

1 **SOLVER COMPOSITION ACROSS THE PDE/LINEAR ALGEBRA**
2 **BARRIER***

3 ROBERT C. KIRBY[†] AND LAWRENCE MITCHELL[‡]

4 **Abstract.** The efficient solution of discretizations of coupled systems of partial differential equa-
5 tions (PDEs) is at the core of much of numerical simulation. Significant effort has been expended on
6 scalable algorithms to precondition Krylov iterations for the linear systems that arise. With few ex-
7 ceptions, the reported numerical implementation of such solution strategies is specific to a particular
8 model setup, and intimately ties the solver strategy to the discretization and PDE, especially when
9 the preconditioner requires auxiliary operators. In this paper, we present recent improvements in the
10 Firedrake finite element library that allow for straightforward development of the building blocks of
11 extensible, composable, preconditioners that decouple the solver from the model formulation. Our
12 implementation extends the algebraic composability of linear solvers offered by the PETSc library by
13 augmenting operators, and hence preconditioners, with the ability to provide any necessary auxiliary
14 operators. Rather than specifying up front the full solver configuration, tied to the model, solvers
15 can be developed independently of model formulation and configured at runtime. We illustrate with
16 examples from incompressible fluids and temperature-driven convection.

17 **Key words.** iterative methods, preconditioning, composable solvers, multiphysics

18 **AMS subject classifications.** 65N22, 65F08, 65F10

19 **1. Introduction.** For over a decade now, domain-specific languages for numer-
20 ical partial differential equations (henceforth PDEs) such as Sundance [30, 29], FEn-
21 iCS [28], and now Firedrake [40] have enabled users to efficiently generate algebraic
22 systems from a high-level description of the variational problems. Both FEniCS and
23 Firedrake make use of the Unified Form Language [1] as a description language for
24 the weak forms of PDEs, converting it into efficient low-level code for form evalua-
25 tion. They also share a Python interface that, for the intersection of their feature
26 sets, is nearly source-compatible. These high-level PDE codes succeed by connecting
27 a rich description language for PDEs to effective lower-level solver libraries such as
28 PETSc [5, 4] or Trilinos [21] for the requisite, and performance-critical, numerical
29 (non)linear algebra.

30 These high-level PDE projects utilize the solver packages in an essentially *unidi-*
31 *rectional* way: the residuals are evaluated, Jacobians formed, and are then handed off
32 to mainly algebraic techniques. Hence, the codes work at their best when (composi-
33 tions of) existing black-box matrix techniques effectively solve the algebraic systems.
34 However, in many situations the best preconditioners require additional structure be-
35 yond a purely algebraic (matrix and vector-level) problem description. Many of the
36 preconditioners for block systems based on block factorizations require discretizations
37 of differential operators not contained in the original problem. These include the
38 pressure-convection-diffusion (PCD) approximation for Navier-Stokes [25, 16], and
39 preconditioners for models of phase separation [24, 19]. An alternate approach to

*
Funding: RCK is supported by the National Science Foundation Computing and Communica-
tions Foundations grant number 1525697 and also acknowledges support from the PRISM Center at
Imperial College, London [ESPRC grant number EP/L000407/1] for sabbatical support. LM is sup-
ported by the Engineering and Physical Sciences Research Council [grant number EP/M011054/1].
This work used the ARCHER UK National Supercomputing Service (<http://www.archer.ac.uk>).

[†]Department of Mathematics, Baylor University, One Bear Place #97328, Waco, TX 76798-7328,
USA (robert_kirby@baylor.edu)

[‡]Department of Computing and Department of Mathematics, Imperial College London, South
Kensington Campus, London SW7 2AZ, UK (lawrence.mitchell@imperial.ac.uk)

40 derive preconditioners for block systems is to use arguments from functional analysis
 41 to arrive at block-diagonal preconditioners. While these are often representable as
 42 the inverse of an assembled operator, in some cases, a mesh and parameter indepen-
 43 dent preconditioner that arises from such an analysis requires the action of the sum
 44 of inverses. An example is the preconditioner suggested in [32, example 4.2] for the
 45 time-dependent Stokes problem.

46 While a high-level PDE engine makes it possible to obtain these new operators
 47 at low user cost, additional care is required to develop a clean, extensible interface.
 48 For example, the PCD preconditioner has been implemented using Sundance and
 49 Playa [23], although the resulting code tightly fused the description of the problem
 50 with a highly specialized specification of the preconditioner. Similarly, in the FEniCS
 51 project, `cbc.block` [31] allows the model developer to write complex block precondi-
 52 tioners as a composition of high-level “symbolic” linear algebra operations; Trilinos
 53 provides similar functionality through Teko [12]. However, in these codes one must
 54 specify up front how to perform the block decomposition. Switching to a different pre-
 55 conditioner requires changing the model code, and there is no high-level manipulation
 56 of variational problems within the blocks. Ideally, one would like a mechanism to im-
 57 plement the specialized preconditioner as a separate component, leaving the original
 58 application code essentially unchanged.

59 *Extensibility* of fundamental types such as solvers, preconditioners, and matrices
 60 has long been a main concern of the PETSc project. For example, the action of a finite
 61 difference stencil applied to a vector can be wrapped behind a matrix “shell” inter-
 62 face and used interchangeably with explicit sparse matrices for many purposes. Users
 63 can similarly provide custom types of Krylov methods or preconditioners. Thanks to
 64 `petsc4py` [13], such extensions can also be implemented in Python as well as C. More-
 65 over, PETSc provides powerful tools to switch between (compositions of) existing and
 66 custom tools either in the application source code or through command-line options.

