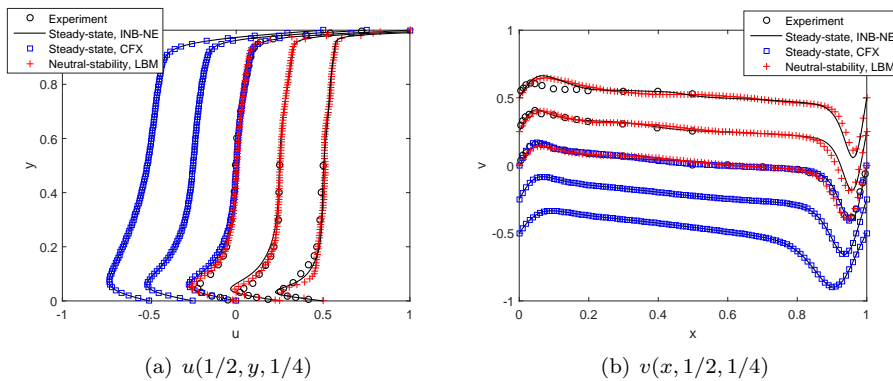(a) $u(1/2, y, 1/4)$ (b) $v(x, 1/2, 1/4)$

FIG. 5. *Velocity profiles of the cavity flow with SAR $0.5 : 1$ at different Re on the midplane $z = 1/4$. Note that the profiles are shifted for visual comparison.* (a) u, from left to right: $Re = 1,000$, $2,000$, $3,200$, $5,000$, and $7,500$; (b) v, from bottom to top: $Re = 1,000$, $2,000$, $3,200$, $5,000$, and $7,500$.
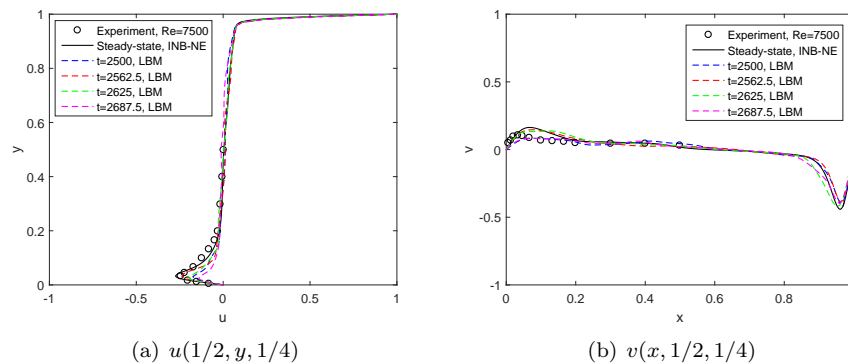

(a) $u(1/2, y, 1/4)$ (b) $v(x, 1/2, 1/4)$

FIG. 6. *Velocity profiles of the cavity flow with SAR $0.5 : 1$ at $Re = 7,500$ on the midplane $z = 1/4$.*

LBM fails to meet the stopping condition at $Re = 7,500$ due to oscillation of the solution when it reaches $\|x^n - x^{n-1}\|_\infty < 10^{-4}$, where x^n is the solution of LBM at the nth time step. Therefore, we take the solution at a snapshot $t = 2562.5$ for this case. Overall, a good agreement is achieved among the methods. The results obtained using INB-NE for cases $Re = 1,000, 2,000$, and $3,200$ match very well with those obtained using CFX, but the latter fails to converge when $Re > 3,200$. The results of INB-NE and LBM coincide with the experimental data at high Reynolds numbers, except for a part of $v(x, 1/2, 1/4)$ nearby the upstream wall at $Re = 7,500$. To investigate the reason for such deviation, we plot several snapshots of the LBM solution at different times in Figure 6. It is found that the velocity profiles change with time, indicating certain unstable perturbations for the time-dependent LBM solution. It is interesting to note that two of the snapshots $t = 2,500$ and $t = 2,687.5$ coincide with the experimental data of $v(x, 1/2, 1/4)$, while the other two coincide with the results of INB-NE, which implies some underlying relationship between the steady-state method and the time-dependent method. However, a thorough investigation of the problem is beyond the scope of this paper.
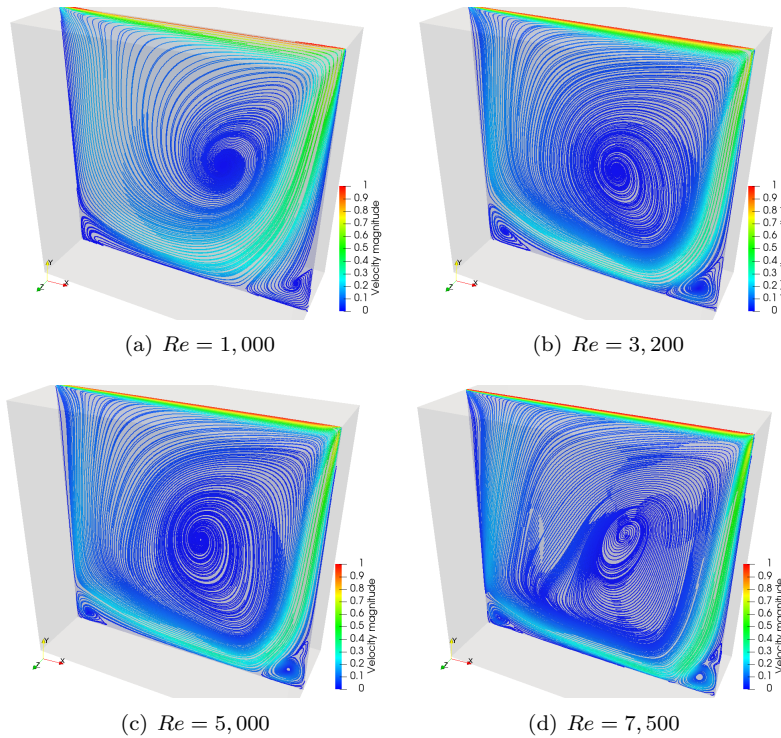
(a) $Re = 1,000$

(b) $Re = 3,200$

(c) $Re = 5,000$

(d) $Re = 7,500$

FIG. 7. *Slice view of streamlines on the midplane $z = 1/4$ for the cavity flow with SAR $0.5 : 1$. The mesh size is $h = 1/256$.*

Figures 7–8 show the slice view of streamlines on the midplane $z = 1/4$ and $x = 1/2$ for the cavity flow with SAR 0.5:1, respectively. The mesh size is $h = 1/256$. The midplane $z = 1/4$ shows clearly the primary and secondary vortices, while the midplane $x = 1/2$ shows a symmetric pattern, demonstrating strong 3D characteristic of the flow. As Re increases, a sequence of eddies with diminishing size are observed at the corner of the cavity, at the same time, flow separation appears symmetrically perpendicular to the driving direction of the lid. In short, for the case SAR 0.5:1, the proposed approach is able to capture the steady-state solutions at high Reynolds numbers. For the rest of the paper, we will use this case as a benchmark to compare various algorithms.

