



Published in final edited form as:

Comput Mech. 2019 September 15; 64: 717–739. doi:10.1007/s00466-019-01678-3.

Performance of preconditioned iterative linear solvers for cardiovascular simulations in rigid and deformable vessels

Jongmin Seo,

Department of Pediatrics and Institute for Computational and Mathematical Engineering(ICME), Stanford University, Stanford, CA, USA, jongminseo@stanford.edu

Daniele E. Schiavazzi,

Department of Applied and Computational Mathematics and Statistics, University of Notre Dame, IN, USA, dschiavazzi@nd.edu

Alison L. Marsden

Department of Pediatrics, Bioengineering and ICME, Stanford University, Stanford, CA, USA, amarsden@stanford.edu

Abstract

Computing the solution of linear systems of equations is invariably the most time consuming task in the numerical solutions of PDEs in many fields of computational science. In this study, we focus on the numerical simulation of cardiovascular hemodynamics with rigid and deformable walls, discretized in space and time through the variational multiscale finite element method. We focus on three approaches: the problem agnostic generalized minimum residual (GMRES) and stabilized bi-conjugate gradient (BICGS) methods, and a recently proposed, problem specific, bi-partitioned (BIPN) method. We also perform a comparative analysis of several preconditioners, including diagonal, block-diagonal, incomplete factorization, multigrid, and resistance based methods. Solver performance and matrix characteristics (diagonal dominance, symmetry, sparsity, bandwidth and spectral properties) are first examined for an idealized cylindrical geometry with physiologic boundary conditions and then successively tested on several patient-specific anatomies representative of realistic cardiovascular simulation problems. Incomplete factorization preconditioners provide the best performance and results in terms of both strong and weak scalability. The BIPN method was found to outperform other methods in patient-specific models with rigid walls. In models with deformable walls, BIPN was outperformed by BICG with diagonal and Incomplete LU preconditioners.

Keywords

Cardiovascular simulation; Iterative linear solvers; Preconditioning; Fluid-structure interaction

1 Introduction

Cardiovascular simulations are increasingly used in clinical decision making, surgical planning, and medical device design. In this context, numerous modeling approaches have been proposed ranging from lumped parameter descriptions of the circulatory system to fully three-dimensional patient-specific representations. Patient-specific models are

generated through a pipeline progressing from segmentation of medical image data, to branch lofting, boolean union, application of physiologic boundary conditions tuned to match patient data, and hemodynamics simulation. In diseased vessels, e.g., characterized by localized stenosis or aneurysms, computational fluid dynamics (CFD) has been widely used to diagnose important clinical indicators, such as pressure drop or flow reduction [34]. Measures of shear stress on the vessel lumen have also been correlated with the risk of endothelial damage and thrombus formation [33, 26]. These quantities are determined by discretization in space and time of the incompressible Navier Stokes equations. Multiscale models have been developed to simulate the local flow field in three-dimensional patient-specific anatomies, while accounting for the presence of the peripheral circulation through closed-loop circuit models providing time dependent boundary conditions [18, 7, 23, 25]. In addition, several approaches for fluid-structure interaction (FSI) have been suggested to account for vessel wall deformability [13, 20, 27]. Recently, hemodynamic models have been used in the solution of complex problems in optimization [24, 22] and uncertainty quantification [29, 30, 32, 31, 36].

Efforts to improve realism in numerical simulations, however, often lead to an increase in the computational cost. Implementation of the coupled-momentum method (CMM) for FSI increases the simulation run time by roughly twice compared to rigid wall assumptions [13], and the cost can be substantially higher for Arbitrary Lagrangian and Eulerian (ALE) FSI. On top of this, optimization and uncertainty quantification studies often require the solution of a large number of simulations to obtain converged solutions. These requirements all point to a pressing need to reduce the computational cost to enable future integration of these tools in the clinical setting.

Preconditioned iterative approaches are widely used to solve linear systems, $\mathbf{A}\mathbf{y} = \mathbf{b}$, resulting from discretizations using variational multiscale finite element methods. However, few studies in the literature have examined in detail how linear solver performance depends on the properties of the coefficient matrix, to provide a concrete grounding for the choice and development of more efficient solvers. In addition, even fewer studies have carried out this analysis in the context of computational hemodynamics, i.e., the specific geometries, boundary conditions, mesh and material properties used to create numerical approximations of blood flow in rigid and deformable vessels. In this study, we investigate reductions in computational cost achievable through solving the discretized linear system efficiently, as this cost is well known to dominate the execution time. The objective of this study is to perform a systematic comparative analysis of existing linear solver strategies, linking their performance with the distributed coefficient matrix characteristics of cardiovascular modeling.

Krylov subspace based iterative solvers are typically preferred for the solution of large linear systems from CFD, due to their superior scalability and memory requirements compared to direct methods [4]. Popular Krylov subspace iterative solvers include the conjugate gradient method (CG) for symmetric positive definite (SPD) coefficient matrices, and the generalized minimum residual method (GMRES) or the bi-conjugate gradient stabilized method (BICGS) in the non-symmetric case. Alternatively, a recently proposed bi-partitioned linear solver (BIPN) [11] leverages the block-structure in the coefficient matrix, separating

contributions from the Navier-Stokes momentum and continuity equations. In BIPN, the coefficient matrix \mathbf{A}_f arising from finite element spatial discretizations and the time discretization for the Navier-Stokes equations consists of four blocks,

$$\mathbf{A}_f = \begin{bmatrix} \mathbf{K} & \mathbf{G} \\ \mathbf{D} & \mathbf{L} \end{bmatrix}, \quad (1)$$

in which \mathbf{K} and \mathbf{G} stem from the momentum equation, \mathbf{D} and \mathbf{L} stem from the continuity equation and stabilization. BIPN solves the matrix block \mathbf{K} using GMRES, while the rest is transformed to a Schur complement form as $\mathbf{L} - \mathbf{D}\mathbf{K}^{-1}\mathbf{G}$, approximated by a SPD matrix, $\mathbf{L}_* + \mathbf{G}_*^T\mathbf{G}_*$, in which the star subscript indicates the symmetric Jacobi scaling with diagonals of \mathbf{K} and \mathbf{L} . The Schur complement form is solved with CG, and the solution time for CG takes more than 90% of the total computing time in benchmark testing [11].

It is also well known that preconditioning plays a key role in accelerating the convergence of Krylov subspace methods [37, 28] by transforming the original linear system $\mathbf{A}\mathbf{y} = \mathbf{b}$ to $\mathbf{M}^{-1}\mathbf{A}\mathbf{y} = \mathbf{M}^{-1}\mathbf{b}$ (left preconditioning), $\mathbf{A}\mathbf{M}^{-1}\mathbf{z} = \mathbf{b}$, $\mathbf{y} = \mathbf{M}^{-1}\mathbf{z}$ (right preconditioning), or $\mathbf{M}_1^{-1}\mathbf{A}\mathbf{M}_2^{-1}\mathbf{y} = \mathbf{M}_1^{-1}\mathbf{b}$, $\mathbf{x} = \mathbf{M}_2^{-1}\mathbf{y}$ (left and right preconditioning). In many cases, \mathbf{M} is constructed in a way that \mathbf{M}^{-1} approximates \mathbf{A}^{-1} . In general, an ideal preconditioner should be relatively cheap to apply and effective to reduce the overall solution time. In its simplest form, left, right or left-right Jacobi (diagonal) preconditioners are effective in shrinking the eigenvalue spectrum of diagonally dominant matrices. Preconditioners based instead on incomplete factorization (ILU) provide an approximate decomposition in the form $\mathbf{M} = \bar{\mathbf{L}}\bar{\mathbf{U}}$, where the sparsity pattern of \mathbf{A} is preserved in the factors. ILUT preconditioners are slightly more general approaches allowing for adjustable inclusion of fill-ins, but require the user to specify an additional threshold parameter. We note that the efficiency of an ILU preconditioner results from a trade off between fewer Krylov iterations needed for convergence and the cost of incomplete factorization [28]. Application-specific preconditioners have also been proposed in cardiovascular hemodynamics to improve performance when the model outlets are coupled through a resistance boundary condition, an RCR circuit, or for more general multi-domain configurations. In what follows, we refer to an in-house implementation of this class of preconditioners as *resistance-based preconditioner* (RPC) [10, 11]. Additional preconditioning techniques for cardiovascular simulations with FSI are suggested in [21, 8]. Finally, algebraic multigrid preconditioners have also received significant recent interest [9].

Despite the availability of open-source implementations of iterative solvers and preconditioners, few studies have systematically compared the performance of these solvers for cardiovascular models with rigid and deformable vessels. In addition, a thorough understanding of the factors affecting the performance of iterative linear solvers (e.g., diagonal dominance, condition number, sparsity pattern, symmetry, and positive-definiteness) is an important prerequisite for optimal choice of solver and for the development of new algorithms with improved performance.

In the current study, we first compare the performance of various iterative linear solvers and preconditioners for an idealized flow through a cylindrical vessel with a resistance outflow boundary condition. We then test our findings using three representative patient-specific cardiovascular models. The Trilinos software library [14], developed at Sandia National Laboratory is coupled with the SimVascular svFSI open source finite element code to provide linear solvers such as GMRES and BICGS, as well as a variety of preconditioners: diagonal (Diag), block-diagonal (BlockD), incomplete LU (ILU), thresholded incomplete LU (ILUT), incomplete Cholesky (IC), and algebraic multigrid (ML). We use Kylov linear solvers and Diag, BlockD, ILU, ILUT preconditioners from the AztecOO package, and IC is implemented via the IF-PACK package. The block-diagonal preconditioner scales the block matrix via the Trilinos Epetra Vbr class. The incomplete factorization methods use Additive Schwarz domain decomposition for the parallelization. The Trilinos ML package for algebraic multigrid is implemented on the AztecOO package. For detailed information on parallelization and preconditioning options, we refer readers to the Trilinos project [14]. BIPN and RPC are integrated and implemented directly in our flow solver with source code available through the SimVascular open source project (www.simvascular.org) [38].

This paper is organized as follows. In section 2, we review the formulation of the coefficient matrices resulting from finite element weak forms in fluid and solid mechanics, before discussing the performance of various linear solvers and preconditioners on a simple pipe benchmark in section 3. In section 4 we report the results of strong and weak scaling for BIPN with RPC and BICG with ILU, while in section 5 we examine the properties of the coefficient matrix. The effect of preconditioning on these properties is reported in section 6. In section 7 we compare performance of linear solvers in patient-specific models. We draw conclusions and discuss future work in section 8.

2 Linear systems in cardiovascular simulation

We begin by introducing the space-time discretization of the equations governing fluid and solid mechanics following an Arbitrary-Lagrangian-Eulerian (ALE) description of the interaction between fluid and structure [15, 3, 11, 39]. These equations are discretized with a variational multiscale finite element method, and are provided in the svFSI solver of the SimVascular open source project [38].

2.1 Linear system for fluid mechanics

Consider a domain $\Omega_f \in \mathbb{R}^3$, occupied by a Newtonian fluid whose evolution in space and time is modeled through the incompressible Navier-Stokes equations in ALE form,

$$\begin{aligned} \rho \frac{\partial \mathbf{u}}{\partial t} \Big|_{\hat{\mathbf{x}}} + \rho \mathbf{v} \cdot \nabla \mathbf{u} &= \rho \mathbf{f} + \nabla \cdot \boldsymbol{\sigma}_f \quad \text{in } \Omega_f, \\ \nabla \cdot \mathbf{u} &= 0 \end{aligned} \quad (2)$$

where ρ , $\mathbf{u} = \mathbf{u}(\mathbf{x}, t)$, and \mathbf{f} are fluid density, velocity vector, and body force, respectively. The fluid stress tensor is $\boldsymbol{\sigma}_f = -p\mathbf{I} + \mu(\nabla \mathbf{u} + \nabla \mathbf{u}^T) = -p\mathbf{I} + \mu \nabla^s \mathbf{u}$, μ is the kinematic viscosity, $p = p(\mathbf{x}, t)$ the pressure, $\mathbf{v} = \mathbf{u} - \hat{\mathbf{u}}$ is the fluid velocity relative to the velocity of the domain boundary $\hat{\mathbf{u}}$. Variables are interpolated in space at time t^n as

$$\begin{aligned}\mathbf{w}(\mathbf{x}) &= \sum_{a \in I_a} N^a(\mathbf{x}) \mathbf{w}^a, \\ \mathbf{q}(\mathbf{x}) &= \sum_{a \in I_a} N^a(\mathbf{x}) \mathbf{q}^a,\end{aligned}\tag{3}$$

$$\begin{aligned}\mathbf{u}(\mathbf{x}, t = t^n) = \mathbf{u}^n(\mathbf{x}) &= \sum_{a \in I_a} N^a(\mathbf{x}) \mathbf{u}^{a,n}, \\ p(\mathbf{x}, t = t^n) = p^n(\mathbf{x}) &= \sum_{a \in I_a} N^a(\mathbf{x}) p^{a,n},\end{aligned}\tag{4}$$

in which \mathbf{I}_a , N^a , w^a , q^a , \mathbf{u}^a , and p^a are the nodal connectivity set, interpolation functions at node a , test function weights, velocity and pressure at node a , respectively. In this study, we employ P1-P1 type (linear and continuous) spatial approximations of the fluid velocity and pressure. We consider a stabilized finite element discretization based on the variational multiscale method [2, 11], leading to the weak form of the Navier-Stokes momentum and continuity residuals

$$\begin{aligned}\mathbf{R}_m^a(\dot{\mathbf{u}}, \mathbf{u}, \mathbf{v}, p) &= \\ & \sum_{e \in \mathbf{I}_e} \int_{\Omega^e} \rho N^a (\dot{\mathbf{u}} - \mathbf{f} + (\mathbf{v} + \mathbf{u}_p) \cdot \nabla \mathbf{u}) d\Omega \\ & + \sum_{e \in \mathbf{I}_e} \int_{\Omega^e} (\nabla N^a)^T (-p \mathbf{I} + \mu \nabla^s \mathbf{u} + \rho \tau_B \mathbf{u}_p \otimes (\mathbf{u}_p \cdot \nabla \mathbf{u}) \\ & \quad - \rho \mathbf{u} \otimes \mathbf{u}_p + \rho \tau_C \nabla \cdot \mathbf{u}) d\Omega - \int_{\Gamma_h} N^a \mathbf{h} d\Gamma,\end{aligned}\tag{5}$$