67 In this work, we enable the rapid development of high-performance precondition-
 68 ers as PETSc extensions using Firedrake and `petsc4py`. To facilitate this, we have
 69 developed a custom matrix type that embeds the complete Firedrake problem de-
 70 scription (UFL variational forms, function spaces, meshes, etc) in a Python context
 71 accessible to PETSc. As a happy byproduct, this enables low-memory matrix-free
 72 evaluation of matrix-vector products. This also allows us to produce PETSc pre-
 73 conditioner in `petsc4py` that act on this new matrix type, accessing the PDE-level
 74 information as needed. For example, a PCD preconditioner can access the meshes
 75 and function spaces to create bilinear forms for, and hence assemble, the needed
 76 mass, stiffness, and convection-diffusion operators on the pressure space and PETSc
 77 `KSP` (linear solver) contexts for the inverses. Moreover, once such preconditioners are
 78 available in a globally importable module, it is now possible to use them instead of
 79 existing algebraic preconditioners by a straightforward runtime modification of solver
 80 configuration options. So, we use our PDE language not only to generate problems
 81 to feed to the solver, but also to extend that solver’s capabilities.

82 Our discussion and implementation will focus on Firedrake as the top-level PDE
 83 library and PETSc as the solver library. Firedrake already relies heavily on PETSc
 84 through `petsc4py` and has a nearly pure Python implementation. Provided one is
 85 content with the Python interface, it should not be terribly difficult to adapt these
 86 techniques for use in FEniCS. Regarding solver libraries, the idiom and usage of
 87 Trilinos and PETSc (if not their actual capabilities) differ considerably, so we make
 88 no speculation as to the difficulties associated with adapting our techniques in that
 89 direction.

90 In the rest of the paper, we set up certain model problems in [section 2](#). After
 91 this, in [section 3](#) we survey certain algorithms that go beyond the current mode of
 92 algebraically preconditioning assembled matrices. These include matrix-free methods,
 93 Schwarz-type preconditioners, and preconditioners that require auxiliary differential
 94 operators. It turns out that a proper implementation of the first of these, matrix-free
 95 methods, provides a very clean way to communicate PDE-level problem information
 96 between PETSc matrices and custom preconditioners, and we describe the implemen-
 97 tation of this and relevant modifications to Firedrake in [section 4](#). Finally, we give
 98 examples demonstrating the efficacy of our approach to the model problems of interest
 99 in [section 5](#).

100 2. Some model applications.

101 **2.1. The Poisson equation.** It is helpful to fix some target applications and
 102 describe things we would like to expedite within our top-level code.

103 A usual starting point is to consider a second-order scalar elliptic equation. Let
 104 $\Omega \subset \mathbb{R}^d$, where $d = 1, 2, 3$, be a domain with boundary Γ . We let $\kappa : \Omega \rightarrow \mathbb{R}^+$ be some
 105 positive-valued coefficient. On the interior of Ω , we seek a function u satisfying

$$106 \quad (1) \quad -\nabla \cdot (\kappa \nabla u) = f$$

107 subject to the boundary condition $u = u_{\Gamma_D}$ on Γ_D and $\nabla u \cdot n = g$ on Γ_N .

108 After the usual technique of multiplying by a test function and integrating by
 109 parts, we reach the weak form of seeking $u \in V_{\Gamma} \subset V$ such that

$$110 \quad (2) \quad (\kappa \nabla u, \nabla v) = (f, v) - \left\langle g, \frac{\partial v}{\partial n} \right\rangle$$

111 for all $v \in V_0 \subset V$, where V is the finite element space, V_0 the subspace with vanishing
 112 trace on Γ_D . Here (\cdot, \cdot) denotes the standard L^2 inner product over Ω , and $\langle \cdot, \cdot \rangle$ that
 113 over Γ .

114 The finite element method leads to a linear system:

$$115 \quad (3) \quad Au = f,$$

116 where A is symmetric and positive-definite (positive semi-definite if $\Gamma_D = \emptyset$), and
 117 the vector f includes both the forcing term and contributions from the boundary
 118 conditions.

119 **2.2. The Navier-Stokes equations.** Moving beyond the simple Poisson oper-
 120 ator, the incompressible Navier-Stokes equations provide challenge.

$$121 \quad (4a) \quad -\frac{1}{\text{Re}} \Delta \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} + \nabla p = 0,$$

$$122 \quad (4b) \quad \nabla \cdot \mathbf{u} = 0$$

124 together with suitable boundary conditions.

125 Among the diverse possible methods, we shall focus here on inf-sup stable mixed
 126 finite element spaces such as Taylor-Hood [9]. This is merely for simplicity of explica-
 127 tion and does not represent a limitation of our approach or implementation. Taking
 128 V_{Γ} to be subset of the discrete velocity space satisfying any strongly imposed bound-
 129 ary conditions and W the pressure space, we have the weak form of seeking \mathbf{u}, p in

130 $V_\Gamma \times W$ such that

$$131 \quad (5a) \quad \frac{1}{\text{Re}} (\nabla \mathbf{u}, \nabla \mathbf{v}) + (\mathbf{u} \cdot \nabla \mathbf{u}, \mathbf{v}) - (p, \nabla \cdot \mathbf{v}) = 0,$$

$$133 \quad (5b) \quad (\nabla \cdot \mathbf{u}, w) = 0$$

134 for all $\mathbf{v}, w \in V_0 \times W$, where V_0 is the velocity subspace with vanishing Dirichlet
135 boundary conditions.