4.2. Comparison of INB and INB-NE. As Re increases, the nonlinear system becomes harder to solve. Figure 9 shows the history of nonlinear residuals obtained using INB and INB-NE for the cavity flow with SAR 0.5:1 and different Re's. The mesh size is $h = 1/256$. For INB-NE, the maximum number of layers is set to be $N_l = 6$, and the preselected parameters are given as $\beta_l = 0.25 \times 10^{-l}$, $l = 0, \ldots, N_l - 1$. The linear solve in the classical INB diverges at the 8th, 6th, 6th, and 6th global Newton step for Re $= 2,000$, 3,200, 5,000, and 7,500, respectively. In contrast, the proposed INB-NE converges well for all cases with Re $= 1,000$–$7,500$. Note that we use a zero vector as the initial guess, i.e., $x^0 = \mathbf{0}$, the first global Newton step is able to compute an approximate solution that satisfies the boundary condition (enforcing the speed of the lid), hence, a significant decrease of the residual norm is observed at this
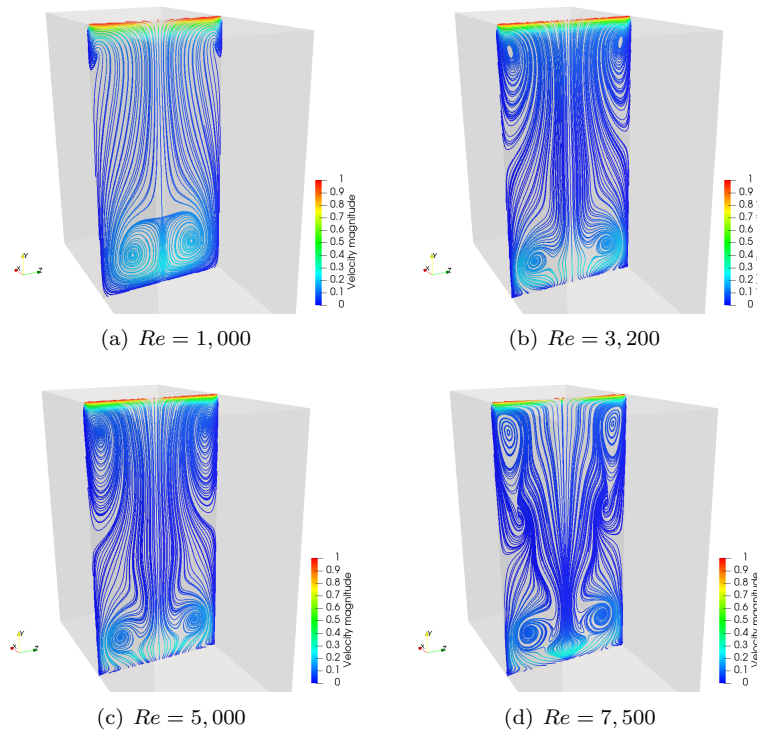
(a) $Re = 1,000$

(b) $Re = 3,200$

(c) $Re = 5,000$

(d) $Re = 7,500$

FIG. 8. *Slice view of streamlines on the midplane $x = 1/2$ for the cavity flow with SAR $0.5 : 1$. The mesh size is $h = 1/256$.*

FIG. 9. *Nonlinear residual history obtained using INB and INB-NE for the cavity flow with SAR $0.5 : 1$. The mesh size is $h = 1/256$. $N_l = 6$, $\beta_l = 0.25 \times 10^{-l}$, $l = 0, \ldots, N_l - 1$.*

step either by INB or INB-NE. A detailed comparison for the numbers of iterations and the total compute times between the two methods are shown in Table 3. We note that the maximum Reynolds number for which INB converges is $Re = 1,739.5$. INB requires less compute time than INB-NE for the case with a small Reynolds number $Re = 1,000$, however, as Re increases, it needs more nonlinear and linear iterations

TABLE 3
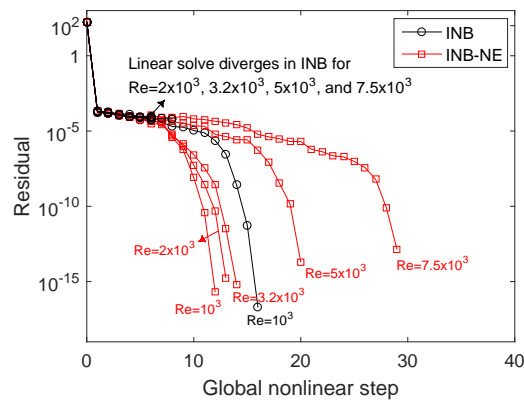*The numbers of iterations and compute times obtained using INB and INB-NE for the cavity flow with SAR $0.5:1$. The mesh size is $h = 1/256$. $N_l = 6$, $\beta_l = 0.25 \times 10^{-l}$, $l = 0, \ldots, N_l - 1$.*

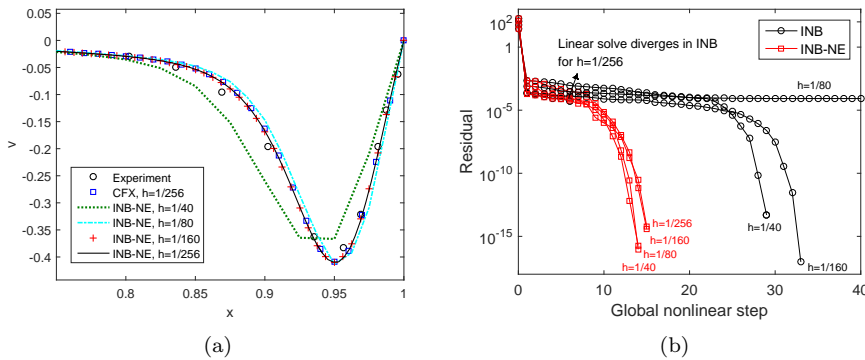| Re	NI_g	LI_g	$T_t(s)$	N_{ne}	NI_{sub}	LI_{sub}	$T_{ne}(s)$
			INB				
1,000	16	210.3	350.87				
1,739.5	18	683.3	853.44				
			INB-NE				
1,000	12	202.3	674.84	3	4.3	11.6	416.88
1,739.5	14	215.4	742.06	2	5.0	14.5	433.51
2,000	13	190.9	975.65	3	4.9	9.8	707.12
3,200	14	200.0	1,419.64	3	4.4	6.6	1,026.90
5,000	20	295.3	1,559.65	7	3.6	12.9	1,026.26
7,500	29	308.5	3,014.03	12	4.1	49.0	2,210.31