$$R_c^a(\dot{\mathbf{u}}, \mathbf{u}, p) = \int_{\Omega} (N^a \nabla \cdot \mathbf{u} - (\nabla N^a)^T \mathbf{u}_p) d\Omega,\tag{6}$$

in which \mathbf{R}_m^a and R_c^a are momentum and continuity residuals at node a , and \mathbf{h} is the surface traction on the Neumann boundary Γ^h . The stabilization parameters are defined as

$$\begin{aligned}
\mathbf{u}_p &= -\tau_M \left(\dot{\mathbf{u}} + \mathbf{v} \cdot \nabla \mathbf{u} + \frac{1}{\rho} \nabla p - \frac{\mu}{\rho} \nabla^2 \mathbf{u} - \mathbf{f} \right), \\
\tau_M &= \left(\frac{4}{\Delta t^2} + \mathbf{v} \cdot \mathbf{G} \mathbf{v} + C_I \left(\frac{\mu}{\rho} \right)^2 \mathbf{G} : \mathbf{G} \right)^{-1/2}, \\
\tau_B &= (\mathbf{u}_p \cdot \mathbf{G} \mathbf{u}_p)^{-1/2}, \quad \tau_C = (\tau_M \mathbf{g} \cdot \mathbf{g})^{-1}, \\
G_{ij} &= \sum_{k=1}^3 \frac{\partial \xi_k}{\partial x_i} \frac{\partial \xi_k}{\partial x_j}, \\
\mathbf{g} \cdot \mathbf{g} &= \sum_{i=1}^3 g_i g_i, \quad g_i = \sum_{k=1}^3 \frac{\partial \xi_k}{\partial x_i},
\end{aligned} \tag{7}$$

where C_I is a constant set to 3, Δt is the time step size, and ξ represents natural coordinates. Integration in time is performed using the unconditionally stable, second order accurate generalized- α method [16], consisting of four steps: predictor, initiator, Newton-Raphson, and corrector. Initial values for accelerations, velocities and pressures at time t^{n+1} are set in the prediction step as

$$\dot{\mathbf{u}}^{a,n+1} = \frac{\gamma-1}{\gamma} \dot{\mathbf{u}}^{a,n}, \quad \mathbf{u}^{a,n+1} = \mathbf{u}^{a,n}, \quad p^{a,n+1} = p^{a,n}, \tag{8}$$

where $\gamma = 0.5 + \alpha_m - \alpha_f$, $\alpha_m = 1/(1 + \rho_\infty)$, and $\alpha_f = (3 - \rho_\infty)/(2 + 2\rho_\infty)$ are the generalized- α method coefficients, while ρ_∞ is the spectral radius set to $\rho_\infty = 0.2$ in this study. In the initiator step, accelerations and velocities are computed at an intermediate stage $n + \alpha_m$ and $n + \alpha_f$

$$\begin{aligned}
\dot{\mathbf{u}}^{a,n+\alpha_m} &= (1 - \alpha_m) \dot{\mathbf{u}}^{a,n} + \alpha_m \dot{\mathbf{u}}^{a,n+1}, \\
\mathbf{u}^{a,n+\alpha_f} &= (1 - \alpha_f) \mathbf{u}^{a,n} + \alpha_f \mathbf{u}^{a,n+1}.
\end{aligned} \tag{9}$$

A Newton-Raphson iteration is performed based on Equations (5) and (6), using $\dot{\mathbf{u}}^{n+\alpha_m}$, $\mathbf{u}^{n+\alpha_f}$, p^{n+1} from (9) by solving a linear system of the form

$$\begin{aligned}
\mathbf{K} \Delta \mathbf{u} + \mathbf{G} \Delta \mathbf{p} &= -\mathbf{R}_m \left(\dot{\mathbf{u}}^{n+\alpha_m}, \mathbf{u}^{n+\alpha_f}, p^{n+1} \right), \\
\mathbf{D} \Delta \mathbf{u} + \mathbf{L} \Delta \mathbf{p} &= -\mathbf{R}_c \left(\dot{\mathbf{u}}^{n+\alpha_m}, \mathbf{u}^{n+\alpha_f}, p^{n+1} \right),
\end{aligned} \tag{10}$$

where the blocks \mathbf{K} , \mathbf{G} , \mathbf{D} , and \mathbf{L} partition the tangent coefficient matrix with blocks for nodes a and b equal to

$$\begin{aligned}\mathbf{K}^{ab} &\approx \frac{\partial \mathbf{R}_m^a}{\partial \Delta \mathbf{u}^b}, \mathbf{G}^{ab} \approx \frac{\partial \mathbf{R}_m^a}{\partial \Delta p^b}, \\ \mathbf{D}^{ab} &\approx \frac{\partial \mathbf{R}_c^a}{\partial \Delta \mathbf{u}^b}, \mathbf{L}^{ab} \approx \frac{\partial \mathbf{R}_c^a}{\partial \Delta p^b}.\end{aligned}\quad (11)$$

We re-write this linear system in matrix form as

$$\mathbf{A}_f \mathbf{y} = -\mathbf{R}_f, \quad (12)$$

where

$$\mathbf{A}_f = \begin{bmatrix} \mathbf{K} & \mathbf{G} \\ \mathbf{D} & \mathbf{L} \end{bmatrix}, \mathbf{y} = \begin{bmatrix} \Delta \mathbf{u} \\ \Delta \mathbf{p} \end{bmatrix}, \mathbf{R}_f = \begin{bmatrix} \mathbf{R}_m \\ \mathbf{R}_c \end{bmatrix}, \quad (13)$$

with blocks \mathbf{K} , \mathbf{G} , \mathbf{D} , \mathbf{L} of size $(3N_{nd} \times 3N_{nd})$, $(N_{nd} \times 3N_{nd})$, $(3N_{nd} \times N_{nd})$, $(N_{nd} \times N_{nd})$, respectively. Here N_{nd} is the total number of nodes, while $\Delta \mathbf{u} \in \mathbb{R}^{3N_{nd}}$ and $\Delta \mathbf{p} \in \mathbb{R}^{N_{nd}}$ contain nodal velocities and pressure increments. We note that the major focus of our study is on solving the linear system in equation (12). Once the momentum and continuity residual norms drop below a given tolerance, the unknowns at the next time step are determined through the corrections

$$\begin{aligned}\dot{\mathbf{u}}^{a,n+1} &\leftarrow \dot{\mathbf{u}}^{a,n+1} + \Delta \mathbf{u}^a, \\ \mathbf{u}^{a,n+1} &\leftarrow \mathbf{u}^{a,n+1} + \gamma \Delta t \Delta \mathbf{u}^a, \\ p^{a,n+1} &\leftarrow p^{a,n+1} + \alpha_f \gamma \Delta t \Delta p^a,\end{aligned}\quad (14)$$

$\forall a \in \mathbf{I}_q$. Finally, detailed expressions for each block of the coefficient matrix from Eq. (5), Eq. (6), Eq. (11), and Eq. (14) are

$$\begin{aligned}\mathbf{K}^{ab} &= \sum_{e \in \mathbf{I}_e} \int_{\Omega^e} \left[\rho \alpha_m N^a N^b \mathbf{I}^{ab} + \rho \tilde{\alpha}_f N^a (\mathbf{v} + \mathbf{u}_p) \cdot \nabla N^b \mathbf{I}^{ab} \right. \\ &\quad \left. + \mu \tilde{\alpha}_f (\nabla N^a \cdot \nabla N^b \mathbf{I}^{ab} + \nabla N^b \otimes \nabla N^a) \right. \\ &\quad \left. + \rho \tilde{\alpha}_f \tau_B \mathbf{u}_p \cdot \nabla N^a \mathbf{u}_p \cdot \nabla N^b \mathbf{I}^{ab} \right. \\ &\quad \left. + \rho \tau_M \mathbf{u} \cdot \nabla N^a (\alpha_m N^b + \tilde{\alpha}_f \mathbf{u} \cdot \nabla N^b) \mathbf{I}^{ab} \right. \\ &\quad \left. + \rho \tilde{\alpha}_f \tau_C \nabla N^a \otimes \nabla N^b \right] d\Omega,\end{aligned}\quad (15)$$

$$\mathbf{G}^{ab} = \sum_{e \in \mathbf{I}_e} \int_{\Omega^e} [-\tilde{\alpha}_f \nabla N^a N^b + \tilde{\alpha}_f \tau_M \mathbf{u} \cdot \nabla N^a \nabla N^b] d\Omega, \quad (16)$$

$$\begin{aligned} \mathbf{D}^{ab} = \sum_{e \in \mathbf{I}_e} \int_{\Omega^e} [\tilde{\alpha}_f N^a \nabla N^b + \tilde{\alpha}_f \tau_M \mathbf{u} \cdot \nabla N^a \nabla N^b \\ + \tau_M \nabla N^a \alpha_m N^b] d\Omega, \end{aligned} \quad (17)$$

$$\mathbf{L}^{ab} = \sum_{e \in \mathbf{I}_e} \int_{\Omega^e} \left[\frac{\tilde{\alpha}_f \tau_M}{\rho} \nabla N^a \cdot \nabla N^b \right] d\Omega, \quad (18)$$

in which $\tilde{\alpha}_f = \gamma \Delta t \alpha_f$ and \mathbf{I}_e is the list of elements containing nodes a and b .

First we observe that the matrix \mathbf{K} is diagonally dominant. Except for entries related to the stabilization terms, which are typically small, the most significant off-diagonal contribution is provided by the viscous term, which is also typically smaller than the acceleration and advection terms in cardiovascular flows. Second, the small magnitude of the stabilization terms suggests that \mathbf{G} is similar to $-\mathbf{D}^T$. We also observe that the matrices \mathbf{K} and \mathbf{A}_f are non-symmetric, while \mathbf{L} is symmetric and singular since it has an identical structure to the matrices arising from the discretization of generalized Laplace operators. We also note that \mathbf{L} is characterized by small entries compared to the other blocks, since it only consists of stabilization terms.

2.2 Linear system for solid mechanics

In the solid domain, we start by introducing measures of deformation induced by a displacement field $\mathbf{d} = \mathbf{x} - \mathbf{X}$, i.e., the difference between the current and material configurations $\mathbf{x} \in \mathbb{R}^3$ and $\mathbf{X} \in \mathbb{R}^3$, respectively

$$\mathbf{F} = \nabla \mathbf{d} + \mathbf{I}, \quad \mathbf{C} = \mathbf{F}^T \mathbf{F}, \quad \mathbf{E} = \frac{1}{2}(\mathbf{C} - \mathbf{I}), \quad (19)$$

where \mathbf{F} , \mathbf{C} , \mathbf{E} , represent the deformation gradient, the Cauchy-Green deformation tensor and the Green strain tensor. The Jacobian is also defined as $J = \det(\mathbf{F})$. We relate the second Piola-Kirchhoff stress tensor \mathbf{S} with the Green strain tensor \mathbf{E} through the Saint Venant-Kirchhoff hyperelastic constitutive model

$$\mathbf{S} = \lambda \operatorname{tr}(\mathbf{E}) \mathbf{I} + 2 \mu \mathbf{E}, \quad (20)$$

where $\lambda = \frac{\nu E_s}{(1+\nu)(1-2\nu)}$, $\mu = \frac{E_s}{2(1+\nu)}$, E_s and ν represent the Young's modulus and Poisson's ratio, respectively. The equilibrium equation is

$$\rho_s \frac{\partial \mathbf{u}}{\partial t} = \rho_s \mathbf{f} + \nabla \cdot \boldsymbol{\sigma}_s \quad \text{in } \Omega_s, \quad (21)$$

where ρ_s and $\boldsymbol{\sigma}_s$ denote the density and solid stress tensor, respectively. This leads to the weak form

$$\int_{\Omega_s^0} [\rho_s^0 \mathbf{w}(\dot{\mathbf{u}} - \mathbf{f}) + \nabla \mathbf{w} : \mathbf{P}] d\Omega = \mathbf{0}, \quad (22)$$

where $\mathbf{P} = \mathbf{F}\mathbf{S}$ is the first Piola-Kirchhoff stress tensor, \mathbf{w} is a virtual displacement and Ω_s^0 is the solid domain in the *reference* configuration. Discretization of (22) leads to the residual

$$\mathbf{R}_m^a(\dot{\mathbf{u}}, \mathbf{d}) = \int_{\Omega_s^0} [\rho_s^0 N^a (\dot{\mathbf{u}} - \mathbf{f}) + \mathbf{F}\mathbf{S} \nabla N^a] d\Omega. \quad (23)$$

Using the generalized- α method, the displacements at time t^{n+1} are predicted as

$$\mathbf{d}^{a,n+1} = \mathbf{d}^{a,n} + \mathbf{u}^{a,n+1} \Delta t + \frac{0.5\gamma - \beta}{\gamma - 1} \dot{\mathbf{u}}^{a,n+1} \Delta t^2, \quad (24)$$

in which $\beta = \frac{1}{4}(1 + \alpha_f - \alpha_m)^2$. In the initiator step, the intermediate displacements are provided by

$$\mathbf{d}^{a,n+\alpha_f} = (1 - \alpha_f) \mathbf{d}^{a,n} + \alpha_f \mathbf{d}^{a,n+1}. \quad (25)$$

Solving (23) with the Newton-Raphson method, we obtain the linear system

$$\mathbf{K}_s \Delta \mathbf{d} = -\mathbf{R}_m(\dot{\mathbf{u}}^{n+\alpha_m}, \mathbf{d}^{n+\alpha_f}), \quad (26)$$

where $\mathbf{K}_s^{ab} \approx \frac{\partial \mathbf{R}_m^a}{\partial \Delta d^b}$, with tangent stiffness matrix

$$\begin{aligned} \mathbf{K}_s^{ab} = & \int_{\Omega_s^0} [\rho_s^0 \alpha_m N^a N^b \mathbf{I} + \hat{\alpha}_f (\mathbf{S} \nabla N^a) \cdot \nabla N^b \mathbf{I} \\ & + \lambda \hat{\alpha}_f (\mathbf{F} \nabla N^a) \otimes (\mathbf{F} \nabla N^b) + \mu \hat{\alpha}_f (\mathbf{F} \nabla N^b) \otimes (\mathbf{F} \nabla N^a) \\ & + \mu \hat{\alpha}_f \mathbf{F} \mathbf{F}^T \nabla N^a \cdot \nabla N^b] d\Omega, \end{aligned} \quad (27)$$

and $\hat{\alpha}_f = \alpha_f \beta \Delta t^2$. At this point, \mathbf{d} is corrected using

$$\mathbf{d}^{a,n+1} \leftarrow \mathbf{d}^{a,n+1} + \beta \Delta t^2 \Delta \mathbf{d}^a, \quad \forall a \in \mathbf{I}_a. \quad (28)$$

Again, we note that the focus of this work is on solving the linear system in equation (26) together with solving equation (12) when including FSI. Finally, we observe that most terms in (27) (i.e., the third, fourth, and fifth terms) contribute to the off-diagonals of \mathbf{K}_s , and that the matrix \mathbf{K}_s is symmetric.