136 Relative to the Poisson equation, we now have several additional challenges. The
137 nonlinearity may be addressed by Newton linearization, and UFL provides automatic
138 differentiation to produce the Jacobian. We also have multiple finite element spaces,
139 one of which is vector-valued. Further, for each nonlinear iteration, the required linear
140 system is larger and more complicated, a block-structured saddle point system of the
141 form

$$142 \quad (6) \quad \begin{bmatrix} F & -B^t \\ B & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ p \end{bmatrix} = \begin{bmatrix} f_1 \\ f_2 \end{bmatrix}.$$

143 We discuss approaches to preconditioning iterative methods for this system in [sec-](#)
144 [tion 3](#).

145 **2.3. Rayleigh-Bénard convection.** Many applications rely on coupling other
146 processes to the Navier-Stokes equations. For example, Rayleigh-Bénard convec-
147 tion [11] includes thermal variation in the fluid, although we take the Boussinesq
148 approximation that temperature variations affect the momentum balance only as a
149 buoyant force. We have, after nondimensionalisation,

$$150 \quad (7a) \quad -\Delta \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} + \nabla p = -\frac{\text{Ra}}{\text{Pr}} T g \hat{z},$$

$$151 \quad (7b) \quad \nabla \cdot \mathbf{u} = 0,$$

$$153 \quad (7c) \quad -\text{Pr} \Delta T + \mathbf{u} \cdot \nabla T = 0,$$

154 where Ra is the Rayleigh number, Pr is the Prandtl number, g is the acceleration due
155 to gravity, and \hat{z} the upward-pointing unit vector. The problem is usually posed on
156 rectangular domains, with no-slip boundary conditions on the fluid velocity. The tem-
157 perature boundary conditions typically involve imposing a unit temperature difference
158 in one direction with insulating boundary conditions in the others.

159 After discretization and Newton linearization, one obtains a block 3×3 system

$$160 \quad (8) \quad \begin{bmatrix} F & -B^t & M_1 \\ B & 0 & 0 \\ M_2 & 0 & K \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ p \\ T \end{bmatrix} = \begin{bmatrix} f_1 \\ f_2 \\ f_3 \end{bmatrix}.$$

161 Here, the F and B matrices are as obtained in the Navier-Stokes equations (with
162 $\text{Re} = 1$). The M_1 and M_2 terms arise from the temperature/velocity coupling, and
163 K is the convection-diffusion operator for temperature.

164 Alternately, letting

$$165 \quad (9a) \quad N = \begin{bmatrix} F & -B^t \\ B & 0 \end{bmatrix},$$

$$166 \quad (9b) \quad \widetilde{M}_1 = \begin{bmatrix} M_1 \\ 0 \end{bmatrix},$$

$$168 \quad (9c) \quad \widetilde{M}_2 = [M_2 \quad 0],$$

169 we can write the stiffness matrix as block 2×2 matrix

$$170 \quad (10) \quad \begin{bmatrix} N & \widetilde{M}_1 \\ \widetilde{M}_2 & K \end{bmatrix}.$$

171 Formulating the matrix in this way allows us to consider composing some (possibly
172 custom) solver technique for Navier-Stokes with other approaches to handle the
173 temperature equation and coupling.

174 **3. Solution techniques.** Via UFL, Firedrake has a succinct, high-level description
175 of these equations and can readily linearize and assemble discrete operators.
176 When efficient techniques for the discrete system exist within PETSc, obtaining the
177 solution is as simple as providing the proper options. When direct methods are applicable,
178 simple options like `-ksp_type preonly -pc_type lu` suffice – possibly augmented with the
179 selection of a package to perform the factorization like MUMPS [2] or UMFPACK [14].
180 Similarly, when iterative methods with black-box preconditioners such as incomplete
181 factorization or algebraic multigrid fit the bill, options such as `-ksp_type cg -pc_type hypre`
182 work. PETSc also provides many block preconditioner mechanisms via `FieldSplit`,
183 allowing users to specify PETSc solvers for inverting the relevant blocks [10].
184 Firedrake automatically enables this by specifying index sets for each function space,
185 passing the information to PETSc when it initializes the solver. A key feature of PETSc
186 is that these choices can be made at runtime via options, *without* modifying the user
187 code that specifies the PDE to solve.

188 As we stated in the introduction, however, many techniques for preconditioning
189 require information beyond what can be learned by an inspection of matrix entries
190 and user-specified options. It is our goal now to survey some of these techniques in
191 more detail, after which we describe our implementation of custom PETSc preconditioners
192 to utilize application-specific problem descriptions in a clean, efficient, and
193 user-friendly way.

194 **3.1. Matrix-free methods.** Switching from a low order method to a higher-order
195 simply requires changing a parameter in the top-level Firedrake application code.
196 However, such a small change can profoundly affect the overall performance footprint.
197 Assembly of stiffness matrices becomes more expensive, both in terms of time and space,
198 as the order increases. An alternative, that does not have the same space and time
199 constraints is to use a *matrix-free* implementation of the matrix-vector product.
200 This is sufficient for Krylov methods, although not for algebraic preconditioners
201 requiring matrix entries.

202 Rather than producing a sparse matrix A , one provides a function that, given
203 a vector x , computes the product Ax . Abstractly, consider a bilinear form a on a
204 discrete space V with basis $\{\psi_i\}_{i=1}^N$. The $N \times N$ stiffness matrix $A_{ij} = a(\psi_j, \psi_i)$
205 can be applied to a vector x as follows. Any vector x is isomorphic to some function
206 $u \in V$ via the identification $x \leftrightarrow u = \sum_{j=1}^N x_j \psi_j$. Then, via linearity,

$$207 \quad (11) \quad \begin{aligned} (Ax)_i &= \sum_{j=1}^N A_{ij} x_j = \sum_{j=1}^N a(\psi_i, \psi_j) x_j \\ &= a(\psi_i, \sum_{j=1}^N x_j \psi_j) = a(\psi_i, u). \end{aligned}$$

208 Just like matrices or load vectors, this can be computed by assembling elementwise
209 contributions in the standard way, considering u to be just some given member of V .

210 In the presence of strongly-enforced boundary conditions, the bilinear form acts
 211 on a subspace $V_0 \subset V$. When a matrix is explicitly assembled, one typically either
 212 edits (or removes) rows and columns to incorporate the boundary conditions. Care
 213 must be taken in enforcing the boundary conditions to ensure that the matrix-free
 214 action agrees with the matrix that would have been assembled.