         (a)   (b)

FIG. 10. *Convergence test using different mesh sizes for the cavity flow with SAR $0.5:1$ at $Re = 3,200$. (a) Enlarged view for the velocity profiles of $v(x, 1/2, 1/4)$. (b) Nonlinear residual history obtained using INB and INB-NE. $N_l = 6$, $\beta_l = 0.25 \times 10^{-l}$, $l = 0, \ldots, N_l - 1$.*

to converge. For $Re = 1,739.5$, INB-NE performs better than INB in terms of the numbers of global iterations and the total compute time. It can also be observed that, with further increase of Re, more applications of NE are required in the new algorithm, thus more compute time is spent on the NE preconditioner, which is about $2/3$ of the total compute time. The efficiency can be improved with dynamic load balancing, since the selection of bad components varies on different layers and at different global Newton iterations. A strategy to dynamically partition and redistribute the bad subset among all the processor cores so that each processor core owns comparable numbers of bad components will be investigated in future implementations.

Next, we conduct a convergence test using different mesh sizes for the cavity flow at $Re = 3,200$, and study the robustness of the proposed method with respect to the mesh size. Results are shown in Figure 10. It is clear from Figure 10(a) that mesh-independent results are achieved when $h \leq 1/160$. From Figure 10(b) we see that the convergence of INB is sensitive to the mesh size; on the other hand, the proposed INB-NE converges for all cases with the number of nonlinear iterations almost independent of the mesh size.

To see how the multilayer NE preconditioner improves the convergence of the global Newton iteration, we show in Figures 11–12 the residual contours of the u-component of the velocity obtained using the classical INB and the proposed INB-NE

(a) 3rd global Newton step using INB

(b) 4th global Newton step using INB

(c) 3rd global Newton step using multilayer INB-NE

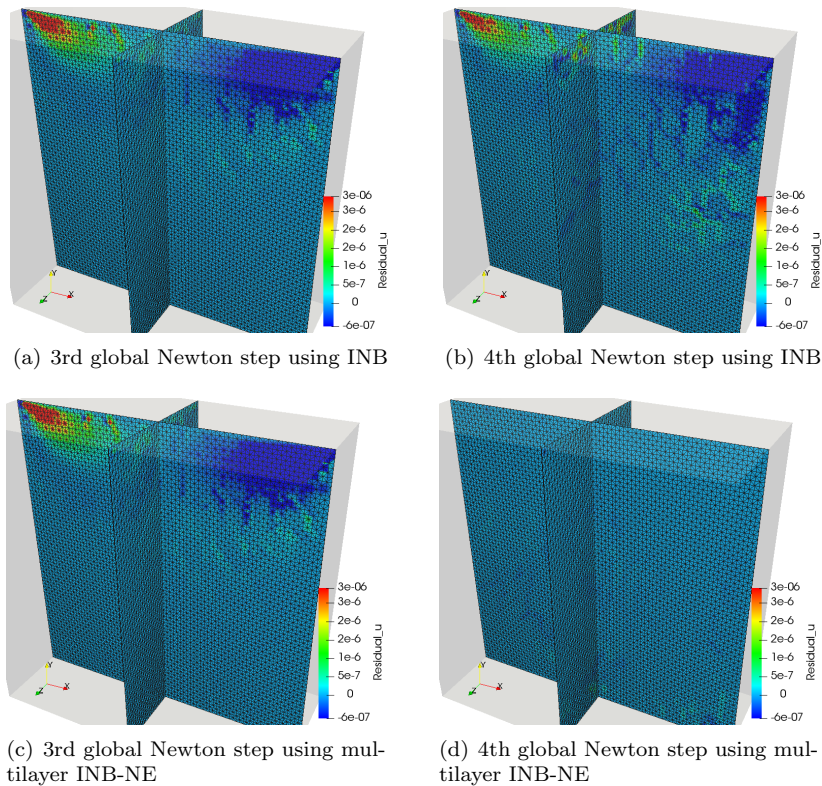(d) 4th global Newton step using multilayer INB-NE

FIG. 11. *Slice view for residual contour of the u-component obtained using the classical INB and the multilayer INB-NE for the cavity flow with SAR 0.5 : 1 at Re = 3, 200. The mesh size is h = 1/80.*

for the cavity flow with $h = 1/80$ at Re $= 3, 200$, respectively. We can see from Figures 11(a)–(b) that a single global Newton step (the 3rd step) in INB barely changes the distribution and strength of the local high nonlinearities characterized by large value of $|F_i|$. In comparison, with the use of the multilayer NE preconditioner (Figures 11(c)–(d)), the overall nonlinearities after the same Newton step are balanced sufficiently and that leads to the convergence of the global Newton iteration. Figure 12 shows the bad subsets colored in red and the residual contours after the subspace correction of NE on the lth layer. We see clearly that the bad subset starts from a small number of mesh points near the top upstream and downstream corners where the highest nonlinearities are located. As l increases, the bad subset is extended to include more surrounding mesh points layer by layer. After the subspace correction is performed, the local high nonlinearities on the lth layer are effectively removed. It can also be observed from the residual contours that some new jumps are produced across the boundary of the bad subset; as expected, they are removed by the next subspace correction step and the strength becomes smaller as evidenced in the figure. The isolated bad components come to the fore because the strongest nonlinearities have already been removed by lower layers. They are scattered due to the complex structure of the 3D cavity flow. However, they can be eliminated effectively layer by layer without degrading the convergence of the outer Newton iteration.