3 Linear solver performance on pipe flow benchmark

All tests discussed in this study are performed using the svFSI finite element solver from the SimVascular open source project, leveraging the message passing interface (MPI) library and optimized to run efficiently on large computational clusters [12]. We note this implementation assigns fluid and solid elements to separate cores and assumes unique interface nodes, i.e., common nodes on the fluid and solid mesh are expected to match. An FSI mesh with matching interface nodes is generated through the freely-available Meshmixer software, with details reported in [39].

We consider a simple pipe benchmark whose size and boundary conditions are chosen to represent the ascending aorta of a healthy subject, having 4 cm diameter, 30 cm length and 0.2 cm thickness, assuming a radius/thickness ratio of 10%. The inflow is steady with a parabolic velocity profile, having a mean flow rate of $Q = 83$ mL/s (5 L/min). A resistance boundary condition is used at the outlet equal to $R = 1600$ g/cm⁴/s, which produces a mean pressure of approximately $P = 100$ mmHg, typical of the systemic circulation of a healthy subject. Simulations are performed with rigid and deformable walls (see Figure 1), using 1,002,436 tetrahedral elements for the fluid domain and 192,668 tetrahedral elements for the wall, generated in SimVascular with the TetGen mesh generator plugin [38, 19]. We measure the wall clock time on the XSEDE Comet cluster for simulations consisting of 10 time steps of 1 millisecond each, using 38 and 48 cores for rigid and deformable simulations, respectively. The XSEDE Comet cluster has 1,944 compute nodes with the Intel Xeon E5-2680v3 cores, 24 cores/node, 2.5GHz clock speed, 960 GFlop/s flop speed, and 120GB/s memory bandwidth. For more information about the Comet cluster, please refer to the XSEDE user portal. We use the restarting scheme for GMRES. We set the restart number of 200 which showed superior performance against smaller restart numbers (See Appendix A). We set the ILUT drop tolerance to 10^{-2} and the fill-in level to 2. In our test, changing the drop tolerance to 10^{-4} and 10^{-6} did not significantly change the performance of the linear solver reported here. For the multigrid preconditioner we selected a maximum level of four, a Gauss-Seidel smoother, and the symmetric Gauss-Seidel for the subsmoother. We confirmed that this setting was superior to other choices for smoother and subsmoother (See Appendix B). Finally, as the node ordering affects the amount of fill-in produced by an ILU decomposition, we apply Reverse-Cuthill-McKee (RCM) reordering prior to the incomplete factorization. RCM has been shown to be effective among many reordering schemes in the solution of non-symmetric sparse linear systems (see, e.g., [5]). From our testing, ILUT with RCM reordering provided superior performance against ILUT without reordering and ILUT with METIS reordering (See Appendix C).

3.1 Rigid wall benchmark

We plot the linear solver performance measured by wall clock time in Figure 2. In Table 2, we report the number of iterations and the portion of total compute time consumed by solving the linear system. Three iterative solver tolerances, $\epsilon = 10^{-3}$, 10^{-6} , 10^{-9} , are tested and compared. Table 1 shows the effect of tolerance on the velocities at $t = 10$ ms, suggesting the velocity error norm is of the same order of the selected tolerance.

Figure 2 shows that incomplete factorization preconditioners are fast and exhibit robust performance across all tolerances, irrespective of the underlying iterative linear solver. Despite similar performance for ILU, ILUT and IC preconditioners across all cases, the slightly worse performance of ILUT with respect to ILU seems to suggest that constructing a more accurate factor may lead to smaller run times than the savings in the factorization cost achieved by dropping additional fill-ins. GMRES with diagonal preconditioners (either diagonal or block-diagonal) appears to be significantly slower than other schemes, particularly as the tolerance ϵ becomes smaller. This degrading performance of standard GMRES for cardiovascular modeling is consistent with previous studies, showing that resistance boundary conditions are responsible for an increase in the condition number [10, 11]. While BICGS seems to perform better than GMRES with diagonal preconditioners, its performance becomes increasingly unstable with smaller tolerance ϵ . Furthermore, while algebraic multigrid preconditioners appear to be superior to diagonal preconditioning, they are inferior to BIPN or GMRES/BICGS with ILU. Finally, the performance of BIPN with RPC preconditioning is comparable to ILU for small tolerances (i.e., $\epsilon = 10^{-3}$) but significantly degrades as the tolerance value decreases.

From Table 2, we see that the time consumed by the linear system solvers constitutes the majority of the total compute time. In BIPN, the compute time of GMRES is significantly smaller than the compute time of CG. As the tolerance value decreases, the relative percentage of compute time for GMRES solve becomes larger. All Trilinos preconditioners show larger iteration numbers with decreasing tolerance. The relatively small percentage of linear solver compute time against the total compute time with IC and ML implies that building such preconditioners is expensive. This suggests that storing and reusing a preconditioner for several time steps could increase efficiency.

3.2 Deformable wall benchmark (FSI)

We illustrate the compute times for the deformable wall case in Figure 3 and summarize the number of iterations and percentage compute time of the linear solvers in Table 3. For FSI simulations with a tolerance $\epsilon = 10^{-3}$, BIPN with RPC shows more than an 8 fold increase in compute time compared to the rigid wall case, which also suggests the limitations of this approach for the deformable wall case. The increase of iteration number for GMRES as tolerance values decrease is notably higher than in the rigid wall case, and the percentage of compute time for the GMRES solve in BIPN is significantly higher, about 80% of the total compute time while the CG part continues to make up a smaller percentage. This suggests directions for future improvement of BIPN in the GMRES part for FSI simulations. Conversely, incomplete factorization preconditioners (both for GMRES and BICGS), exhibit good performance across all tolerances and a limited increase in compute time compared to

the rigid case. Diagonal preconditioners show a comparable performance to incomplete factorization schemes at large tolerances, but the performance degrades for smaller tolerances. Among all algorithms implemented in Trilinos, the algebraic multigrid preconditioners is the slowest, while BICG with ILU appears to be the best solution scheme overall.

4 Parallel scalability

Parallel scalability is investigated for two preconditioned linear solvers, BIPN-RPC and BICG-ILU, in terms of speedup (see, e.g., [6]), defined as the computing speed of multiple cores compared to a single core calculation, i.e., $S_p = T_1/T_p$, where T_p is the compute time on N_p cores. The ideal strong scalability performance between S_p and N_p is linear, however, sublinear scaling is expected due to communication cost.

4.1 Strong scaling

In this section, we monitor the compute time for a model with a fixed number of degrees of freedom, while progressively increasing the number of cores. We first test the strong scalability by varying the number of cores, testing 1, 2, 4, 8, 24, 48, and 96 cores for an ≈ 1 million element mesh (Mesh1 in Table 4) (Figure 4). We chose the number of cores as multiples of 24, to use all cores in any given node. We note, however, that one should use only between 2/3 and 3/4 of the total number of cores in a given node since the local memory bandwidth is often a bottleneck resulting in higher overhead and reduced speed improvement. In the rigid wall model, BIPN-RPC and BICG-ILU show similar performance across all core numbers as shown in Figure 2. The parallel speedup shows that both methods scale well up to 24 cores, leading to about 40,000 elements per core, while their parallel performance is reduced when running on more than 48 cores. In Figure 2, the zig-zag behavior of BIPN-RPC (FSI) reveals excessive inter-core communications and memory references. In the FSI problem, the total compute time of BIPN-RPC is larger, and the speed up is worse than BICG-ILU. BICG-ILU shows consistently good scalability for both rigid and FSI models whereas the scalability of BIPN-RPC degrades significantly in the FSI case.

We conduct an additional scaling study on a refined mesh with 8M elements (Mesh5 in Table 4). We use 24, 48, 96, 192, and 384 cores and we use the $N_p = 24$ case as a reference for S_p . As shown in Figure 5, the BICG-ILU speedup scales almost linearly up to 192 cores, i.e., $\sim 40,000$ elements per core. In the 8M mesh, the compute time of BIPN-RPC becomes significantly larger due to a poor weak scalability, as discussed in the next section.

4.2 Weak scaling

In this section, we solve models of increasing size and report the simulation time by keeping approximately the same number of elements per core. The number of nodes, elements and non-zeros in the coefficient matrix is summarized for each mesh in Table 4. We use 20, 40, 80, 160, 320 cores for meshes 1 to 5, resulting in 27,000 elements and about 4,500 nodes per a core.

BICG-ILU shows increased compute time but does not show any significant changes in scaling when increasing the number of cores. For the rigid model, BIPN-RPC shows poor

scalability as the number of cores exceeds 160. This explains the loss of superior performance of BIPN-RPC against BICG-ILU in the strong scalability study with 8M mesh. Similar to the rigid case, BIPN-RPC shows performance loss after 160 cores for the FSI model.

5 Characteristics of unpreconditioned matrices

In this section, we investigate the properties of the coefficient matrices discussed in section 2.1 to better understand the performance results. With reference to the pipe benchmark, we visualize the matrix sparsity pattern and investigate its properties including bandwidth, symmetry, diagonal dominance and spectrum. We also investigate the characteristics of both *global* and *local* matrices. The *global* matrix contains all mesh nodes from all cores, while a *local* matrix contains only a subset of the nodes in the *global* matrix assigned to a single core upon partitioning. We report single-core, *local* matrix characteristics with $N_{lnd} \sim 5000$ nodes to represent the typical case of distributed discretizations consisting of ~ 25000 tetrahedral elements per core. This way we focus on detailed *local* information in a region of specific interest (e.g. resistance boundary), and also calculate properties of the matrix such as eigenvalues and condition numbers in a cost-effective way. We discuss properties for two groups of coefficient matrices, i.e., matrices associated with fluid flow and matrices associated with solid mechanics.

5.1 Matrix properties for fluid flows in rigid vessels

Sparsity pattern.—The structure of \mathbf{A}_f is determined by element nodal connectivity, i.e., the non-zero columns for node a are from nodes belonging to the element *star* associated to a . The *star* of elements associated to a given node a is the set of all elements connected to a . An example of *global* sparsity pattern for the pipe benchmark unstructured mesh is reported in Figure 7. The node ordering starts from the outer surface and goes to the interior of the pipe, as seen from the reduced connectivity between the upper-left block and the remaining inner nodes. The density is less than 1 percent, and sparsity is more than 99.99 percent (Table 4).

A Reverse-Cuthill-McKee (RCM) bandwidth minimizing permutation of the same raw matrix shows a banded sparse structure illustrated in Figure 7. We report quantitative estimates for bandwidth and number of non-zeros, showing that bandwidth is approximately 1000 with a maximum of about 1500, and most nodes are connected to 15-16 other nodes.

Diagonal dominance.—A closer look at the magnitudes of entries in the *global* matrix \mathbf{A}_f reveals a clear block-structure (Figure 8). The 4-by-4 block structure corresponds to the matrix blocks in equation (13). This shows that diagonal entries are larger than the off-diagonals in each block of \mathbf{K} , suggesting that magnitudes of entries in the blocks of \mathbf{G} , \mathbf{D} , \mathbf{L} are small compared to diagonals in \mathbf{K} . This relates to the dominant contribution of the acceleration and advection terms, over the stabilization and viscous terms in (15) (see, e.g., [11]). The matrix \mathbf{A}_f is, however, not diagonally dominant (i.e. $|A_{ii}| < \sum_{j=1} |A_{ij}|$). To show this, we quantified the relative magnitudes of off-diagonal and diagonal entries (Figure 9). Specifically, we counted the number of elements that have a certain percentage of absolute magnitudes compared to the diagonal values, showing that most off-diagonal values are less

than 20% of the associated diagonal. To report quantitative estimates of diagonal dominance, we measure the mean of the relative magnitude of the sum of off-diagonal values to the diagonal,

$$D(\mathbf{K}) = \frac{1}{N_r} \sum_{i=1}^{N_r} \left[|K_{ii}| / \left(\sum_{\substack{j=1 \\ i \neq j}}^{N_c} |K_{ij}| \right) \right], \quad (29)$$

in which N_r and N_c is the number of rows and columns in \mathbf{K} respectively. D increases with the diagonal dominance. $D(\mathbf{A}_\rho)$ for the *global* matrix is 0.514 and $D(\mathbf{K})$ for the *global* matrix is 0.678.

Symmetry.—We use an index, S , to quantify how close a matrix is to symmetric. We first obtain off-diagonal elements of \mathbf{A} by subtracting its diagonal as $\tilde{\mathbf{A}} = (\mathbf{A} - \text{diag}(\mathbf{A}))$. We then decompose the $\tilde{\mathbf{A}}$ into the symmetric part, $\tilde{\mathbf{A}}_{sym} = (\tilde{\mathbf{A}} + \tilde{\mathbf{A}}^T) / 2$, and the skew-symmetric part, $\tilde{\mathbf{A}}_{skew} = (\tilde{\mathbf{A}} - \tilde{\mathbf{A}}^T) / 2$. The index S is defined as

$$S(\mathbf{A}) = \frac{|\tilde{\mathbf{A}}_{sym}| - |\tilde{\mathbf{A}}_{skew}|}{|\tilde{\mathbf{A}}_{sym}| + |\tilde{\mathbf{A}}_{skew}|}, \quad (30)$$

in which we use the 2-norm for $|\mathbf{A}|$. The index equals -1 for a perfectly skew-symmetric matrix and 1 for a perfectly symmetric matrix. As shown in Section 2.1, the matrix is nonsymmetric in the \mathbf{K} , \mathbf{G} , \mathbf{D} blocks due to stabilization and convective terms, with $S(\mathbf{A}_\rho)$ equal to 0.9859. That is, $S(\mathbf{A}_\rho)$ is a nearly symmetric matrix in the analyzed regime (ideal aortic flow). Finally, \mathbf{L} is a symmetric and semi-positive definite matrix with $S(\mathbf{L}) = 1$.