215 Typically, such an approach has a much lower startup cost than an explicit sparse
 216 matrix since no assembly is required. Forgoing an assembled matrix also saves con-
 217 siderably on memory usage. Moreover, the arithmetic intensity (ai) of matrix-free
 218 operator application is almost always higher than that of an assembled matrix (sparse
 219 matrix multiplication has $\text{ai} \approx 1/6$ [20]). Matrix-free methods are therefore an in-
 220 creasingly good match to modern memory bandwidth-starved hardware, where the
 221 balanced arithmetic intensity is $\text{ai} \approx 10$. The algorithmic complexity is either the same
 222 ($\mathcal{O}(p^{2d})$ for degree p elements in d dimensions), or better ($\mathcal{O}(p^{d+1})$) if a structured
 223 basis can be exploited through sum factorisation. On simplex elements, the latter op-
 224 timisation is not currently available through the form compiler in Firedrake. Thus we
 225 will expect our matrix-free operator applications to have the same algorithmic scaling
 226 as assembled matrices (though with different constant factors). If we can exploit the
 227 vector units in modern processors effectively, we can expect that matrix-free applica-
 228 tions will be at least competitive with, and often faster than, assembled matrices (for
 229 example [33] demonstrate significant benefits, relative to assembled matrices, for Q_2
 230 operator application on hexahedra).

231 **3.2. Preconditioning high-order discretizations: additive Schwarz.** Matrix-
 232 free methods preclude algebraic preconditioners such as incomplete factorization or
 233 algebraic multigrid. Depending on the available smoothers, if a mesh hierarchy is
 234 available, geometric multigrid is a possibility [7, 8]. Here, we discuss a family of ad-
 235 ditive Schwarz methods. Originally proposed by Pavarino in [37, 38], these methods
 236 fall within the broad family of subspace correction methods [44].

237 These two-level methods decompose the finite element space into a low order space
 238 on the original mesh and the high-order space restricted to local pieces of the mesh,
 239 such as patches of cells around each vertex. Any member of the original finite element
 240 space can be written as a combination of items from this collection of subspaces,
 241 although the decomposition in this case is certainly not orthogonal. One obtains a
 242 preconditioner for the original finite element operator by additively combining the
 243 (possibly approximate) inverses of the restrictions of the original operator to these
 244 spaces. Schöberl [42] showed for the symmetric coercive case that the preconditioned
 245 system has eigenvalue bounds independent of both mesh size and polynomial degree
 246 and gave computational examples from elasticity confirming the theory. Although not
 247 covered by Schöberl’s analysis, these methods have also been applied with success to
 248 the Navier-Stokes equations [39].

249 This approach is *generic* in that it can be attempted for any problem. Given a
 250 bilinear form over a function space of degree k , one can programmatically build the
 251 lowest-order instance of the function space and set up the vertex patches for the mesh.
 252 Then, one can easily modify the bilinear form to operate on the new subspaces and
 253 perform the subspace correction. We have developed such a generic implementation,
 254 parametrized over the UFL problem description.

255 One drawback of this method is the relatively high memory cost of storing the
 256 patch-wise Cholesky or LU factors, especially at high order and in 3D. One may further
 257 decompose the local patch spaces through “spider vertices” to reduce the memory
 258 required and still retain a powerful method [42]. Such refinements are possible within

our framework, although we have not pursued them to date.

3.3. Block preconditioners and Schur complement approximations. Having motivated matrix-free methods and preconditioners for higher-order discretizations in the simple case of the Poisson operator, we now return to the Navier-Stokes equations introduced earlier. In particular, we are interested in preconditioners for the Jacobian stiffness matrix (6).

Block factorization of the system matrix provides a starting point for many powerful preconditioners [6, 16, 18]. Consider the block LDU factorization of the system matrix in (6) as

$$(12) \quad \begin{bmatrix} F & -B^t \\ B & 0 \end{bmatrix} = \begin{bmatrix} I & 0 \\ BF^{-1} & I \end{bmatrix} \begin{bmatrix} F & 0 \\ 0 & S \end{bmatrix} \begin{bmatrix} I & -F^{-1}B^t \\ 0 & I \end{bmatrix},$$

where I is the identity matrix of the proper size and $S = BF^{-1}B^t$ is the Schur complement. The inverse of this matrix is then given by

$$(13) \quad \begin{bmatrix} F & -B^t \\ B & 0 \end{bmatrix}^{-1} = \begin{bmatrix} I & F^{-1}B^t \\ 0 & I \end{bmatrix} \begin{bmatrix} F^{-1} & 0 \\ 0 & S^{-1} \end{bmatrix} \begin{bmatrix} I & 0 \\ -BF^{-1} & I \end{bmatrix}.$$

Since this is the exact inverse, applying it in a preconditioning phase leads to Krylov convergence in a single iteration if all blocks are inverted exactly. Note that inverting the Schur complement matrix S either requires assembling it as a dense matrix or else using a Krylov method where the matrix action is computed implicitly via two matrix-vector products and an inner solve to produce F^{-1} .

Two kinds of approximations lead to more practical methods. For one, it is possible to neglect either or both of the triangular factors. This gives a structurally simpler preconditioner, at the cost (assuming exact inversion of S) of a slight increase in the iteration count. For example, it is common to use only the lower triangular part of the matrix, giving a preconditioning matrix of the form

$$(14) \quad P = \begin{bmatrix} F & 0 \\ B & S \end{bmatrix}$$

which has inverse

$$(15) \quad P^{-1} = \begin{bmatrix} F^{-1} & 0 \\ 0 & S^{-1} \end{bmatrix} \begin{bmatrix} I & 0 \\ -BF^{-1} & I \end{bmatrix}.$$

Using P as a left preconditioner, $P^{-1}A$ is readily seen to give a unit upper triangular matrix, and it is known that GMRES will converge in two (very expensive) iterations since the resulting preconditioned matrix system has a quadratic minimal polynomial [36].