(a) $l = 1$ (b) $l = 1$

(c) $l = 2$ (d) $l = 2$

(e) $l = 3$ (f) $l = 3$

(g) $l = 4$ (h) $l = 4$

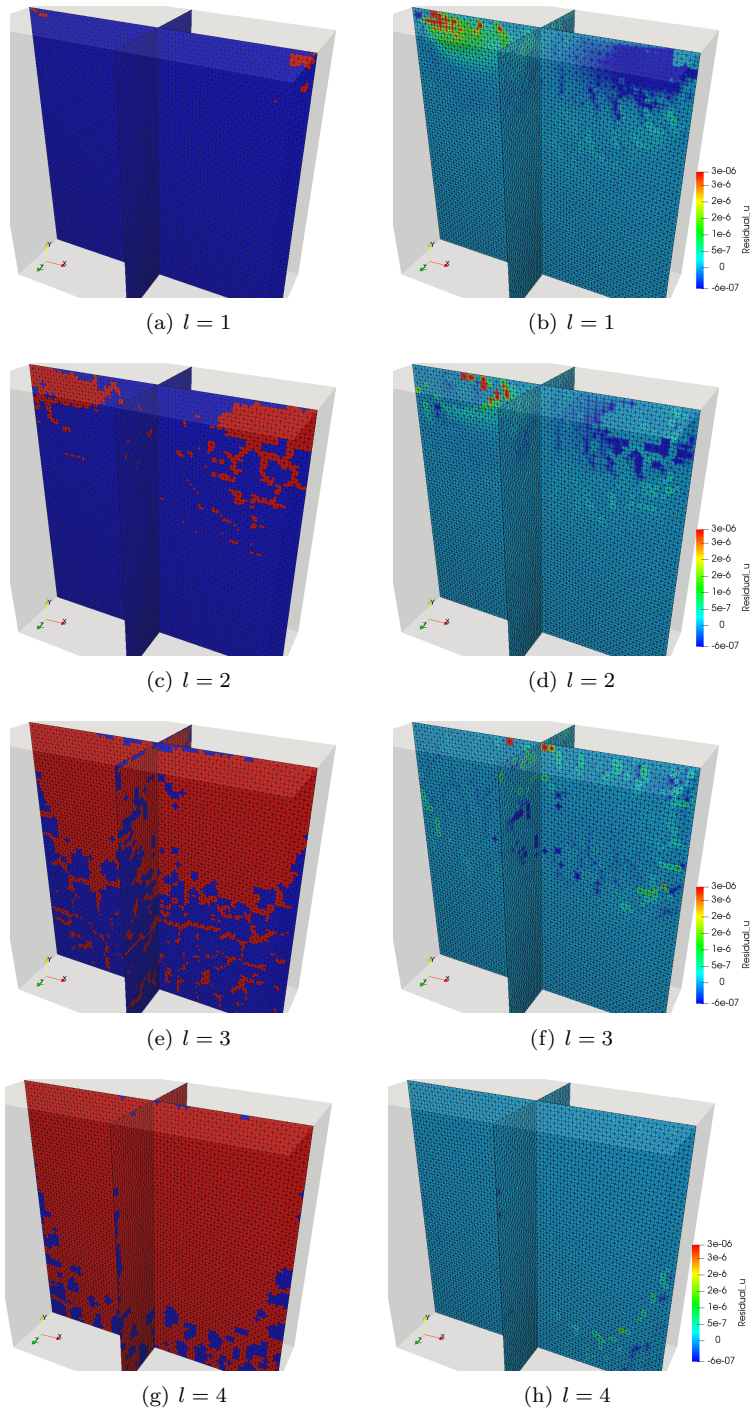FIG. 12. *Slice view on the lth layer at the 3rd global Newton step obtained using the multilayer INB-NE for the cavity flow with SAR $0.5 : 1$ at $\mathrm{Re} = 3,200$. (a), (c), (e), (g): bad subsets colored in red; (b), (d), (f), (h): residual contours of the u-component after the subspace correction. The mesh size is $h = 1/80$. $N_l = 6$, $\beta_l = 0.5 \times 10^{-l}$, $l = 0, \ldots, N_l$.*

TABLE 4

The numbers of iterations and compute times obtained using INB-NE with different preselected parameters for the cavity flow with SAR 0.5 : 1 at Re = 3, 200. The mesh size is h = 1/256. $\beta_l = \beta \times 10^{-l}$, $l = 0, \ldots, N_l - 1$. "–" indicates that the case fails to converge.

β	0.75	0.5	0.25	0.1	0.75	0.5	0.25	0.1
	$N_l = 5$				$N_l = 6$			
NI_g	21	23	19	20	–	20	14	15
LI_g	242.9	345.5	242.7	206.9	–	192.6	200.0	215.0
T_t(s)	1,369.92	1,402.32	1,236.33	1,022.27	–	863.79	1,419.64	1,021.98
N_{ne}	13	6	4	3	–	4	3	3
NI_{sub}	2.2	3.0	4.0	4.4	–	3.4	4.4	4.5
LI_{sub}	3.3	9.9	5.6	6.9	–	9.6	6.6	5.7
T_{ne}(s)	869.85	740.36	791.02	585.09	–	448.20	1,026.90	680.24
	$N_l = 7$				$N_l = 8$			
NI_g	–	14	16	13	22	14	16	13
LI_g	–	183.0	230.7	313.6	353.5	183.0	230.7	313.6
T_t(s)	–	772.69	831.75	774.04	2,095.96	779.31	829.2	778.98
N_{ne}	–	3	3	2	5	3	3	2
NI_{sub}	–	4.0	3.8	4.3	4.1	4.0	3.8	4.3
LI_{sub}	–	11.8	9.2	8.7	7.7	11.8	9.2	8.7
T_{ne}(s)	–	489.36	466.77	419.66	1,442.49	494.53	465.14	423.23

4.3. Performance with respect to the preselected parameters. To understand the impact of the parameters on the performance of the NE preconditioner, we test the case SAR 0.5:1 with $h = 1/256$ at Re $= 3, 200$ using different values of N_l and β_l. For simplicity, we use $\beta_l = \beta \times 10^{-l}$ and change the value of β in the tests. The resulting numbers of iterations and the compute times are shown in Table 4. When N_l increases, more subspace correction steps may be performed in the NE preconditioner, which results in fewer global Newton iterations. Moreover, the time spent on the overall NE preconditioner decreases because the number of NE steps is reduced. On the other hand, the performance of NE depends also on the choice of β_l, which controls the dimension of the subspace problem on each layer. From the table we find that the pair of values $(N_l, \beta_l) = (7, 0.5 \times 10^{-l})$ is suitable to minimize the total compute time and guarantees the convergence of the overall algorithm. We remark that the choice of these parameters depends on the value of Reynolds number and the mesh size used for the problem.