Eigenvalues.—Spectral properties are widely used to characterize the convergence and robustness of iterative solvers. It is well known, for example, that the rate of convergence of CG depends on the spectral radius of a left-hand-side SPD matrix. Despite eigenvalues clustered around 1 leading to rapid convergence of iterative solvers for well-conditioned SPD matrices, the eigenvalues may not be solely responsible for the convergence rate of these solvers and other matrix characteristics may play a role [28]. Calculation of all eigenvalues (λ_j) of the *global* matrix for a typical cardiovascular model with order 1 million mesh elements is prohibitively expensive. In this paper we therefore report the spectrum of *local* matrices instead of the *global* matrix. For a smaller size system, we also demonstrate that the distribution of eigenvalues from *local* matrices is a good approximation to the distribution of eigenvalue of the *global* matrix (See appendix D).

In Figure 10, we plotted the spectrum of *local* \mathbf{A}_f matrices from the pipe benchmark with rigid walls. The eigenvalues of \mathbf{A}_f are complex with small magnitudes of the imaginary part up to $\mathcal{O}(10^{-8})$, while the magnitudes of the real part ranges from $\mathcal{O}(10^{-9})$ to $\mathcal{O}(10^{-1})$. In Figure 10, there are three distinct groups of eigenvalues with different ranges of the real part. The first group contains those with the real part less than 10^{-5} . The second group contains those with real part ranging between 10^{-5} and 10^{-2} . The spectrums of \mathbf{K} and \mathbf{L} in Figure 10

show that the third group, with eigenvalues larger than 10^{-5} in \mathbf{A}_f is attributed to the \mathbf{K} block, while the eigenvalues of \mathbf{L} , \mathbf{G} and \mathbf{D} are responsible for the group of smallest real eigenvalues in \mathbf{A}_f . We list several minimum and maximum eigenvalues of a *local* matrix without resistance boundary condition (BC) in Table 3. The maximum eigenvalues of \mathbf{K} and \mathbf{A}_f appear to be the same, suggesting that block \mathbf{K} dominates the high portion of the spectrum, with the smallest eigenvalues provided by the blocks \mathbf{L} , \mathbf{G} and \mathbf{D} . This suggests that the large condition number of $\mathbf{A}_f \sim \mathcal{O}(10^6)$, obtained from MATLAB *condest*, relates to the inhomogeneous eigenvalue spectrum observed in the momentum, continuity and coupling blocks. Additionally, the small condition number of \mathbf{K} justifies the idea behind the BIPN approach, i.e., to solve the \mathbf{K} block separately, while expressing the other blocks in Schur complement form [11]. \mathbf{L} is singular, and thus has zero eigenvalue and an extremely large condition number. Lastly, the resistance boundary condition is responsible for the few largest eigenvalues order of $\mathcal{O}(10^{-1})$ in Figure 10, as we discuss in the next section.

Effects of the resistance boundary condition.—A resistance boundary condition perturbs the condition number of the coefficient matrix \mathbf{A}_f and may be responsible for a significant increase in the solution time for the tangent linear system [10]. The boundary traction \mathbf{h} , is given, in this case by

$$\mathbf{h}(\mathbf{u}, p, \mathbf{x}, t) = -P_i \mathbf{n}, \quad \mathbf{x} \in \Gamma_i, \quad (31)$$

in which P_i is the pressure at surface i , evaluated as $P_i = R_i Q_i$, i.e., proportional to the flow rate Q_i across the surface

$$Q_i(t) = \int_{\Gamma_i} \mathbf{u} \cdot \mathbf{n} d\Gamma, \quad (32)$$

through the prescribed resistance R_i . Thus, the contribution of the resistance boundary condition to the coefficient matrix is

$$\mathbf{K}^{bc} = \sum_{i=1}^{n^{bc}} \tilde{R}_i \mathbf{S}_i \otimes \mathbf{S}_i, \quad \mathbf{S}_i = \int_{\Gamma_i} N_a \mathbf{n} d\Gamma, \quad (33)$$

where n^{bc} is the number of resistance boundaries, and $\tilde{R}_i = \gamma \Delta t R_i$. \mathbf{K}^{bc} is finally added to the \mathbf{K} sub-matrix, resulting in the coefficient matrix $\tilde{\mathbf{K}} = \mathbf{K} + \mathbf{K}^{bc}$. Generalization from an outlet resistance to a coupled lumped parameter network model is accomplished using a slightly more general expression for \mathbf{K}^{bc} , i.e.

$$\mathbf{K}^{bc} = \sum_{k=1}^{n^{bc}} \sum_{l=1}^{n^{bc}} \gamma \Delta t M_{kl} \int_{\Gamma_k} N^a \mathbf{n}_i d\Gamma \int_{\Gamma_l} N^b \mathbf{n}_j d\Gamma, \quad (34)$$

$$M_{kl} = \frac{\partial P_k^{n+1}}{\partial Q_l^{n+1}},$$

where the resistance matrix M_{kl} is obtained by coupling pressures and flow rates at different outlets [25].

Addition of a resistance boundary condition alters the topology of the coefficient matrix due to the rank one contribution $\mathbf{S}_j \otimes \mathbf{S}_j$. This, in practice, couples all velocity degrees of freedom on a given outlet, significantly affecting the performance of matrix multiplication for the \mathbf{K} block and the fill-in generated by its LU decomposition. Thus, the vector \mathbf{S}_j is stored separately, to improve the efficiency of matrix multiplication and for RPC preconditioning. Figure 8 shows how the *global* matrix entries are affected by the presence of a resistance boundary condition, i.e., large magnitude components in the z-directional velocity block (lower-right block in \mathbf{K}) arise. To better highlight this effect, we show two *local* matrices with and without a resistance boundary condition in Figure 11. Addition of \mathbf{K}^{bc} increases the contribution of off-diagonal entries moving the matrix further away from diagonal dominance.

The resistance BC perturbs the spectrum and increases the condition number of \mathbf{A}_F . The largest few eigenvalues in the spectrum of \mathbf{A}_F and \mathbf{K} in Figure 10 are calculated from *local* matrices in partitioned domains interfacing the resistance boundary. The change of spectral properties of $\tilde{\mathbf{K}}$ due to rank one contribution \mathbf{K}^{bc} (see, e.g., [11]) is measured with maximum and minimum eigenvalues reported in Table 6. The maximum eigenvalue of $\tilde{\mathbf{K}}$ is significantly larger than \mathbf{K} , leading to a $\sim \mathcal{O}(10)$ -fold increase in the condition number of $\tilde{\mathbf{K}}$, thus increasing the spectral radius of the whole spectrum as shown in Figure 10. Additionally, our tests confirm that the first maximum eigenvalue \mathbf{A}_F increases linearly with the assigned resistance. Thus, in a more general case, we expect several large eigenvalues to be added to \mathbf{A}_F for models with multiple outlet resistances.

5.2 Matrix properties for fluid flow in deformable vessels

Sparsity pattern.—The *global* sparsity pattern for the FSI mesh is illustrated in Figure 12, where nodes in the solid-fluid interface are ordered first, followed by nodes in the fluid region next to the interface, and nodes in the solid domain. As shown in Figure 7 when the connectivity matrix is reordered by RCM, the global sparsity pattern has a banded sparse matrix structure with the larger maximum bandwidth of ≈ 1750 compared to the rigid case, while the number of non-zeros in a row is mostly clustered around 15 to 16, similar to the rigid case. Magnitudes of entries associated with solid nodes appear to be significantly larger than those in the fluid domain. For example, the magnitude of \mathbf{K}_s is order one (Figure 13), whereas the magnitude of \mathbf{K} is order 10^{-3} (Figure 11). In what follows, we focus on the *local* matrix \mathbf{K}_s from Eq. (27) since the characteristics of \mathbf{A}_F in FSI are similar to the rigid wall case.

Diagonal dominance.—The metric introduced above to quantify diagonal dominance drops to $D(\mathbf{A}_{\text{FSI}}) = 0.4775$ for the *global* matrix \mathbf{A}_{FSI} , i.e., the additional FSI terms reduce the diagonal dominance of the system. The magnitudes of a *local* \mathbf{K}_s matrix before and after diagonal scaling is compared, in Figure 13, to a *local* \mathbf{K} from the fluid domain. The diagonally scaled block \mathbf{K}_s qualitatively shows how the off-diagonals in \mathbf{K}_s are larger than

those in \mathbf{K} . The diagonal dominance metric for the *local* blocks \mathbf{K}_s and \mathbf{K} are found to be $D(\mathbf{K}_s) = 0.379$ and $D(\mathbf{K}) = 0.678$, respectively.

Symmetry and positive definiteness.—The symmetry metric for matrix \mathbf{K}_s is one, i.e., $S(\mathbf{K}_s) = 1$ and positive definite as expected from its properties and confirmed numerically.

Eigenvalues.—We calculated and plotted the eigenvalue spectrum of *local* matrices from the FSI benchmark in Figure 14. All eigenvalues of \mathbf{K}_s are real due to the symmetry of \mathbf{K}_s . As listed in Table 7, the magnitudes of eigenvalues in \mathbf{K}_s is significantly larger than the magnitudes of eigenvalues from \mathbf{A}_f .

5.3 Discussion

Results from the previous sections suggest the following conclusions. First, the condition number of both the fluid and solid tangent matrices \mathbf{A}_f and \mathbf{K}_s appears to be large, and therefore preconditioning is necessary. Second, the fluid matrix \mathbf{A}_f is more diagonally dominant than the solid matrix \mathbf{K}_s . This suggests that diagonal preconditioning is expected to be more effective for rigid wall simulations, but incomplete factorization preconditioners are expected to work better under fluid-structure interaction, consistent with the results obtained in the pipe benchmark. Third, resistance and coupled multidomain boundary conditions need a special treatment for preconditioning, due to their effect on the maximum eigenvalue and condition number.

6 Effect of preconditioning

In this section we investigate how application of various preconditioners affects the spectral properties of the coefficient matrix in both the rigid and deformable case by explicitly computing the preconditioned matrix $\mathbf{M}_l^{-1}\mathbf{A}\mathbf{M}_r^{-1}$, where \mathbf{M}_l^{-1} is a left preconditioner and \mathbf{M}_r^{-1} is a right preconditioner.

Consider a left and right Jacobi preconditioning for \mathbf{A}_f :

$$\begin{aligned} \mathbf{W}^m &= \text{diag}(\mathbf{K})^{-1/2}, \mathbf{W}^c = \text{diag}(\mathbf{L})^{-1/2}, \\ \mathbf{K}^* &\leftarrow \mathbf{W}^m \mathbf{K} \mathbf{W}^m, \mathbf{G}^* \leftarrow \mathbf{W}^m \mathbf{G} \mathbf{W}^c, \\ \mathbf{D}^* &\leftarrow \mathbf{W}^c \mathbf{D} \mathbf{W}^m, \mathbf{L}^* \leftarrow \mathbf{W}^c \mathbf{L} \mathbf{W}^c, \\ \Delta \mathbf{u}^* &\leftarrow \mathbf{W}^m \Delta \mathbf{u}, \Delta \mathbf{p}^* \leftarrow \mathbf{W}^c \Delta \mathbf{p}, \end{aligned} \quad (35)$$

resulting in the linear system $\mathbf{A}_f^* \mathbf{y}^* = -\mathbf{R}^*$. Spectral properties of the preconditioned matrix \mathbf{A}_f^* and \mathbf{K}_s^* are reported in Figure 15 and Table 8.

Table 8 shows how diagonal preconditioning is effective in improving the conditioning of \mathbf{A}_f , particularly for \mathbf{K} , without resistance BC. The condition number of \mathbf{K}^* is reduced to ~ 10 , and only a few linear solver iterations are expected to be sufficient to substantially reduce the approximated residual. This again justifies the approach followed by the BIPN solver, where the linear system involving the momentum block \mathbf{K} is solved separately, thus shifting the computational cost to the iterative solution of its Schur-complement. The

condition number of the approximated Schur complement block $\mathbf{L}^* + \mathbf{G}^{*T}\mathbf{G}^*$ is, in this example, equal to 120, consistent with previous findings in [11]. This also justifies the fact that ~80 percent of total compute time in BIPN is dedicated to the solution of the Schur complement linear system. Conversely, a symmetric Jacobi preconditioning does not seem to significantly reduce the condition number of the solid block \mathbf{K}_s . This is attributed to the presence of large off-diagonal values in \mathbf{K}_s that are only marginally affected by diagonal preconditioning. As a result, the eigenvalues of \mathbf{K}_s^* range from $\mathcal{O}(10^{-3})$ to $\mathcal{O}(1)$ while the eigenvalues of \mathbf{A}_f^* range $\mathcal{O}(10^{-2})$ to $\mathcal{O}(1)$, shown in Figure 15, with the exception of a few eigenvalues from the resistance BC. This, in turn, justifies the superiority of incomplete factorization preconditioners for FSI simulations.

Application of a symmetric Jacobi preconditioner to a *local* coefficient with resistance boundary condition leads to the eigenvalues reported in Table 9. Despite a reduction by three orders of magnitude, the condition number is still one order of magnitude larger than in the case without resistance BC. Note that the maximum eigenvalue of the preconditioned matrix is one order of magnitude larger than the second maximum eigenvalue, consistent with previous observations. Thus, the RPC preconditioning proposed in [10] seeks a preconditioning matrix \mathbf{H} such that $\mathbf{H} \approx (\mathbf{K}^{total})^{-1}$. The idea is to construct \mathbf{H} by combining the diagonal components of \mathbf{K} with the resistance contributions stored in \mathbf{S}_j as

$$\mathbf{H} = (\mathbf{K}^d)^{-1} - \sum_{j=1}^{n^{bc}} \left[\frac{R_j ((\mathbf{K}^d)^{-1} \mathbf{S}_j) \otimes ((\mathbf{K}^d)^{-1} \mathbf{S}_j)}{1 + R_j \left\| (\mathbf{K}^d)^{-\frac{1}{2}} \mathbf{S}_j \right\|^2} \right], \quad (36)$$

where \mathbf{K}^d is $\text{diag}(\mathbf{K})$. The preconditioned matrix $\mathbf{H}\tilde{\mathbf{K}}$ has a small condition number (Table 7) and smaller off-diagonal entries (Figure 16).