Given the cost of inverting S , it is also desirable to devise a suitable approximation. A simple approach is to use a pressure mass matrix, which gives mesh-independent but rather large eigenvalue bounds [17]. More sophisticated approximations are well-documented in the literature [16]. For our purposes, we will consider one in particular, the *pressure convection-diffusion* (hence PCD) preconditioner [25, 15]. It is based on the approximation

$$(16) \quad S^{-1} = (BF^{-1}B^t)^{-1} \approx K_p^{-1}F_pM_p^{-1} \equiv X^{-1},$$

296 where K_p is the Laplace operator acting on the pressure space, M_p is the mass matrix,
 297 and F_p is a discretization of the convection-diffusion operator

$$298 \quad (17) \quad \mathcal{L}p \equiv -\frac{1}{Re}\Delta p + \mathbf{u}_0 \cdot \nabla p,$$

299 with \mathbf{u}_0 the velocity at the current Newton iterate. Although this requires solving
 300 linear systems, the mass and stiffness matrices are far cheaper to invert than F .

301 While one could use this approximation to precondition a Krylov solver for S ,
 302 it is far more typical to replace S^{-1} with X^{-1} . For example, using the triangular
 303 preconditioner (14) gives the further approximation in a block preconditioner:

$$304 \quad (18) \quad \tilde{P}^{-1} = \begin{bmatrix} F^{-1} & 0 \\ 0 & X^{-1} \end{bmatrix} \begin{bmatrix} I & 0 \\ -BF^{-1} & I \end{bmatrix} = \begin{bmatrix} F^{-1} & 0 \\ 0 & K_p^{-1}F_pM_p^{-1} \end{bmatrix} \begin{bmatrix} I & 0 \\ -BF^{-1} & I \end{bmatrix}.$$

305 Although bypassing the solution of the Schur complement system increases the outer
 306 iteration count, it typically results in a much more efficient overall method. Also,
 307 note that only the action of the off-diagonal blocks is required for the preconditioner
 308 so that a matrix-free treatment is appropriate.

309 Preconditioning strategies for the Navier-Stokes equations can quickly find their
 310 way into problems coupling other processes to fluids. We return now to the Bénard
 311 convection stiffness matrix (10), where N is itself the Navier-Stokes stiffness matrix
 312 in (6). Block preconditioners based on this formulation, replacing N^{-1} with a very
 313 inexact solve via PCD-preconditioned GMRES, proved more effective than techniques
 314 based on 3×3 preconditioners [22]. Here, we present a lower-triangular block pre-
 315 conditioner rather than the upper-triangular one in [22] with similar practical results.

316 A block Gauss-Seidel preconditioner for (10) can be taken as

$$317 \quad (19) \quad P = \begin{bmatrix} N & 0 \\ \widetilde{M}_2 & K \end{bmatrix},$$

318 the inverse of which requires evaluation of N^{-1} and K^{-1} :

$$319 \quad (20) \quad P^{-1} = \begin{bmatrix} N^{-1} & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} I & 0 \\ -\widetilde{M}_2 & I \end{bmatrix} \begin{bmatrix} I & 0 \\ 0 & K^{-1} \end{bmatrix}.$$

320 Replacing these inverses with approximations/preconditioners \tilde{N}^{-1} and \tilde{K}^{-1} gives

$$321 \quad (21) \quad \tilde{P}^{-1} = \begin{bmatrix} \tilde{N}^{-1} & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} I & 0 \\ -\widetilde{M}_2 & I \end{bmatrix} \begin{bmatrix} I & 0 \\ 0 & \tilde{K}^{-1} \end{bmatrix}.$$

322 At this point, replacing \tilde{N}^{-1} with the block preconditioner (18) recovers a block
 323 lower-triangular 3×3 preconditioner:

$$324 \quad (22) \quad \tilde{P}^{-1} = \begin{bmatrix} F^{-1} & 0 & 0 \\ 0 & K_p^{-1}F_pM_p^{-1} & 0 \\ 0 & 0 & I \end{bmatrix} \begin{bmatrix} I & 0 & 0 \\ -BF^{-1} & I & 0 \\ 0 & 0 & I \end{bmatrix} \begin{bmatrix} I & 0 & 0 \\ 0 & I & 0 \\ -M_2 & 0 & I \end{bmatrix} \begin{bmatrix} I & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & \tilde{K}^{-1} \end{bmatrix}.$$

325 **4. Implementation.** The core object in our implementation is an appropriately
 326 designed “implicit” matrix that provides matrix-vector actions and also makes PDE-
 327 level discretization information available to custom preconditioners within PETSc.
 328 Here, we describe this class, how it interacts with both Firedrake and PETSc, and
 329 how it provides the requisite functionality. Then, we demonstrate how it cleanly
 330 provides the proper information for custom preconditioners.

331 **4.1. Implicit matrices.** First, we note that Firedrake deals with matrices at
 332 two different levels. A Firedrake-level `Matrix` instance maintains symbolic information
 333 (the bilinear form, boundary conditions). It in turn contains a PETSc `Mat` (typically
 334 in some sparse format), which is used when creating solvers.

335 Our implicit matrices mimic this structure, adding an `ImplicitMatrix` sibling
 336 class to the existing `Matrix`, lifting shared features into a common `MatrixBase` class.
 337 Where the `ImplicitMatrix` differs is that its PETSc `Mat` now has type `python` (rather
 338 than a normal sparse format such as `aij`). To provide the appropriate matrix-vector
 339 actions, the `ImplicitMatrix` instance provides an `ImplicitMatrixContext` instance
 340 to the PETSc `Mat`¹. This context object contains the PDE-level information – the
 341 bilinear form and boundary conditions – necessary to implement matrix-vector prod-
 342 ucts. Moreover, this context object enables building custom preconditioners since it
 343 is available from within the “low-level” PETSc `Mat`.