4.4. Comparison of the single layer approach and the multilayer approach. If only one subspace correction step is performed in Step 3 of Algorithm 2.1, the multilayer NE preconditioner reduces to the classical pointwise approach. In this paper, we call this approach "single layer INB-NE" and compare it with the proposed multilayer approach in terms of the robustness and efficiency. In the single layer approach, the bad subset of mesh points I_b^k is defined as

$$(4.1) \qquad I_b^k = \left\{ i \,|\, \text{If } \max_c\{|F_{i_c}(x^k)|\} > \beta_0 \|F(x^k)\|_\infty, \ c = 0, \ldots, m-1 \right\}.$$

Following the idea of restricted elimination, we define the restricted bad subset $I_b^{k,\varepsilon}$ by excluding some mesh points near the boundary of I_b^k, such that

$$(4.2) \qquad I_b^{k,\varepsilon} = \left\{ i \,|\, \text{If } \max_c\{|F_{i_c}(x^k)|\} > (\beta_0 + \varepsilon_0)\|F(x^k)\|_\infty, \ c = 0, \ldots, m-1 \right\},$$

where ε_0 is a preselected constant. We refer to [33, 50, 62] for more details of the single layer approach. Table 5 shows the numbers of iterations and compute times obtained using INB-NE with the single layer approach and the multilayer approach for the cavity flow with SAR 0.5:1 and different Re's. The residual histories are

TABLE 5

The numbers of iterations and compute times obtained using INB-NE with the single layer approach and the multilayer approach for the cavity flow with SAR 0.5 : 1. The mesh size is $h = 1/256$.

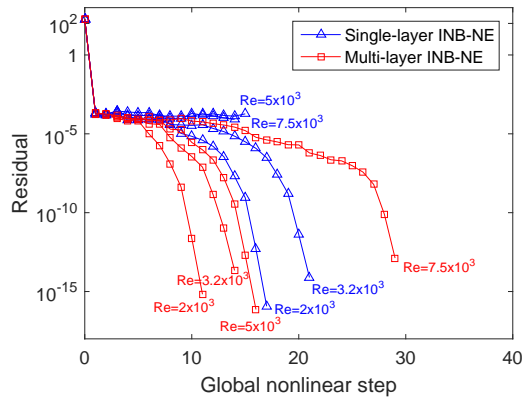| Re	NI_g	LI_g	$T_t(s)$	N_{ne}	NI_{sub}	LI_{sub}	$T_{ne}(s)$
			Single layer INB-NE				
2,000	17	195.5	667.33	5	5.0	8.9	306.52
3,200	21	218.4	854.04	7	4.4	6.8	376.47
			Multilayer INB-NE				
2,000	11	241.4	691.80	2	3.9	9.3	434.44
3,200	14	183.0	772.69	3	4.0	11.8	489.36
5,000	16	273.9	1,379.15	4	4.4	6.3	970.27
7,500	29	308.5	3,014.03	12	4.1	49.0	2,210.31

FIG. 13. *Nonlinear residual history obtained using INB-NE with the single layer approach and the multilayer approach for the cavity flow with SAR 0.5 : 1. The mesh size is $h = 1/256$.*

plotted in Figure 13. The mesh size is $h = 1/256$. Optimal parameters are used for different cases. For the single layer approach, we choose $\beta_0 = \varepsilon_0 = 0.5 \times 10^{-3}$; for the multilayer approach, $(N_l, \beta_l) = (7, 0.5 \times 10^{-l})$ is used for the case Re = 2,000 and 3,200, $(N_l, \beta_l) = (7, 0.25 \times 10^{-l})$ is used for the case Re = 5,000, and $(N_l, \beta_l) = (6, 0.25 \times 10^{-l})$ is used for the case Re = 7,500, $l = 0, \ldots, N_l - 1$. The multilayer approach is expected to have a stronger elimination effect on the overall nonlinearities, at the cost of more compute time for each NE step. It is seen from the table that the multilayer approach results in fewer NE steps and fewer global Newton steps than the single layer approach. For a small Reynolds number Re = 2,000, the single layer approach is slightly better in terms of the total compute time; however, as RE increases, it requires more total compute time to converge and even diverges for cases Re > 3,200. In contrast, the multilayer approach requires a smaller total compute time at Re = 3,200 and ensures the convergence for cases Re = 5,000 and Re = 7,500.

Figure 14 shows the step length λ^k with respect to the global Newton step k. The application of multilayer NE results in $\lambda^k = 1$ for almost every Newton step, except for the most difficult case Re = 7,500. The ability to restore the full step length along the Newton direction s^k implies fast convergence of the Newton iteration. From Table 5 and Figures 13–14, we see that the proposed multilayer approach is superior to the single layer approach in terms of the robustness and efficiency for cases with high Reynolds numbers. We also remark that the field-based approach [59, 60, 61] fails to
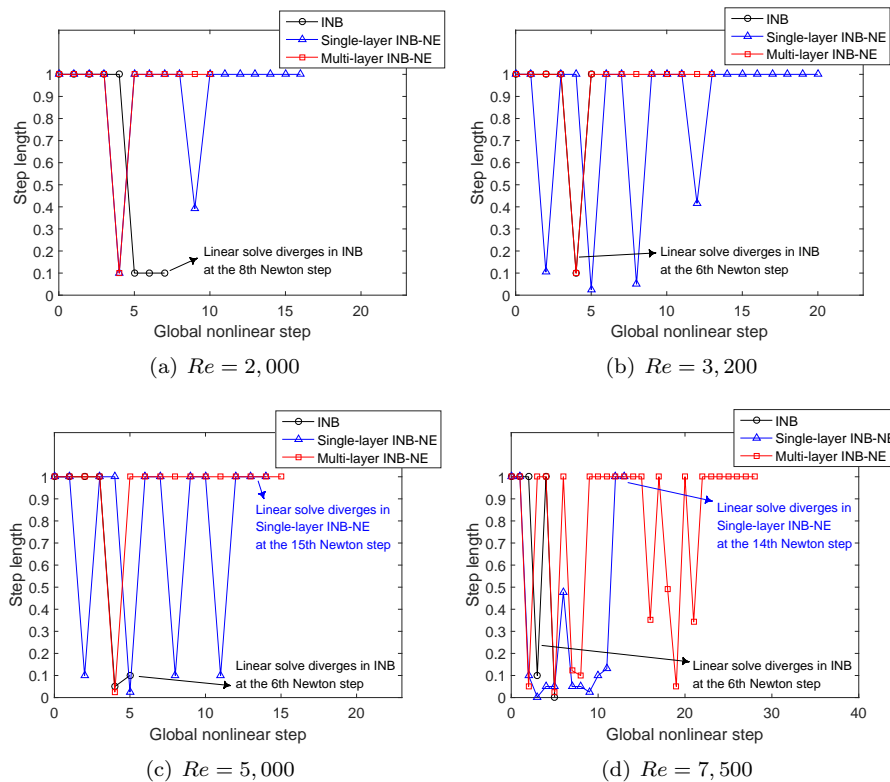
(a) $Re = 2,000$