7 Performance of linear solvers in patient-specific models

In this section we tested the performance of linear solvers on patient-specific cardiovascular models, in an effort to extrapolate the results obtained for the pipe benchmark to more realistic problems. We use three different models with a wide range of boundary conditions (i.e., resistance, RCR, coronary BC, closed-loop multidomain), with and without wall deformability and covering various patient-specific anatomies. All anatomic models were constructed from medical image data using Sim-Vascular.

7.1 Pulmonary hypertension

The first model represents the left and right pulmonary arteries with associated branches and is used to investigate the effects of pulmonary hypertension (PH). The finite element mesh contains 3,223,072 tetrahedral elements to represent the pulmonary lumen, has rigid walls and 88 outlets with Windkessel (RCR) boundary conditions, prescribed through a coupled 0-D multi-domain approach [25, 40] (Figure 17). A pulsatile inflow waveform extracted from PC-MRI was imposed at the pulmonary artery inlet. This model is solved using a time step

of 0.46 milliseconds and 120 cores ($\approx 25,000$ elements per core). The tolerance on the Newton-Raphson residual is set to $\epsilon = 10^{-4}$.

We briefly report *global* matrix characteristics in the PH model. The diagonal dominance metric is $D(\mathbf{A}_f) = 0.5598$, and $D(\mathbf{K}) = 0.7397$. The metric value for $D(\mathbf{A}_f)$ is similar to values from the pipe model with rigid walls. The matrix is nearly-symmetric, $S(\mathbf{A}_f) = 0.9903$.

Results in Figure 17 compare the performance of diagonal, block-diagonal and ILU preconditioning. BIPN with RPC shows the best performance, followed by ILU-BICG. This is expected due to the large number of resistance boundary conditions (i.e., 88) at the model outlets. Diagonal preconditioners with GMRES instead perform poorly, consistent with our observations in the pipe benchmark.

7.2 Coronary artery bypass graft

Second, we consider a model of coronary artery bypass graft surgery (see, e.g., [27, 36]) with 4,199,945 tetrahedral elements and rigid walls, coupled with a closed-loop 0D lumped parameter network (LPN), including coupled heart, coronary and systemic circulation models (Figure 18). Simulations were performed using 168 cores ($\sim 24,000$ elements per core), with a time step of 0.87 millisecond and a non linear iteration tolerance of $\epsilon = 10^{-3}$.

The diagonal dominance metric for the *global* matrix in the CABG model is $D(\mathbf{A}_f) = 0.5200$, and $D(\mathbf{K}) = 0.6914$. The matrix in the CABG model is less diagonally dominant than the previous cylinder or pulmonary hypertension model. The matrix is also near symmetric, $S(\mathbf{A}_f) = 0.9938$.

As expected, due to the presence of coupled multidomain boundary conditions [10], BIPN results in the best performance, followed closely by BICG with ILU, while GMRES with a diagonal preconditioner performs poorly. The relative performance of ILU against the diagonal preconditioner is better than seen in previous models. The smaller diagonal dominance metric in CABG model confirms superiority of ILU over the diagonal preconditioner.

7.3 Left coronary

Next, we tested performance of linear solvers in a left coronary model. The left coronary model was extracted from a full coronary artery model used in Tran *et al.*[36]. The pulsatile flow waveform at the inlet of the left coronary branch was extracted from the full model simulation and imposed at the inlet of the model. The model has six outlets, each with applied open-loop coronary outlet boundary conditions [17]. All resistance and compliance values as well as inflow pulsatile waveforms are determined to produce a normal physiologic response of the left coronary artery following our prior work. We ran simulations with rigid and deformable walls with a lumen mesh containing 486,066 tetrahedral elements and a vessel wall mesh with 206,369 tetrahedral elements, time step of 1 millisecond and tolerance of $\epsilon = 10^{-4}$. We used 20 cores for the rigid wall simulation and 24 cores for the deformable wall simulation.

The diagonal dominance metric for the left coronary model with rigid wall is $D(\mathbf{A}_f) = 0.5238$, and $D(\mathbf{K}) = 0.6972$, which are similar to the CABG model. The matrix is near symmetric, $S(\mathbf{A}_f) = 0.9855$, however, this model is furthest from symmetric among all models considered.

For the ALE FSI model, the diagonal dominance is reduced as $D(\mathbf{A}_{FSI}) = 0.4323$, $D(\mathbf{K}) = 0.5320$. Note that the number of elements in a wall mesh is 40 percent of the fluid mesh, so the effect of adding solid mechanics in the linear system is more significant than the pipe case where only 20 percent of elements were in the in wall mesh. As a result, the symmetry metric is very close to 1, $S(\mathbf{A}_{FSI}) = 0.99998$, since \mathbf{K}_s is symmetric.

As shown in Figure 20, performance test results are consistent with previous findings. RPC-BIPN is the fastest method for the rigid wall simulation. In FSI, the performance of BIPN is poor, while diagonal and ILU preconditioners with GMRES perform better.

7.4 Discussion

From the performance results in the patient-specific models, we find that RPC-BIPN is the fastest method for rigid wall simulations with many resistance BCs in agreement with the pipe model. ILU-BICG is only slightly slower than RPC-BIPN while the standard diagonal scaled GMRES fails. The performance degradation of RPC-BIPN in FSI models is consistent with the pipe model and suggests the need for future improvements of BIPN for ALE FSI.

8 Summary and conclusions

In this paper we study the performance of preconditioned iterative linear solvers for cardiovascular simulation in rigid and deformable walls. To this end, we implement various iterative linear solvers - GMRES, BICGS, and BIPN - and preconditioners - diagonal, block-diagonal, ILU, ILUT, ML, and RPC in a single flow solver. Standard iterative solvers and preconditioners are employed from the Trilinos library and compared against RPC-BIPN, implemented in our in house solver. Simulation wall clock time is measured and compared in a benchmark pipe flow with a resistance BC.

ILU preconditioned BICG provides the best overall performance in both rigid and deformable wall simulations. RPC-BIPN in the FSI simulation shows ~ 8 fold increase in compute time compared to the rigid wall case. Strong and weak scalings of ILU-BICG and RPC-BIPN are reported.

To better understand the observed performance, characteristics of the left-hand-side matrix in the linear system are examined. We report sparsity patterns, diagonal dominance, symmetry, eigenvalues and condition numbers in *global* and *local* matrices. Results show that the sparse matrix structure has a narrow banded structure after Reverse-Cuthill-McKee reordering. The matrix from the fluid domain has larger diagonal values than off-diagonals and is nearly symmetric. Eigenvalues and the condition number of the matrix from the fluid domain show that the \mathbf{K} block has a significantly smaller condition number compared to the \mathbf{A}_f matrix, supporting the main premise of BIPN. Effects of preconditioning on matrix

characteristics are investigated by explicitly forming the preconditioned matrix. A diagonal preconditioner is shown to be effective to reduce the range of eigenvalues in the fluid domain, especially for the \mathbf{K} matrix.

Adding wall deformability to the fluid simulation increases the bandwidth of the matrix and decreases the relative magnitudes of the diagonal values compared to the off-diagonal values. Due to the reduction of diagonal dominance, a diagonal preconditioner does not significantly reduce the condition number of the original matrix.

The resistance boundary condition disturbs the sparsity and diagonal dominance of the original fluid matrix, and causes an ill-conditioned system by adding an eigenvalue which is larger than the maximum eigenvalue of matrix without resistance BC. The resistance based preconditioner successfully reduces the condition number of the system with a resistance boundary condition by four orders of magnitude, while a diagonal preconditioner only reduces the condition number by two orders of magnitudes.

The performance of various preconditioned linear solvers is evaluated in four patient-specific models. In these models, RPC-BIPN is best for rigid wall models with multiple resistance or coupled LPN outlet boundary conditions. In the deformable wall simulation, RPC-BIPN shows significant performance degradation and diagonal preconditioners or ILU with BICG achieve the best performance.

This study motivates several new research directions to develop new preconditioned linear solver strategies. The effectiveness of BIPN for the solution of fluids problems with rigid walls has been proven in the current study. Currently our in-house code (RPC-BIPN) uses the bi-partitioned approach for ALE FSI, forming a linear system from the momentum equations for the fluid and the solid domains together, and another system for the continuity equation for the fluid domain. However, since the characteristics of the matrix in the solid domain is different from the fluid domain, most notably diagonal dominance, linear systems from these two domains should be separately solved (i.e. Tripartitioning). The inefficiency of solving FSI in BIPN stems from adding off-diagonal dominance to the left-hand-side matrix block \mathbf{K} . Since RPC is based on a simple diagonal preconditioner, solving the system \mathbf{K} becomes less efficient. We suggest solving \mathbf{K}_s separately with an Incomplete Cholesky preconditioner, exploiting its symmetric property, rather than a simple diagonal preconditioning. Exploration of this idea is the subject of future work.

Additionally, we point out that the Schur complement block in BIPN is not preconditioned. Since a major portion of the computational cost in BIPN is consumed when solving the Schur complement block [11], acceleration of the linear solver performance by a proper preconditioning technique could significantly reduce the compute time. To form a preconditioner for the Schur complement block, one would need an efficient sparse matrix-matrix multiplication scheme as well as explicit formation of the Schur complement block. The open-source Trilinos library provides this option as well as various preconditioners so combining a partitioning approach with Trilinos is expected to provide consistent performance in both rigid and deformable wall simulations of cardiovascular hemodynamics. Implementation and testing of this approach is left for future investigation.

Testing of linear solver performance on more complex patient-specific disease with large wall deformations (e.g. aortic dissection) are warranted, and would likely lead to further insights. Future studies are also warranted to further assess solver performance and matrix characteristics, towards development of new solver and preconditioner strategies.

Acknowledgements

This work was supported by NIH grant (NIH R01-EB018302), NSF SSI grants 1663671 and 1339824, and NSF CDSE CBET 1508794. This work used the Extreme Science and Engineering Discovery Environment (XSEDE) [35], which is supported by National Science Foundation grant number ACI-1548562. We thank Mahidhar Tatineni for assisting on building Trilinos on Comet cluster, which was made possible through the XSEDE Extended Collaborative Support Service (ECSS) program [1]. The authors also thank Michael Saunders, Michael Heroux, Mahdi Esmaily, Ju Liu, and Vijay Vedula, for fruitful discussions that helped in the preparation of this paper. The authors would like to thank the two anonymous reviewers whose comments greatly contributed to improve the completeness of the present study. We also acknowledge support from the open source SimVascular project at www.simvascular.org.

9: Appendix

A GMRES restart

We tested different GMRES restart numbers in the pipe benchmark problems. In Figure 20, we plot compute times for preconditioned GMRES using a pipe model with rigid wall and a pipe with deformable wall. Our test shows that decreasing restart number increases the compute time of linear solver in the rigid pipe model. The FSI model does not show a notable difference.

B Choice of smoother and subsmoother for ML

In the ML package, multiple options are available for the smoother and the subsmoother. As shown in Figure 21, the Gauss-Seidel smoother works best among Chebyshev, symmetric Gauss-Seidel, and ILUT. For the subsmoother, the symmetric Gauss-Seidel is the best among Chebyshev and MLS.

C Effect of reordering in ILU

We evaluated and compared compute times of linear solvers with different reordering methods. RCM and METIS reordering for ILUT via the Trilinos IFPACK are implemented. We use 2 level fill-in and 10^{-2} dropping tolerance for this test. Figure 23 shows performance differences between ILUT with different reordering schemes. From the testing, we confirm that the RCM is the fastest method against METIS and no reordering. The superior performance of RCM is notable when GMRES is used with ILUT.

D Eigenvalue spectrums of the *local* and *global* matrix.

In this section we compare the spectrum of eigenvalues in the *local* and *global* matrices and investigate how our analysis on *local* eigenvalues can be generalized to the *global* matrix. We use a pipe model in the same dimension shown in Figure 1 meshed with 24,450 elements with $N_{nd} = 5462$. We use one core to extract the *global* matrix, and four cores to examine *local* matrices. As shown in Figure 24, the eigenvalue distributions of the *global* and *local*

matrices are similar. Although the eigenvalues from the *global* and *local* matrices are not exactly the same, the distribution of eigenvalues of *local* matrices is a good approximation to the distribution of eigenvalues in the *global* matrix.