344 UFL’s `adjoint` function, which reverses the test and trial function in a bilinear
 345 form, also makes it straightforward to provide the action of the matrix transpose,
 346 needed in some Krylov methods [41, §7.1]. The implicit matrix constructor simply
 347 stashes the action of the original bilinear form and its adjoint, and the multiplication
 348 and transposed multiplication are nearly identical using Firedrake’s `assemble` method
 349 with boundary conditions appropriately enforced.

350 We enable `FieldSplit` preconditioners on implicit matrices by means of over-
 351 loading submatrix extraction. The PETSc interface to submatrix extraction does not
 352 presuppose any particular block structure. Instead, the function receives integer index
 353 sets corresponding to rows and columns to be extracted into the submatrix. Since
 354 the PDE-level description operates at the level of fields, we only support extraction of
 355 submatrices that correspond to (some subset of) the fields that the matrix contains.
 356 Our method determines whether a provided index set is a concatenation of a subset
 357 of the index sets defining the different fields and returns the list of integer labels of
 358 the fields in the subset. While this implementation compares index sets by value
 359 and therefore increases in expense as the number of per-process degrees of freedom
 360 increases, it must only be carried out once per solve (be it linear or non-linear), since
 361 the index set structure does not change. We have not found it to be a measureable
 362 fraction of the solution time in our implementation.

363 Splitting implicit matrices offers an efficient alternative to splitting already-assembled
 364 sparse matrices. Currently, splitting a standard assembled matrix into blocks requires
 365 the allocation and copying of the subblocks. While PETSc also includes a “nested”
 366 matrix type (essentially an array of pointers to matrices), collecting multiple fields
 367 into a single block (e.g. the pressure and velocity in Bénard convection) requires that
 368 the user code state up front the order in which nesting occurs. This would mean
 369 that editing/recompilation of the code is necessary to switch between precondition-
 370 ing approaches that use different variable splittings, contrary to our goal of efficient
 371 high-level solver configuration and customization.

372 The typical user interface in Firedrake involves interacting with PETSc via a
 373 `VariationalSolver`, which takes charge of configuring and calling the PETSc linear
 374 and nonlinear solvers. It allocates matrices and sets the relevant callback functions
 375 for Jacobian and residual evaluation to be used inside `SNES` (PETSc’s nonlinear solver
 376 object). Switching between implicit and standard sparse matrices is now facilitated
 377 through additional PETSc database options, so that the type of Jacobian matrix

¹Owing to the cross-language issues and lack of proper inheritance mechanisms in C, this is the standard way of implementing new types from Python in PETSc.

378 is set with `-mat_type` and the, possibly different, preconditioner matrix type with
 379 `-pmat_type`. This latter option facilitates using assembled matrices for the matrix-
 380 vector product, while still providing PDE-level information to the solver. In this way,
 381 enabling matrix-free methods simply requires an options change in Firedrake and no
 382 other user modification.

383 **4.2. Preconditioners.** It is helpful to briefly review certain aspects of the
 384 PETSc formalism for preconditioners. One can think of (left) preconditioning as
 385 converting a linear system

$$386 \quad (23) \quad b - Ax = 0$$

387 into an equivalent system

$$388 \quad (24) \quad \hat{P}(b - Ax) = 0,$$

389 where $\hat{P}(\cdot)$ applies an approximation of the inverse of the preconditioning matrix P
 390 to the residual².

391 Then, given a current iterate x_i , we have the residual

$$392 \quad (25) \quad r_i = b - Ax_i.$$

393 PETSc preconditioners are specified to act on residuals, so that $\hat{P}(r_i)$ then gives an
 394 approximation to the error $e_i = x - x_i$. This enables sparse direct methods to act as
 395 preconditioners, converting the residual into the exact (up to roundoff error) residual,
 396 and direct solvers nonetheless conform to the KSP interface (e.g. `-ksp_type preonly`
 397 `-pc_type lu`).

398 PETSc preconditioners are built in terms of both the system matrix A and a possi-
 399 bly different “preconditioning matrix” A_p (for example, preconditioning a convection-
 400 diffusion operator with the Laplace operator). So then, $\hat{P} = \hat{P}(A, A_p)$ is a method for
 401 constructing an (approximation to) the inverse of A . Preconditioner implementations
 402 must provide PETSc with an `apply` method that computes $y \leftarrow \hat{P}x$. Creation of the
 403 data (for example, an incomplete factorization) necessary to apply the preconditioner
 404 is carried out in a `setUp` method.

405 Firedrake now provides Python-level scaffolding to expedite the implementation
 406 of preconditioners that act on implicit matrices. Instead of manipulating matrix
 407 entries like ILU or algebraic multigrid, these preconditioners use the UFL problem
 408 description from the Python context contained in the incoming matrix P to do what
 409 is needed. Hence, these preconditioners can be parametrized not over particular
 410 matrices, but over bilinear forms. To demonstrate the generality of our approach, we
 411 have implemented several such examples.

412 **4.2.1. Assembled preconditioners.** While one can readily define block pre-
 413 conditioner using implicit matrices, the best methods for inverting the diagonal
 414 blocks may in fact be algebraic. This illustrates a critical use case of our simplest
 415 preconditioner acting on implicit matrices. We have defined a generic preconditioner
 416 `AssembledPC` whose `setUp` method simply forces the assembly of an underlying bi-
 417 linear form and then sets up a sub-preconditioner (typically an algebraic one) acting
 418 on the sparse matrix. Then, the `apply` method simply forwards to that of the sub-
 419 preconditioner. For example, to use an implicit matrix-vector product but incomplete
 420 factorization on an assembled matrix for the preconditioner, one might use options
 421 like

²We use this notation since it possible that \hat{P} is not a linear operator.