(b) $Re = 3,200$

(c) $Re = 5,000$

(d) $Re = 7,500$

FIG. 14. *Step length λ^k obtained using INB and INB-NE with the single layer approach and the multilayer approach for the cavity flow with SAR $0.5:1$. The mesh size is $h = 1/256$.*

converge for all these cases by choosing either the velocity or the pressure as the field to eliminate.

4.5. Study of the parallel scalability. Finally, we report the parallel performance of the proposed multilayer INB-NE method in this subsection. The cavity flow with SAR 0.5:1 at Re = 2,000 is tested using two fixed mesh sizes $h = 1/160$ and $h = 1/256$. To obtain the best performance, we choose $(N_l, \beta_l) = (7, 0.5 \times 10^{-l})$, $l = 0, \ldots, N_l - 1$, and use ILU(2) as the subdomain solver. The relative tolerance for the global nonlinear solver is given as $\gamma_r = 10^{-12}$. We consider two scalability tests with the change of the number of processor cores np: (1) the strong scalability where the overall problem size is fixed; and (2) the weak scalability where the number of unknowns per processor core is fixed. Table 6 shows the scalability results in detail. Figure 15 shows the change of total compute times with respect to np. We see that as np increases, the numbers of global Newton and GMRES iterations remain stable, and the total compute times decrease correspondingly. The results show reasonably good performance in terms of the strong scalability. On the other hand, when fixing the number of unknowns per processor core (i.e., when $np = 256$ for the case $h = 1/160$ and $np = 1,024$ for the case $h = 1/256$), the total compute times are kept essentially unchanged as np has a fourfold increase, resulting in good performance in terms of the weak scalability.

TABLE 6

Scalability results of the multilayer INB-NE for the cavity flow with SAR $0.5:1$ at $Re = 2,000$. $N_l = 7$, $\beta_l = 0.5 \times 10^{-l}$, $l = 0,\ldots,N_l - 1$. ILU(2) is used as the subdomain solver. The relative tolerance for the global nonlinear solver is $\gamma_r = 10^{-12}$.

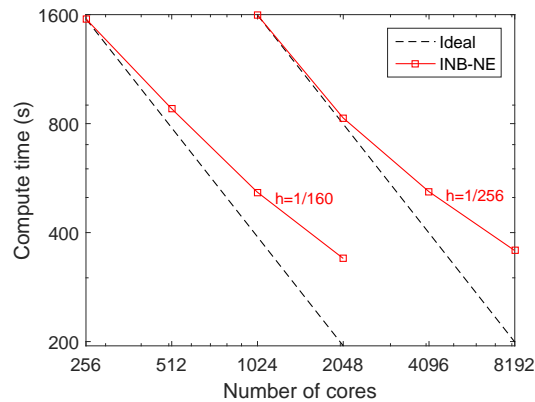| np	NI_g	LI_g	$T_t(s)$	N_{ne}	NI_{sub}	LI_{sub}	$T_{ne}(s)$
\multicolumn{8}{c}{$h = 1/160$, #unknowns=8,398,404}								
256	10	74.9	1,557.11	2	4.4	3.7	1,258.8
512	12	72.0	878.91	3	3.8	4.7	686.98
1,024	10	86.0	516.86	2	4.4	4.1	423.61
2,048	11	101.5	339.33	3	4.0	5.3	273.42
\multicolumn{8}{c}{$h = 1/256$, #unknowns=34,081,284}								
1,024	13	123.4	1,595.51	2	3.9	22.9	1,051.51
2,048	10	151.7	828.89	2	3.9	8.7	579.18
4,096	10	153.4	518.84	2	4.1	12.7	374.65
8,192	11	183.2	356.69	2	3.8	12.2	248.21

FIG. 15. *Scalability results of the multilayer INB-NE for the cavity flow with SAR $0.5:1$ at $Re = 2,000$. $N_l = 7$, $\beta_l = 0.5 \times 10^{-l}$, $l = 0,\ldots,N_l - 1$. ILU(2) is used as the subdomain solver. The relative tolerance for the global nonlinear solver is $\gamma_r = 10^{-12}$.*

5. Concluding remarks. The 3D steady-state driven cavity flow problem is defined on a simple computational domain, but it is actually one of the most difficult problems for algebraic solvers when the Reynolds number is high. When a classical INB is used to solve the resulting nonlinear system, it fails to converge when Re is larger than 1,739.5. When a single layer INB-NE is used it converges well for Re $\leq 3,200$. In this paper, we proposed a multilayer NE preconditioner. The key idea is to perform sweeps of subspace corrections in a cascadic manner to remove the local high nonlinearities that cause difficulty for the convergence of the inexact Newton method. We tested the algorithm using rectangular cavities with two SARs. Results of numerical experiments show that the proposed method is more robust and faster than the classical INB and the single layer INB-NE with respect to high Reynolds numbers $(3,200 \leq Re \leq 7,500)$, and is scalable on a supercomputer with over 8,000 processor cores. The focus of the paper was on the driven cavity flow problems, but the algorithm is expected to work for other highly nonlinear problems.

REFERENCES

[1] S. Albensoeder, H. C. Kuhlmann, and H. J. Rath, *Three-dimensional centrifugal-flow in-stabilities in the lid-driven cavity problem*, Phys. Fluids, 13 (2001), pp. 121–135.

[2] S. Albensoeder and H. C. Kuhlmann, *Accurate three-dimensional lid-driven cavity flow*, J. Comput. Phys., 206 (2005), pp. 536–558.

[3] V. F. de Almeida and J. J. Derby, *Construction of solution curves for large two-dimensional problems of steady-state flows of incompressible fluids*, SIAM J. Sci. Comput., 22 (2000), pp. 285–311.

[4] ANSYS, ANSYS CFX-Solver Theory Guide, Release 16.2, (2015), https://www.ansys.com/products/fluids/ansys-cfx.