References

1. High Performance Computer Applications 6th International Conference, volume 595, Germany, 1 2016.
2. Bazilevs Y, Calo VM, Hughes TJR, and Zhang Y. Isogeometric fluid–structure interaction: theory, algorithms and computations. *Comput. Mech*, 43:3–37, 2008.
3. Bazilevs Y, Hsu MC, Benson DJ, Sankaran S, and Marsden AL. Computational fluid-structure interaction: methods and application to a total cavopulmonary connection. *Comput. Mech*, 45(1): 77–89, 12 2009.
4. Benzi M. Preconditioning techniques for large linear systems: A survey. *Journal of Computational Physics*, 182:418–477, 2002.
5. Benzi M, Szyld DB, and van Duin A. Orderings for incomplete factorization preconditioning of nonsymmetric problems. *Siam J. Sci. Comput*, 20(5):1652–1670, 1999.
6. Chen W and Poirier B. Parallel implementation of efficient preconditioned linear solver for grid-based applications in chemical physics. i: Block jacobi diagonalization. *Journal of Computational Physics*, pages 185–197, 2005.
7. Corsini C, Cosentino D, Pennati G, Dubini G, Hsia T-Y, and Migliavacca F. Multiscale models of the hybrid palliation for hypoplastic left heart syndrome. *Journal of Biomechanics*, 44:767–770, 2011. [PubMed: 21092964]
8. Deparis S, Forti D, Grandperrin G, and Quarteroni A. Facsi: A block parallel preconditioner for fluid-structure interaction in hemodynamics. *Journal of Computational Physics*, 327:700–718, 2016.
9. dos Santos RW, Plank G, Bauer S, and Vigmond EJ. Parallel multigrid preconditioner for the cardiac bidomain model. *IEEE Transactions on Biomedical Engineering*, 51(11):1960–1967, 2004. [PubMed: 15536898]
10. Esmaily-Moghadam M, Bazilevs Y, and Marsden AL. A new preconditioning technique for implicitly coupled multidomain simulations with applications to hemodynamics. *Comput. Mech*, pages DOI 10.1007/s00466-013-0868-1, 2013.
11. Esmaily-Moghadam M, Bazilevs Y, and Marsden AL. A bi-partitioned iterative algorithm for solving linear systems arising from incompressible flow problems. *Computer Methods in Applied Mechanics and Engineering*, 286:40–62, 2015.
12. Esmaily-Moghadam M, Bazilevs Y, and Marsden AL. Impact of data distribution on the parallel performance of iterative linear solvers with emphasis on cfd of incompressible flows. *Comput. Mech*, 55:93–103, 2015.
13. Figueroa CA, Vignon-Clementel IE, Jansen KE, Hughes TJ, and Taylor CA. A coupled momentum method for modeling blood flow in three-dimensional deformable arteries. *Comput. Meth. Appl. Mech. Engrg*, 195(41-43):5685–5706, 2006.
14. Heroux MA, Bartlett RA, Howle VE, Hoekstra RJ, Hu JJ, Kolda TG, Lehoucq RB, Long KR, Pawlowski RP, Phipps ET, Salinger AG, Thornquist HK, Tuminaro RS, Willenbring JM, Williams A, and Stanley KS. An overview of the trilinos project. *ACM Trans. Math. Softw*, 31(3):397–423, 2005.
15. Hughes TJ, Liu WK, and Zimmermann TK. Lagrangian-eulerian finite element formulation for incompressible viscous flows. *Comput. Meth. Appl. Mech. Engrg*, 29(3):329–349, 1981.
16. Jansen KE, Whiting CH, and Hulbert GM. A generalized- α method for integrating the filtered Navier–Stokes equations with a stabilized finite element method. *Comput. Meth. Appl. Mech. Engrg*, 190(3-4):305–319, 2000.
17. Kim HJ, Vignon-Clementel IE, Figueroa CA, LaDisa JF, Jansen KE, Feinstein JA, and Taylor CA. On Coupling a Lumped Parameter Heart Model and a Three-Dimensional Finite Element Aorta Model. *Ann. Biomed. Eng*, 37(11):2153–2169, 2009. [PubMed: 19609676]

18. Lagana K, Dubini G, Migliavacca F, Pietrabissa R, Pennati G, Veneziani A, and Quarteroni A. Multiscale modelling as a tool to prescribe realistic boundary conditions for the study of surgical procedures. *Biorheology*, 39:359–364, 2002. [PubMed: 12122253]
19. Lan H, Updegrove A, Wilson NM, Maher GD, Shadden SC, and Marsden AL. A re-engineered software interface and workflow for the open source simvascular cardiovascular modeling. *Journal of Biomechanical Engineering*, 140(2):024501:1–11, 2017.
20. Long CC. Fluid–structure interaction: Physiologic simulation of pulsatile ventricular assist devices using isogeometric analysis. Dissertation in University of California, San Diego, 2013.
21. Manguoglu M, Takizawa K, Sameh AH, and Tezduyar TE. Solution of linear systems in arterial fluid mechanics computations with boundary layer mesh refinement. *Comput. Mech*, 46:83–89, 2010.
22. Marsden AL. Optimization in cardiovascular modeling. *Annual Review of Fluid Mechanics*, 46:519–546, 2014.
23. Marsden AL and Esmaily-Moghadam M. Multiscale modeling of cardiovascular flows for clinical decision support. *Applied Mechanics Reviews*, 67:030804, 2015.
24. Marsden AL, Feinstein JA, and Taylor CA. A computational framework for derivative-free optimization of cardiovascular geometries. *Comput. Meth. Appl. Mech. Engrg*, 197:1890–905, 2008.
25. Moghadam ME, Vignon-Clementel I, Figliola R, and Marsden A. A modular numerical method for implicit 0d/3d coupling in cardiovascular finite element simulations. *Journal of Computational Physics*, 244:63–79, 2013.
26. Nesbitt WS, Westein E, Tovar-Lopez FJ, Tolouei E, Mitchell A, Fu J, Carberry J, Fouras A, and Jackson SP. A shear gradient–dependent platelet aggregation mechanism drives thrombus formation. *Nature Medicine*, 15(6):665–673, 2009.
27. Ramachandra AB, Kahn AM, and Marsden AL. Patient-specific simulations reveal significant differences in mechanical stimuli in venous and arterial coronary grafts. *J. of Cardiovasc. Trans. Res*, 2016.
28. Saad Y. *Iterative Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics, ISBN: 978-0-89871-534-7, 2003.
29. Sankaran S, Kim HJ, Choi G, and Taylor CA. Uncertainty quantification in coronary blood flow simulations: Impact of geometry, boundary conditions and blood viscosity. *Journal of Biomechanics*, 49:2540–2547, 2016. [PubMed: 26803339]
30. Schiavazzi DE, Arbia G, Baker C, Hlavacek AM, Hsia TY, Marsden AL, Vignon-Clementel IE, and T. M. of Congenital Hearts Alliance (MOCHA) Investigators. Uncertainty quantification in virtual surgery hemodynamics predictions for single ventricle palliation. *Int. J. Numer. Meth. Biomed. Engrg*, e02737:1–25, 2016.
31. Schiavazzi DE, Baretta A, Pennati G, Hsia T-Y, and Marsden AL. Patient-specific parameter estimation in single-ventricle lumped circulation models under uncertainty. *Int. J. Numer. Meth. Biomed. Engrg*, e02799, 2017.
32. Schiavazzi DE, Doostan A, Iaccarino G, and Marsden A. A generalized multi-resolution expansion for uncertainty propagation with application to cardiovascular modeling. *Comput. Methods Appl. Mech. Engrg*, 314:196–221, 2017.
33. Sengupta D, Kahn AM, Burns JC, Sankaran S, Shadden SC, and Marsden AL. Image-based modeling of hemodynamics in coronary artery aneurysms caused by kawasaki disease. *Biomech Model Mechanobiol*, 11:915–932, 2012. [PubMed: 22120599]
34. Taylor CA, Fonte TA, and Min JK. Computational fluid dynamics applied to cardiac computed tomography for noninvasive quantification of fractional flow reserve. *Journal of the American College of Cardiology*, 61(22):2233–41, 2013. [PubMed: 23562923]
35. Towns J, Cockerill T, Dahan M, Foster I, Gathier K, Grimshaw A, Hazlewood V, Lathrop S, Lifka D, Peterson GD, Roskies R, Scott JR, and Wilkins-Diehr N. Xsede: Accelerating scientific discovery. *Computing in Science and Engineering*, 16(5):62–74, Sep-Oct 2014.
36. Tran JS, Schiavazzi DE, Ramachandra AB, Kahn AM, and Marsden AL. Automated tuning for parameter identification and uncertainty quantification in multiscale coronary simulations. *Comput. Fluids*, 142:128–138, 2017. [PubMed: 28163340]

37. Trefethen LN and Bau D. Numerical Linear Algebra. SIAM, 1997.
38. Updegrove A, Wilson NM, Merkow J, Lan H, Marsden AL, and Shadden SC. Simvascular: An open source pipeline for cardiovascular simulation. *Annals of Biomedical Engineering*, 45(3):525–541, 2016. [PubMed: 27933407]
39. Vedula V, Lee J, Xu H, Kuo C-CJ, Hsiai TK, and Marsden AL. A method to quantify mechanobiologic forces during zebrafish cardiac development using 4-d light sheet imaging and computational modeling. *PLoS Comput Biol*, 13(10):e1005828, 2017. [PubMed: 29084212]
40. Yang W, Marsden AL, Ogawa MT, Sakarovitch C, Hall KK, Rabinovitch M, and Feinstein JA. Right ventricular stroke work correlates with outcomes in pediatric pulmonary arterial hypertension. *Pulmonary Circulation*, 2018.




Fig. 1: Schematic representations and mesh for cylindrical pipe benchmark, (top) a rigid model with parabolic inflow and outlet resistance boundary condition, (bottom) an FSI model with same boundary conditions. For each model we show a magnified view of the tetrahedral finite element mesh. The fluid and solid domains are colored in red and gray, respectively.

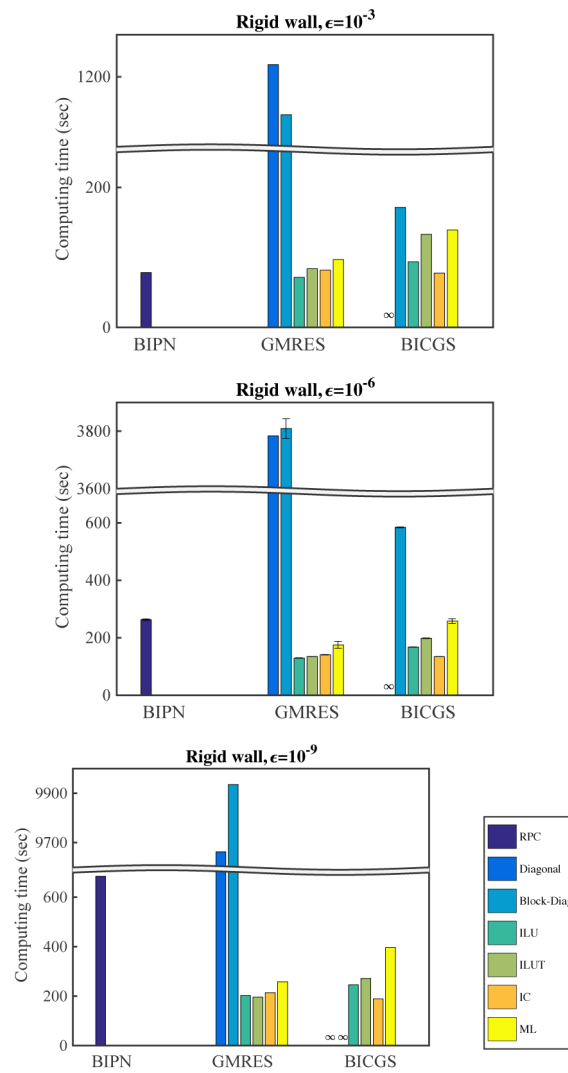


Fig. 2: Compute times for linear solvers and preconditioners using a rigid pipe model with tolerances (top) $\epsilon = 10^{-3}$, (center) $\epsilon = 10^{-6}$, (bottom) $\epsilon = 10^{-9}$. For $\epsilon = 10^{-6}$, error bars are plotted by taking standard deviations from two repeated simulations. Differences between repeated simulations are caused by different computing nodes assigned by the scheduler on the Comet cluster.

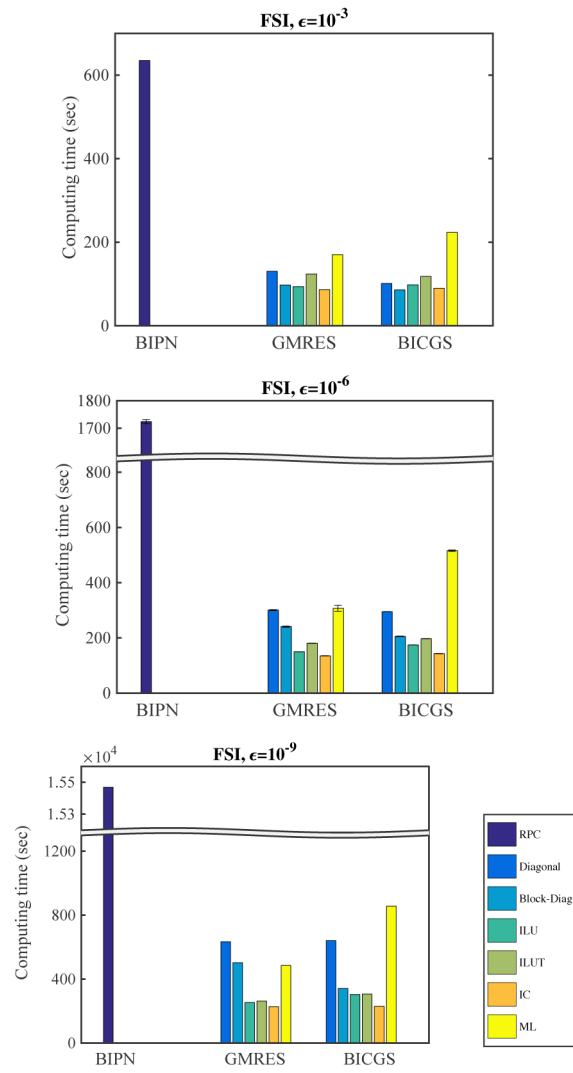


Fig. 3: Compute times for linear solvers and preconditioners using a deformable wall pipe model (FSI) with tolerances (top) $\epsilon = 10^{-3}$, (middle) $\epsilon = 10^{-6}$, (bottom) $\epsilon = 10^{-9}$. For $\epsilon = 10^{-6}$, error bars are plotted by taking standard deviations from two repeated simulations.

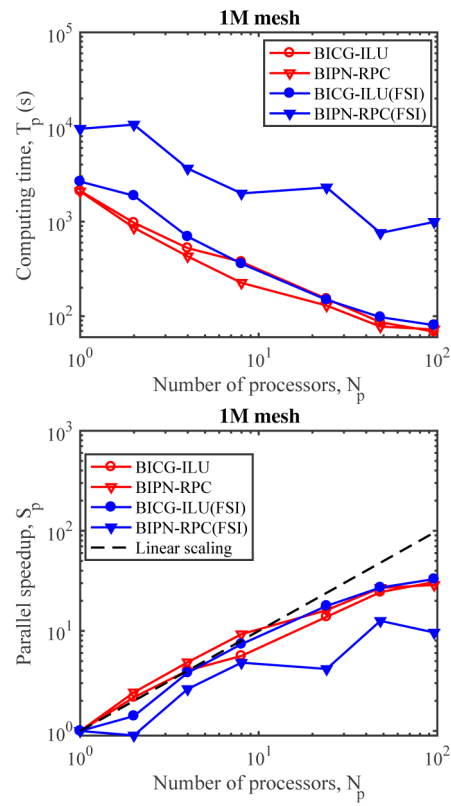


Fig. 4: Strong scaling of BIPN-RPC and BICG-ILU for pipe benchmark with rigid and deformable walls on the 1M lumen mesh model (Mesh2). (top) compute time, T_p versus number of cores, N_p , (bottom) speedup, $S_p = T_1/T_p$, versus N_p .