422
423
424
425
426

```

-mat_type matfree
-pc_type python
-pc_python_type firedrake.AssembledPC
-assembled_pc_type ilu

```

428
429
430
431

As mentioned, `FieldSplit` preconditioners provide a critical use case, enabling one to leave the overall matrix implicit, and assemble only those blocks that are required. In particular, the off-diagonal blocks never require assembly, and this can result in significant memory savings.

432
433
434
435
436
437
438
439
440
441
442
443

4.2.2. Schur complement approximations. Our next example, Schur complement approximations, is more specialized but very relevant to the problems in fluid mechanics expressed above. PETSc provides two pathways to define preconditioners for the Schur complement, such as (16). Within the source code, one may pass to the function `PCFieldSplitSetSchurPre` a matrix which will be used by a preconditioner to construct an approximation to the Schur complement. Alternatively, PETSc can automatically construct some approximations that may be obtained by algebraic manipulations of the original operator (such as the SIMPLE or LSC approximations [16]). While the latter may be configured using only runtime options, the former requires that the user pick apart the solver and call `PCFieldSplitSetSchurPre` on the appropriate PC object. Enabling this preconditioning option, or incorporating it into larger coupled systems requires modification of the model source code.

444
445
446
447
448
449
450
451
452
453
454

Since our implicit matrices and their subblocks contain the UFL problem specification, a preconditioner acting on the Schur complement block has complete freedom to utilize the UFL bilinear form to set up auxiliary operators. We have implemented two Schur complement approximations suitable for incompressible flow, an inverse mass matrix and the PCD preconditioner, both of which follow a similar pattern. The `setUp` function extracts the pressure function space from the UFL bilinear form and defines and assembles bilinear forms for the auxiliary operators. It also defines user-configurable KSP contexts as needed (e.g. for the K_p and M_p operators in (16)). The PCD preconditioner also requires a subsequent update phase in which the F_p matrix is updated as the Jacobian evolves. The `apply` method simply performs the correct combination of matrix-vector products and linear solves.

455
456
457
458

The high-level Python syntax of `petsc4py` and `Firedrake` combine to allow a very concise implementation in these cases. In the case of PCD, we specify the initial and subsequent setup methods plus application method in less than 150 lines of code, including Python doc strings and hooks into the PETSc viewer system.

459
460
461
462
463
464
465
466
467

User data. The PCD preconditioner requires a very slight modification of the application code. In particular, UFL does not expose named parameters. That is, one may not ask the variational problem what the Reynolds number is. Also, it is not obvious to the preconditioner which piece of the current Newton state corresponds to the velocity, which is needed in constructing F_p . To address such difficulties, `Firedrake`'s `VariationalSolver` classes can take an arbitrary Python dictionary of user data, which is available inside the implicit matrix, and hence to the preconditioners. This facility requires documentation, but fits with the general PETSc idiom of allowing all callbacks to user code to provide a generic "application context".

468
469
470
471
472

4.2.3. Additive Schwarz. Our additive Schwarz implementation requires both more involved UFL manipulation and low-level implementation details. We have implemented it as a Python preconditioner that defers to a PETSc `PCCOMPOSITE` to perform the composition, but extracts and manipulates the symbolic description of the problem to create two Python preconditioners, one for the P_1 subproblem and

473 one for the local, high-degree, patch problems.

474 The P_1 preconditioner requires us to construct the P_1 discretization of the given
 475 operator, plus restriction and prolongation operators between the global P_k and P_1
 476 spaces. UFL provides a utility to make the first of these straightforward – we just
 477 replace the test and trial functions in the original expression graph with test and
 478 trial functions taken from the P_1 space on the same mesh. The second is a bit
 479 more involved. We rely on the fact that the P_1 basis functions on a cell are naturally
 480 embedded in the P_k space, and hence their interpolant in P_k is exact. Using FIAT [26]
 481 to construct this interpolant on a single cell, we then generate a cell kernel that is
 482 called for every coarse element in the mesh to produce the prolongation operator as
 483 a sparse matrix. Optionally, this can also occur in a matrix-free fashion.

484 Setting up and solving the patch problems presents more complications. During
 485 a startup phase, we must query the mesh to discover and store the cells in each vertex
 486 patch. At this time, we also construct the sets of global degrees of freedom involved
 487 in each patch, setting up indirections between patch-level and processor-level degrees
 488 of freedom.

489 Our implementation leverages PETSc’s DMPlex representation of computational
 490 meshes [27], that Firedrake uses, to iterate over and query the mesh to construct
 491 this information. Due to the repeated low-level instructions required for this, we
 492 have implemented this in C as a normal PETSc preconditioner. Our implementation
 493 requires that the high-level “problem aware” preconditioner, in Python, initialise the
 494 patch preconditioner with the problem-specific data. This includes the function space
 495 description, identification of any Dirichlet nodes in the space, along with a callback
 496 to construct the patch operator. This callback is effectively the low-level code created
 497 when calling `assemble` on a UFL form. As is usual with PETSc objects, all aspects
 498 of the subsolves are configurable at runtime. Application of the patch inverses can
 499 either store and reuse matrices and factorizations (at the cost of high memory usage)
 500 or build, invert, and discard matrices patch-by-patch. This has much lower memory
 501 usage, but is computationally more expensive without access to either fast patch
 502 inverses or fast patch assembly routines.