[5] V. Babu and S. A. Korpela, *Numerical solutions of the incompressible, three-dimensional Navier-Stokes equations*, Comput. Fluids, 23 (1994), pp. 675–691.

[6] S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, D. A. May, L. C. McInnes, R. T. Mills, T. Munson, K. Rupp, P. Sanan, B. F. Smith, S. Zampini, H. Zhang, and H. Zhang, *PETSc Users Manual*, Argonne National Laboratory, Argonne, IL, 2019.

[7] F. Bassi, A. Crivellini, D. A. Di Pietro, and S. Rebay, *An artificial compressibility flux for the discontinuous Galerkin solution of the incompressible Navier-Stokes equations*, J. Comput. Phys., 218 (2006), pp. 794–815.

[8] M. Benzi, G. H. Golub, and J. Liesen, *Numerical solution of saddle point problems*, Acta Numer., 14 (2005), pp. 1–137.

[9] F. A. Bornemann and P. Deuflhard, *The cascadic multigrid method for elliptic problems*, Numer. Math., 75 (1996), pp. 135–152.

[10] L. Botti and D. A. Di Pietro, *A pressure-correction scheme for convection-dominated in-compressible flows with discontinuous velocity and continuous pressure*, J. Comput. Phys., 230 (2011), pp. 572–585.

[11] X.-C. Cai, W. D. Gropp, D. E. Keyes, R. G. Melvin, and D. P. Young, *Parallel Newton–Krylov–Schwarz algorithms for the transonic full potential equation*, SIAM J. Sci. Comput., 19 (1998), pp. 246–265.

[12] X.-C. Cai and M. Sarkis, *A restricted additive Schwarz preconditioner for general sparse linear systems*, SIAM J. Sci. Comput., 21 (1999), pp. 792–797.

[13] X.-C. Cai and D. E. Keyes, *Nonlinearly preconditioned inexact Newton algorithms*, SIAM J. Sci. Comput., 24 (2002), pp. 183–200.

[14] X.-C. Cai and X. Li, *Inexact Newton methods with restricted additive Schwarz based nonlinear elimination for problems with high local nonlinearity*, SIAM J. Sci. Comput., 33 (2011), pp. 746–762.

[15] T. S. Coffey, C. T. Kelley, and D. E. Keyes, *Pseudotransient continuation and differential-algebraic equations*, SIAM J. Sci. Comput., 25 (2003), pp. 553–569.

[16] A. B. Cortes and J. D. Miller, *Numerical experiments with the lid-driven cavity flow prob-lem*, Comput. Fluids, 23 (1994), pp. 1005–1027.

[17] J. E. Dennis and R. B. Schnabel, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, Classics Appl. Math. 16, SIAM, Philadelphia, 1996.

[18] M. Deville, T.-H. Lê, and Y. Morchoisne, *Numerical Simulation of 3-D Incompressible Un-steady Viscous Laminar Flows*, Notes on Numerical Fluid Mechanics 36, Vieweg, Braun-schweig, Germany, 1992.

[19] H. Ding, C. Shu, K. S. Yeo, and D. Xu, *Numerical computation of three-dimensional in-compressible viscous flows in the primitive variable form by local multiquadric differential quadrature method*, Comput. Methods Appl. Mech. Engrg., 195 (2006), pp. 516–533.

[20] V. Dolean, M. J. Gander, W. Kheriji, F. Kwok, and R. Masson, *Nonlinear precondition-ing: How to use a nonlinear Schwarz method to precondition Newton's method*, SIAM J. Sci. Comput., 38 (2016), pp. 3357–3380.

[21] S. C. Eisenstat and H. F. Walker, *Choosing the forcing terms in an inexact Newton method*, SIAM J. Sci. Comput., 17 (1996), pp. 16–32.

[22] E. Erturk, T. C. Corke, and C. Gokcol, *Numerical solutions of 2-D steady incompressible driven cavity flow at high Reynolds numbers*, Internat. J. Numer. Method Fluids, 48 (2005), pp. 747–774.

[23] E. Erturk, *Discussions on driven cavity flow*, Internat. J. Numer. Method Fluids, 60 (2009), pp. 275–294.

[24] Y. Feldman and A. Y. Gelfgat, *On pressure-velocity coupled time-integration of incom-pressible Navier-Stokes equations using direct inversion of Stokes operator or accelerated multigrid technique*, Comput. Struct., 87 (2009), pp. 710–720.

[25] Y. FELDMAN AND A. Y. GELFGAT, *Oscillatory instability of a 3D lid-driven flow in a cube*, Phys. Fluids, 22 (2010), 093602.

[26] Y. FELDMAN AND A. Y. GELFGAT, *From multi- to single-grid CFD on massively parallel computers: Numerical experiments on lid-driven flow in a cube using pressure-velocity coupled formulation*, Comput. Fluids, 46 (2011), pp. 218–223.

[27] A. Y. GELFGAT, *Linear instability of the lid-driven flow in a cubic cavity*, Theor. Comput. Fluid Dyn., 33 (2019), pp. 59–82.

[28] U. GHIA, K. N. GHIA, AND C. T. SHIN, *High-Re solutions for incompressible flow using the Navier-Stokes equations and a multigrid method*, J. Comput. Phys., 48 (1982), pp. 387–411.

[29] F. GIANNETTI, P. LUCHINI, AND L. MARINO, *Linear stability analysis of three-dimensional lid-driven cavity flow*, in Atti del XIX Congresso AIMETA di Meccanica Teorica e Applicata, Aras Edizioni, Ancona, Italy, 2009, pp. 14–17.

[30] S. GONG AND X.-C. CAI, *A nonlinear elimination preconditioned inexact Newton method for heterogeneous hyperelasticity*, SIAM J. Sci. Comput., 41 (2019), pp. S390–S408.

[31] W. D. GROPP, D. K. KAUSHIK, D. E. KEYES, AND B. F. SMITH, *High-performance parallel implicit CFD*, Parallel Comput., 27 (2001), pp. 337–362.

[32] E. HACHEM, B. RIVAUX, T. KLOCZKO, H. DIGONNET, AND T. COUPEZ, *Stabilized finite element method for incompressible flows with high Reynolds number*, J. Comput. Phys., 229 (2010), pp. 8643–8665.

[33] J. HUANG, C. YANG, AND X.-C. CAI, *A nonlinearly preconditioned inexact Newton algorithm for steady state lattice Boltzmann equations*, SIAM J. Sci. Comput., 38 (2016), pp. A1701–A1724.