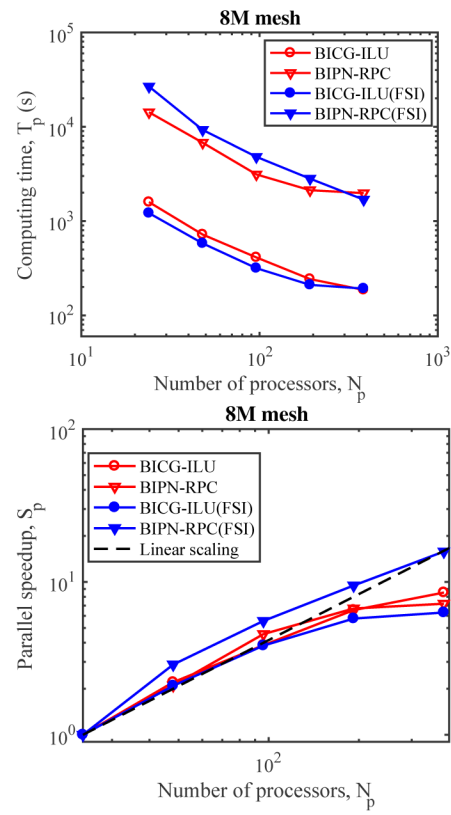


Fig. 5: Strong scaling of BIPN-RPC and BICG-ILU for pipe benchmark on the 8M lumen mesh model (Mesh5). (top) T_p versus N_p , (bottom) $S_p = T_{24}/T_p$ versus N_p .

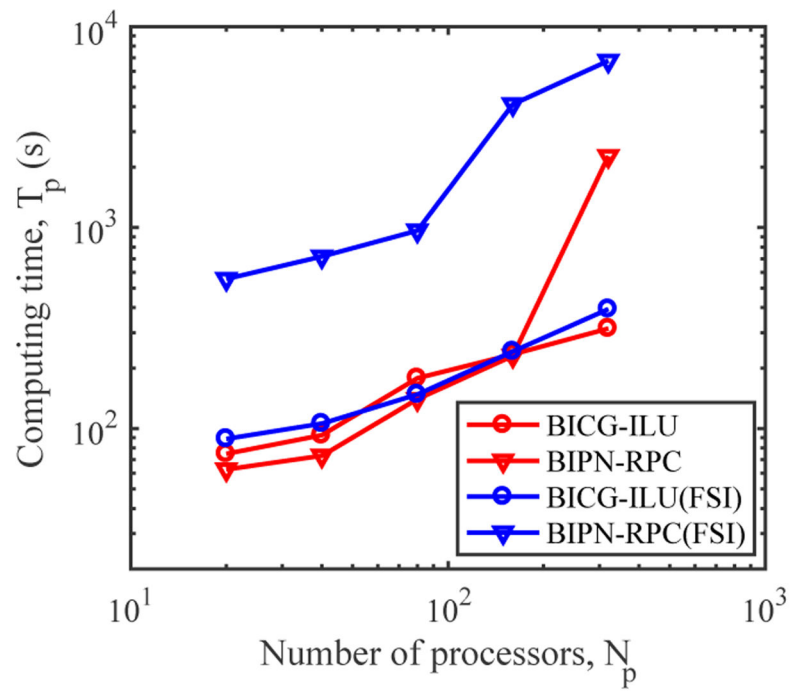


Fig. 6: Weak scaling of BICG-ILU and BIPN-RPC for pipe benchmark with rigid and deformable walls.

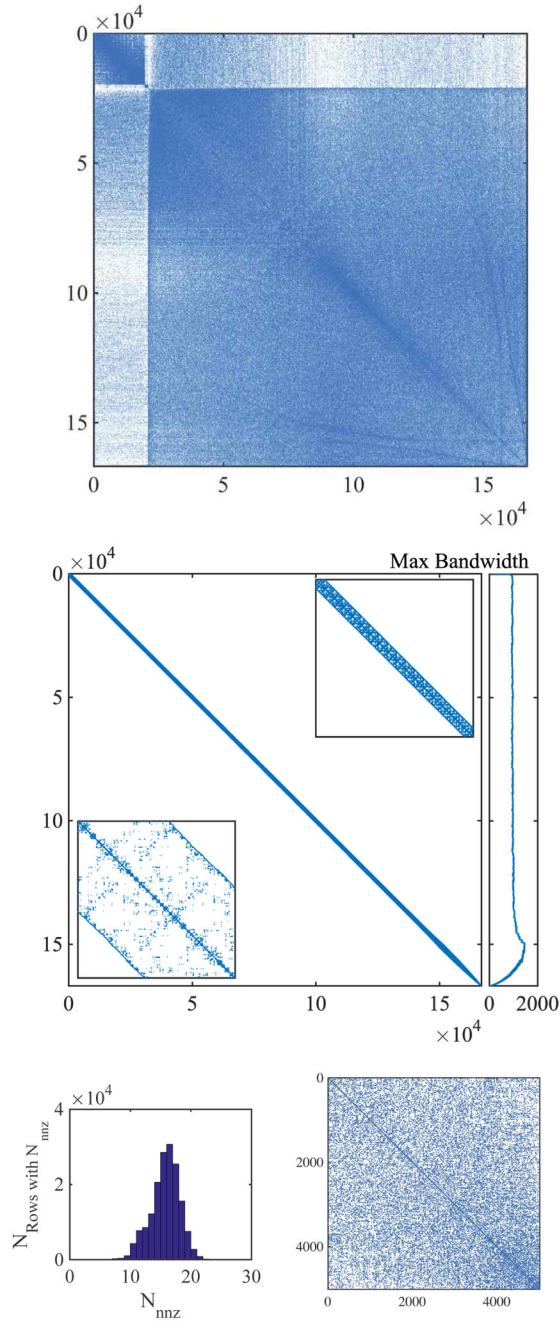


Fig. 7: Sparsity pattern of a connectivity matrix for the 1M pipe mesh. Among all matrix elements, only non-zero values are colored in the plot. (top) The *global* sparsity pattern for the whole pipe model with $N_{nd} = 166,899$. (center) Reverse Cuthill-McKee reordering of the *global* connectivity matrix. The upper-right inset is 100 magnification of the diagonal of the reordered matrix. The lower-left inset is 10,000 magnification. For each row, the left exterior plot shows the maximum bandwidth in that row. (bottom left) Frequencies for the number of non-zero elements in a row (N_{nnz}). (bottom right) A *local* sparsity pattern from a core with $N_{Ind} = 5014$.

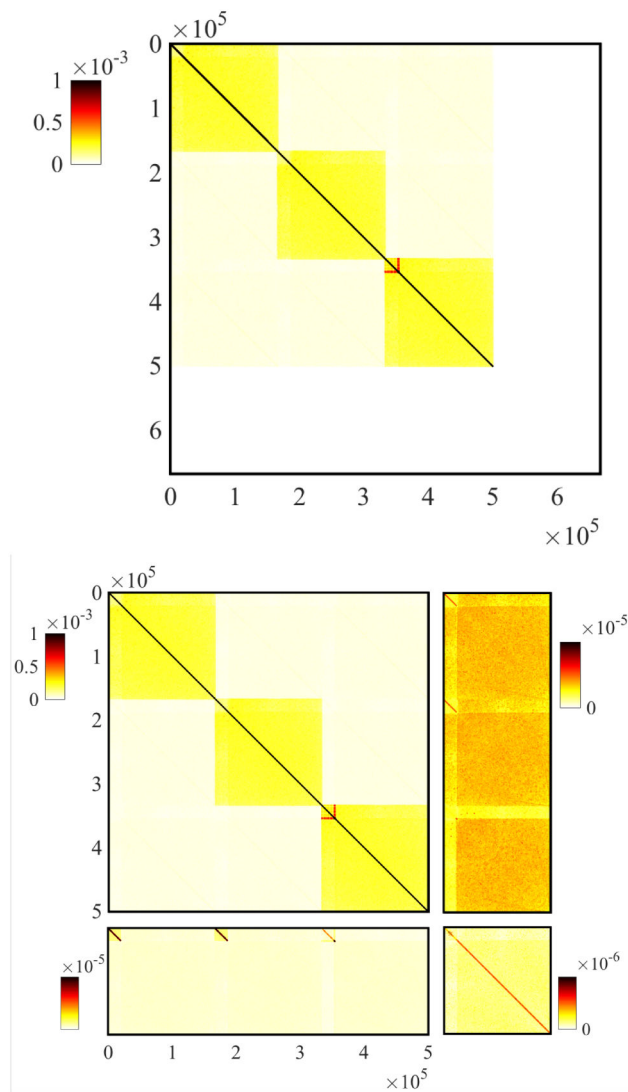


Fig. 8: A visual representation of the *global* sparse matrix \mathbf{A}_f with entries colored by absolute magnitude. (top) The full matrix \mathbf{A}_f colored by magnitudes ranging from 0 to 10^{-3} . (bottom) A decomposed matrix colored by magnitudes of each sub-block. In \mathbf{K} in the upper-left block, the colorbar ranges from 0 to 10^{-3} . In \mathbf{G} in the upper-right block and \mathbf{D} , in the lower-left block, the colorbar ranges from 0 to 10^{-5} . In \mathbf{L} in the lower-right block the colorbar ranges from 0 to 10^{-6} .




Fig. 9: Measures of diagonal dominance. (top left) The magnitude of diagonal values in each row of \mathbf{A}_f . (top right) The sum of absolute magnitude of off-diagonal values in each row of \mathbf{A}_f . (bottom) Histogram of the number of matrix elements $N(A_{ij})$ with magnitudes that correspond to the certain percentage of the associated diagonal entry. Only elements that are more than 1% of diagonal entries are counted.




Fig. 10: Spectrum of (top) *local* fluid matrices \mathbf{A}_f (bottom left) \mathbf{L} blocks, (bottom right) \mathbf{K} blocks. All *local* eigenvalue spectrums from 38 cores are plotted together with different colors.




Fig. 11: A visual representation of (top) a *local* matrix \mathbf{A}_f without resistance BC (bottom) a *local* matrix $\mathbf{A}_{f,res}$ with a resistance BC, $R=1600\text{g}/\text{cm}^4/\text{s}$. Matrix elements are colored by their absolute magnitude. For both figures color ranges from 0 to 10^{-3} .




Fig. 12: (Top) *Global* sparsity pattern for the FSI pipe benchmark with $N_{nd} = 198,128$. (bottom) Visual representation of the entry magnitudes for the *global* matrix \mathbf{A}_{FSI} .




Fig. 13: Sparsity pattern of a $local \mathbf{K}_s$ colored by absolute magnitude of each entry, for $N_{Ind} = 4892$. (top) A raw matrix, K_s , in the solid domain (center) A scaled matrix with its diagonal \mathbf{K}_s^* in the solid domain (bottom) A scaled $local \mathbf{K}^*$ with its diagonal in the fluid domain.




Fig. 14: Spectrum of (red) *local* A_f in the fluid domain and (blue) K_s in the solid domain. All *local* eigenvalue spectrums from 48 cores are plotted together.




Fig. 15: Spectrum of (red) *local* \mathbf{A}_f^* in the fluid domain and (blue) \mathbf{K}_s^* in the solid domain. All *local* eigenvalue spectrums from 48 cores are plotted together.




Fig. 16: *local* sparse matrix structure of preconditioned \mathbf{K} block. (top) $\tilde{\mathbf{K}}^*$ without resistance BC, (center) $\tilde{\mathbf{K}}^*$ with a resistance BC, (bottom) $\mathbf{H}\tilde{\mathbf{K}}$. The colorbar is same for all figures.

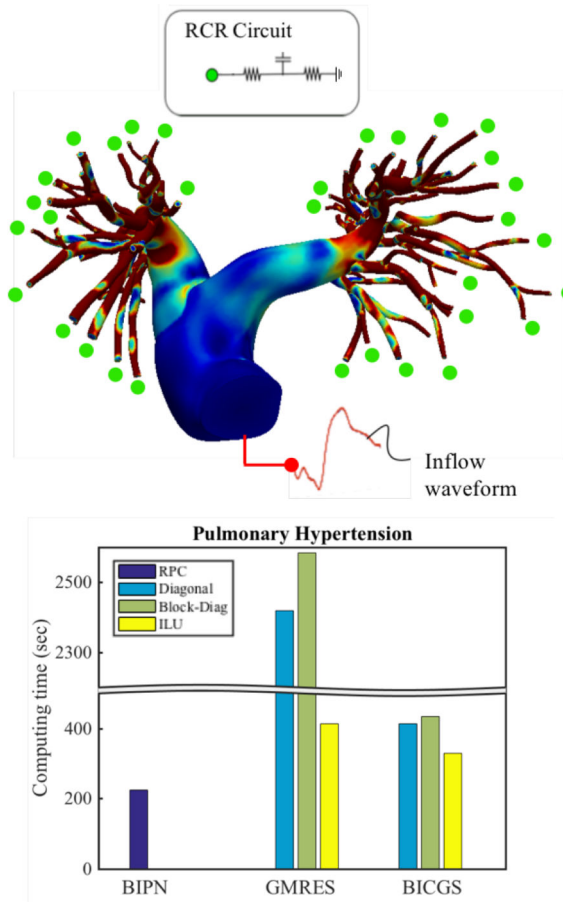


Fig. 17: (Top) Patient-specific model for pulmonary hypertension with schematics of boundary conditions. The model is colored by instantaneous wall shear stress. (bottom) compute times for preconditioned iterative linear solvers.