[34] F.-N. HWANG AND X.-C. CAI, *A parallel nonlinear additive Schwarz preconditioned inexact Newton algorithm for incompressible Navier-Stokes equations*, J. Comput. Phys., 204 (2005), pp. 666–691.

[35] F.-N. HWANG AND X.-C. CAI, *A class of parallel two-level nonlinear Schwarz preconditioned inexact Newton algorithms*, Comput. Methods Appl. Mech. Engrg., 196 (2007), pp. 1603–1611.

[36] F.-N. HWANG, H.-L. LIN, AND X.-C. CAI, *Two-level nonlinear elimination-based preconditioners for inexact Newton methods with application in shoched duct flow calculation*, Electron. Trans. Numer. Anal., 37 (2010), pp. 239–251.

[37] F.-N. HWANG, Y.-C. SU, AND X.-C. CAI, *A parallel adaptive nonlinear elimination preconditioned inexact Newton method for transonic full potential equation*, Comput. Fluids, 110 (2015), pp. 96–107.

[38] B. N. JIANG, T. L. LIN, AND L. A. POVINELLI, *Large-scale computation of incompressible viscous flow by least-squares finite element method*, Comput. Methods Appl. Mech. Engrg., 114 (1994), pp. 213–231.

[39] G. KARYPIS, Family of Graph and Hypergraph Partitioning Software, http://glaros.dtc.umn. edu/gkhome/views/metis (2019).

[40] D. A. KNOLL AND D. E. KEYES, *Jacobian-free Newton-Krylov methods: A survey of approaches and applications*, J. Comput. Phys., 193 (2004), pp. 357–397.

[41] J. R. KOSEFF AND R. L. STREET, *On end wall effects in a lid-driven cavity flow*, J. Fluids Eng., 106 (1984), pp. 385–389.

[42] H. C. KU, R. S. HIRSH, AND T. D. TAYLOR, *A pseudospectral method for solution of the three-dimensional incompressible Navier-Stokes equations*, J. Comput. Phys., 70 (1987), pp. 439–462.

[43] H. C. KUHLMANN AND S. ALBENSOEDER, *Stability of the steady three-dimensional lid-driven flow in a cube and the supercritical flow dynamics*, Phys. Fluids, 26 (2014), 024104.

[44] H. C. KUHLMANN AND F. ROMANO, *The lid-driven cavity*, in Computational Modelling of Bifurcations and Instabilities in Fluid Dynamics, A. Gelfgat, ed., Springer, Cham, Switzerland, 2018, pp. 233–310.

[45] W. LAYTON, H. K. LEE, AND J. PETERSON, *Numerical solution of the stationary Navier–Stokes equations using a multilevel finite element method*, SIAM J. Sci. Comput., 20 (1998), pp. 1–12.

[46] L. LIU AND D. E. KEYES, *Field-split preconditioned inexact Newton algorithms*, SIAM J. Sci. Comput., 37 (2015), pp. A1388–A1409.

[47] L. LIU AND D. E. KEYES, *Convergence analysis for the multiplicative Schwarz preconditioned inexact Newton algorithm*, SIAM J. Numer. Anal., 54 (2016), pp. 3145–3166.

[48] L. LIU, D. E. KEYES, AND R. KRAUSE, *A note on adaptive nonlinear preconditioning techniques*, SIAM J. Sci. Comput., 40 (2018), pp. A1171–A1186.

[49] L. LUO, W.-S. SHIU, R. CHEN, AND X.-C. CAI, *A nonlinear elimination preconditioned inexact Newton method for blood flow problems in human artery with stenosis*, J. Comput. Phys., 399 (2019), 108926.

[50] L. Luo, R. Chen, X.-C. Cai, and D. E. Keyes, *A Nonlinear Elimination Preconditioned Inexact Newton Algorithm for Steady State Incompressible Flow Problems on* 3D *Unstructured Meshes*, Proceedings of Domain Decomposition Methods in Science and Engineering XXV, Springer Nature, 2020.

[51] S. J. Owen and J. F. Shepherd, CUBIT, http://cubit.sandia.gov (2019).

[52] A. K. Prasad and J. R. Koseff, *Reynolds number and end-wall effects on a lid-driven cavity flow*, Phys. Fluids A, 1 (1989), pp. 208–218.

[53] Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd ed., SIAM, Philadelphia, 2003.

[54] J. N. Shadid, R. S. Tuminaro, and H. F. Walker, *An inexact Newton method for fully coupled solution of the Navier-Stokes equations with heat and mass transport*, J. Comput. Phys., 137 (1997), pp. 155–185.

[55] P. N. Shankar and M. D. Deshpande, *Fluid mechanics in the driven cavity*, Annu. Rev. Fluid Mech., 32 (2000), pp. 93–136.

[56] T. W. H. Sheu and S. F. Tsai, *Flow topology in a steady three-dimensional lid-driven cavity*, Comput. Fluids, 31 (2002), pp. 911–934.

[57] E. M. Wahba, *Steady flow simulations inside a driven cavity up to Reynolds number* 35000, Comput. Fluids, 66 (2012), pp. 85–97.

[58] C. H. Whiting and K. E. Jansen, *A stabilized finite element method for the incompressible Navier-Stokes equations using a hierarchical basis*, Internat. J. Numer. Methods Fluids, 35 (2001), pp. 93–116.

[59] H. Yang, F.-N. Hwang, and X.-C. Cai, *Nonlinear preconditioning techniques for full-space Lagrange–Newton solution of PDE-constrained optimization problems*, SIAM J. Sci. Comput., 38 (2016), pp. A2756–A2778.

[60] H. Yang, C. Yang, and S. Sun, *Active-set reduced-space methods with nonlinear elimination for two-phase flow problems in porous media*, SIAM J. Sci. Comput., 38 (2016), pp. B593–B618.

[61] H. Yang, S. Sun, and C. Yang, *Nonlinearly preconditioned semismooth Newton methods for variational inequality solution of two-phase flow in porous media*, J. Comput. Phys., 332 (2017), pp. 1–20.

[62] H. Yang and F.-N. Hwang, *An adaptive nonlinear elimination preconditioned inexact Newton algorithm for highly local nonlinear multicomponent PDE systems*, Appl. Numer. Math., 133 (2018), pp. 100–115.

[63] H. Yu, L.-S. Luo, and S. S. Girimaji, *LES of turbulent square jet flow using an MRT lattice Boltzmann model*, Comput. Fluids, 35 (2006), pp. 957–965.