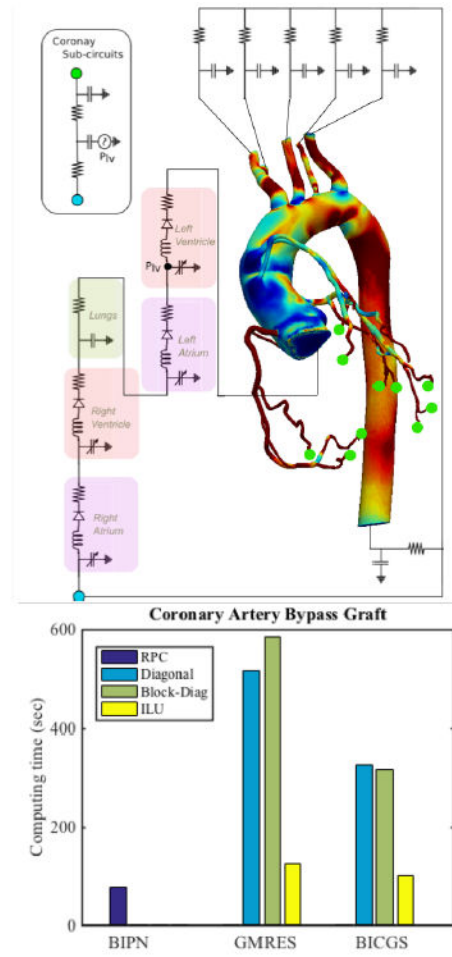


Fig. 18: (Top) Patient-specific model for coronary bypass graft model with schematics of boundary conditions. The model is colored by instantaneous wall shear stress. (bottom) compute times for preconditioned iterative linear solvers.

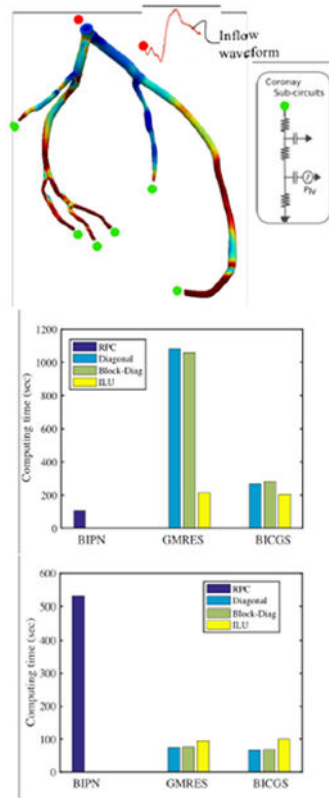


Fig. 19: (Top) Patient-specific model for left coronary artery model with schematics of boundary conditions. The model is colored by instantaneous wall shear stress. (center) compute times for preconditioned iterative linear solvers in the rigid wall simulation. (bottom) compute times for preconditioned iterative linear solvers in the deformable wall simulation.




Fig. 20: Compute times for GMRES and preconditioners using (top) a rigid pipe model, and (bottom) a pipe model with deformable wall using different the GMRES restart numbers.




Fig. 21: Compute times for linear solvers preconditioned with ML using a (a) rigid and (b) an FSI pipe model with tolerance $\epsilon = 10^{-3}$, with different smoothers. The symmetric Gauss-Seidel subsolver is used. For the rigid wall model, 38 cores are used. For the FSI model, 48 cores are used.

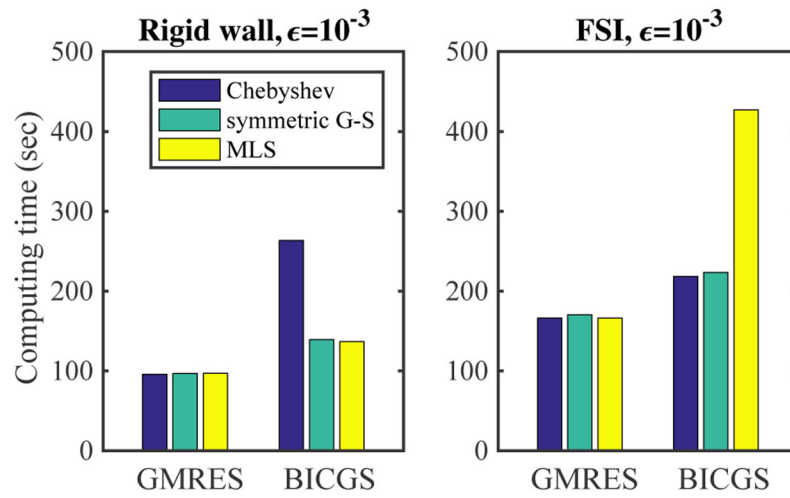


Fig. 22: Compute times for linear solvers preconditioned with ML using a (a) rigid and (b) FSI pipe model with tolerance $\epsilon = 10^{-3}$, with different subsolvers. The Gauss-Seidel smoother is used. For the rigid wall model, 38 cores are used. For the FSI model, 48 cores are used.




Fig. 23: Compute times for linear solvers preconditioned with ILUT using a (a) rigid and (b) FSI pipe model with tolerance $\epsilon = 10^{-3}$, with different reorderings. For the rigid wall, 38 cores are used. For the FSI, 48 cores are used.




Fig. 24: The spectrum of eigenvalues for a rigid pipe model with Neumann BC at the outlet. (top) eigenvalues obtained from four *local* matrices. Different colors are used to represent eigenvalues from different *local* matrices. (bottom) eigenvalues obtained from the *global* matrix.

Table 1:

L^2 -norm of velocity errors from different residual norm tolerances. The solution error is obtained from nodal velocity errors after 10 time steps, using a reference simulation with a tolerance up to the machine precision, $\epsilon = 10^{-12}$.

	$\epsilon = 10^{-1}$	$\epsilon = 10^{-3}$	$\epsilon = 10^{-6}$	$\epsilon = 10^{-9}$
<i>Err</i>	0.2235	0.0028	1.618×10^{-7}	1.107×10^{-9}

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 2:

The number of iterations and the portion of compute time consumed by the linear solver for the pipe benchmark problem with rigid walls. The number of linear solver iterations is counted for 10 time step calculations. t_T is the total compute time, t_{LS} is the compute time consumed by solving the linear system.

BIPN-RPC / $\epsilon = 10^{-3}$					
N_{bipn}	t_{LS}/t_T	N_{gmres}	t_{gmres}/t_{LS}	N_{cg}	t_{cg}/t_{LS}
129	89.5%	1475	5.3%	38582	77.6%
BIPN-RPC / $\epsilon = 10^{-6}$					
N_{bipn}	t_{LS}/t_T	N_{gmres}	t_{gmres}/t_{LS}	N_{cg}	t_{cg}/t_{LS}
847	95.1%	22699	17.25%	91470	55.7%
BIPN-RPC / $\epsilon = 10^{-9}$					
N_{bipn}	t_{LS}/t_T	N_{gmres}	t_{gmres}/t_{LS}	N_{cg}	t_{cg}/t_{LS}
1923	95.0%	247142	43.2%	791817	49.4%

$\epsilon = 10^{-3}$	GMRES		BICG	
	N_{gmres}	t_{LS}/t_T	N_{bicgs}	t_{LS}/t_T
PC				
Diag	155544	98.8%	N/A	N/A
Block-D	147768	98.7%	30027	94.4%
ILU	2768	85.3%	3104	89.6%
ILUT	1493	88.0%	2360	92.1%
IC	3636	40.5%	3290	41.3%
ML	1080	51.9%	1054	66.7%
$\epsilon = 10^{-6}$	GMRES		BICGS	
	N_{gmres}	t_{LS}/t_T	N_{bicgs}	t_{LS}/t_T
PC				
Diag	492752	99.9%	N/A	N/A
Block-D	491185	99.9%	110237	96.3%
ILU	6632	89.6%	6770	91.7%
ILUT	3382	89.8%	4045	92.9%
IC	8191	54.3%	7710	51.3%
ML	2526	64.6%	2376	76.6%
$\epsilon = 10^{-9}$	GMRES		BICGS	
	N_{gmres}	t_{LS}/t_T	N_{bicgs}	t_{LS}/t_T
PC				
Diag	1252448	100%	N/A	N/A
Block-D	1251025	100%	N/A	N/A
ILU	10066	91.0%	9766	92.2%
ILUT	5132	90.6%	5468	93.0%
IC	13728	59.8%	11022	53.4%
ML	3992	69.6%	3711	80.4%

Table 3:

The number of iterations and the portion of compute time consumed by the linear solver for the pipe benchmark problem with deformable walls. The number of linear solver iterations is counted for 10 time step calculations. t_T is the total compute time, t_{LS} is the compute time consumed by solving the linear system.

BIPN-RPC / $\epsilon = 10^{-3}$					
N_{bipn}	t_{LS}/t_T	N_{gmres}	t_{gmres}/t_{LS}	N_{cg}	t_{CG}/t_{LS}
608	78.6%	117377	77%	64865	15.6%
BIPN-RPC / $\epsilon = 10^{-6}$					
N_{bipn}	t_{LS}/t_T	N_{gmres}	t_{gmres}/t_{LS}	N_{cg}	t_{CG}/t_{LS}
1580	83.9%	318272	79.4%	145183	13.4%
BIPN-RPC / $\epsilon = 10^{-9}$					
N_{bipn}	t_{LS}/t_T	N_{gmres}	t_{gmres}/t_{LS}	N_{cg}	t_{CG}/t_{LS}
14872	83.3%	4087258	79.6%	874018	11.0%

$\epsilon = 10^{-3}$	GMRES		BICGS	
	N_{gmres}	t_{LS}/t_T	N_{bicgs}	t_{LS}/t_T
PC				
Diag	11283	71.9%	12094	66.3%
Block-D	9127	65.1%	9438	59.9%
ILU	3900	61.2%	3655	64.2%
ILUT	3493	70.6%	3333	70.9%
IC	4089	17.2%	3873	20.02%
ML	3735	31.74%	3799	44.43%
$\epsilon = 10^{-6}$	GMRES		BICGS	
	N_{gmres}	t_{LS}/t_T	N_{bicgs}	t_{LS}/t_T
PC				
Diag	33455	86.2%	45623	85.3%
Block-D	26311	81.7%	29589	77.3%
ILU	7699	70.1%	7338	74.9%
ILUT	5904	75.1%	5537	77.1%
IC	8241	28.7%	7819	32.0%
ML	7109	49.9%	7610	63.9%
$\epsilon = 10^{-9}$	GMRES		BICGS	
	N_{gmres}	t_{LS}/t_T	N_{bicgs}	t_{LS}/t_T
PC				
Diag	71884	90.25%	100431	90.6%
Block-D	55392	88.2%	49695	82.1%
ILU	13801	75.1%	13093	80.0%
ILUT	9848	76.7%	9404	79.3%
IC	14615	42.2%	13858	43.4%
ML	12178	59.8%	9587	78.4%

Table 4:

Number of nodes, elements and non-zero entries in tangent stiffness matrix for selected meshes in scaling studies.

	Fluid		
	# Nodes	# Tetrahedra	# Non-zeros
Mesh1	93,944	551,819	1,496,375
Mesh2	166,899	1,002,436	2,718,385
Mesh3	349,469	2,131,986	5,784,902
Mesh4	716,298	4,415,305	11,990,977
Mesh5	1,314,307	8,117,892	22,144,741
	Structure		Total
	# Nodes	# Tetrahedra	# Non-zeros
Mesh1	31,223	96,147	1,834,624
Mesh2	50,909	192,668	3,331,744
Mesh3	100,759	412,408	7,028,673
Mesh4	97,206	893,452	14,565,414
Mesh5	378,089	1,752,270	27,188,079

Table 5:

The 1-norm condition number estimates and five maximum and minimum eigenvalues (λ_j) of a *local* matrix \mathbf{A}_f and blocks \mathbf{K} , and \mathbf{L} without resistance BC.

A_f : Condition number= 1.293×10^6					
λ_j	1 st	2 nd	3 rd	4 th	5 th
Max($\times 10^{-3}$)	3.304	3.301	3.298	3.282	3.270
Min($\times 10^{-8}$)	0.748	1.042	1.348	1.524	1.672
K : Condition number=163					
λ_j	1 st	2 nd	3 rd	4 th	5 th
Max($\times 10^{-3}$)	3.304	3.301	3.298	3.282	3.270
Min($\times 10^{-5}$)	4.615	4.621	5.958	6.381	6.614
L : Condition number= 1.555×10^{18}					
λ_j	1 st	2 nd	3 rd	4 th	5 th
Max($\times 10^{-7}$)	9.053	8.180	8.085	8.083	7.983
Min($\times 10^{-9}$)	0.000	0.456	1.224	1.765	2.214

Table 6:

The 1-norm condition number estimates and five maximum and minimum eigenvalues (λ_j) of a *local* matrix $\tilde{\mathbf{A}}_f$ and $\tilde{\mathbf{K}}$ with a resistance BC.

$\tilde{\mathbf{A}}_f$: Condition number= 3.299×10^7					
λ_j	1 st	2 nd	3 rd	4 th	5 th
Max($\times 10^{-3}$)	86.525	3.253	3.244	3.239	3.185
Min($\times 10^{-8}$)	0.671	1.036	1.350	1.427	1.504
$\tilde{\mathbf{K}}$: Condition number= 5.883×10^3					
λ_j	1 st	2 nd	3 rd	4 th	5 th
Max($\times 10^{-3}$)	86.525	3.253	3.244	3.239	3.185
Min($\times 10^{-5}$)	2.878	2.886	3.457	3.981	4.002

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 7:

The 1-norm condition number estimate and five maximum and minimum eigenvalues (λ_j) of the raw *local* \mathbf{K}_s matrix.

\mathbf{K}_s : Condition number= 3.169×10^4					
λ_i	1 st	2 nd	3 rd	4 th	5 th
Max	7.706	7.469	7.389	7.325	7.294
Min($\times 10^{-3}$)	1.374	1.404	1.468	1.510	1.529

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 9:

The 1-norm condition number estimates and five maximum and minimum eigenvalues (λ_i) of the preconditioned *local* matrix $\tilde{\mathbf{A}}_f^*$, $\tilde{\mathbf{K}}^*$ and \mathbf{HK} .

$\tilde{\mathbf{A}}_f^*$: Condition number=9243					
λ_i	1 st	2 nd	3 rd	4 th	5 th
Max	55.15	2.262	2.247	2.247	2.243
Min($\times 10^{-1}$)	0.232	0.358	0.467	0.492	3.99
$\tilde{\mathbf{K}}^*$: Condition number=1862					
λ_i	1 st	2 nd	3 rd	4 th	5 th
Max	55.16	2.248	2.247	2.244	2.241
Min	0.217	0.384	0.392	0.396	0.399
\mathbf{HK} : Condition number=76					
λ_i	1 st	2 nd	3 rd	4 th	5 th
Max	2.248	2.247	2.244	2.241	2.241
Min	0.217	0.384	0.392	0.396	0.399