503 **5. Examples and results.** We now present some examples and weak scaling
 504 results using Firedrake, and the new preconditioning framework we have developed.
 505 All results in this study were conducted on ARCHER, a Cray XC30 hosted at the
 506 University of Edinburgh. Each compute node contains two 2.7 GHz, 12-core E5-
 507 2697v2 (Ivy Bridge) processors, for a total of 24 cores per node, with a guaranteed
 508 not to exceed floating point performance of 518.4 Gflop/s. The spec sheet memory
 509 bandwidth is 119.4 GB/s per node, and we measured a STREAM triad [35] bandwidth
 510 of 74.1 GB/s when using 24 pinned MPI processes³. All experiments were performed
 511 with 24 MPI ranks per node (i.e. fully populated) with processes pinned to cores. For
 512 all experiments, we use regular simplicial meshes⁴ of the unit d -cube with piecewise
 513 linear coordinate fields.

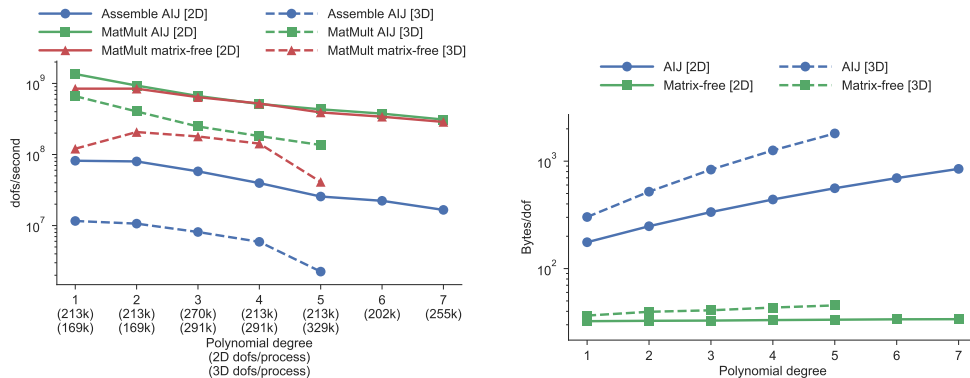
514 **5.1. Operator application.** Without access to fast, sum-factored algorithms,
 515 forming element tensors has complexity $\mathcal{O}(p^{3d})$ for Jacobian matrices, and $\mathcal{O}(p^{2d})$ for
 516 residual evaluation. Similarly, matrix-vector products for assembled sparse matrices
 517 require $\mathcal{O}(p^{2d})$ work, as do matrix-free applications (although the constants can be
 518 very different). Since Firedrake does not currently implement sum-factored algorithms

³The compiler did not generate non-temporal stores for this code.

⁴These meshes are nonetheless treated as unstructured by Firedrake.

519 on simplices, we expect that our matrix-free implementation to have the same time
 520 complexity as assembled sparse matrix-vector application. An advantage is that we
 521 have constant memory usage per degree of freedom (modulo surface-to-volume effects).

522 **Figure 1** shows performance of our implementation for a Poisson operator discretised
 523 with piecewise polynomial Lagrange basis functions. We see that we broadly
 524 observe the expected algorithmic behavior (barring in three dimensions, as explained
 525 in the figure). Assembled matrix-vector multiplication is faster than matrix-free applica-
 526 tion, although not by much for the two-dimensional case, at the cost of higher
 527 memory consumption per degree of freedom and the need to first assemble the matrix
 (costing approximately 10 matrix-free actions).



(a) Degrees of freedom per second processed for matrix assembly and matrix-vector products. The performance of matrix-free operator action and assembly at degree 5 in 3D becomes noticeably worse because the data for tabulated basis functions spills from the fastest cache.

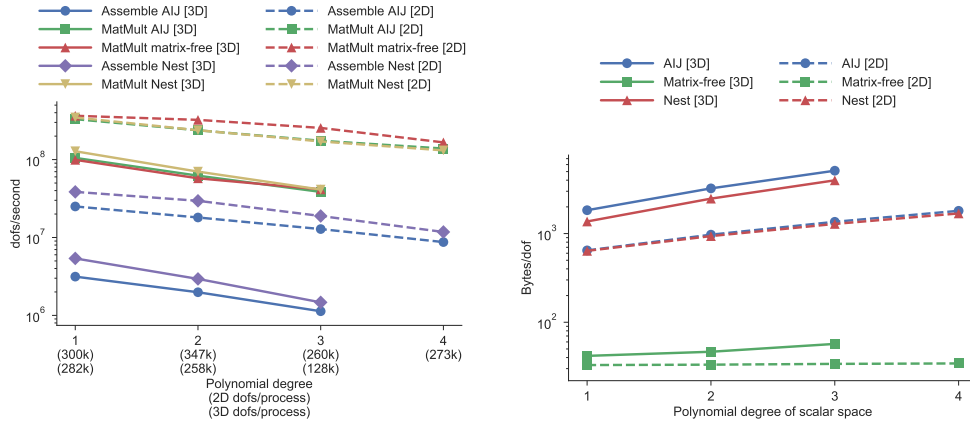
(b) Bytes of memory per degree of freedom. For the matrix-free case, memory usage is not quite constant, since Firedrake stores the ghosted representation, and so a surface-to-volume term appears in the memory per dof (more noticeable in three dimensions).

FIG. 1. Performance of matrix-vector products for a Poisson operator discretised on simplices in two and three dimensions (48 MPI processes).

528

529 The same story appears for more complex problems, and we show one example,
 530 the operator for Rayleigh-Bénard convection discretised using $P_{k+1}-P_k-P_k$ elements, in
 531 **Figure 2**. In two dimensions, the matrix-free action is faster than assembled operator
 532 application, and in three dimensions the cost is less than a factor 1.5 greater (even
 533 at lowest order). Given the high cost of matrix assembly, any iterative method that
 534 requires fewer than 10 matrix-vector products will be better off matrix-free, even
 535 before memory savings are considered.

536 To determine if these timings are good in absolute terms, we use a roofline model
 537 [43]. The arithmetic intensity for assembled matrix-vector products is calculated
 538 following [20]. For matrix assembly and matrix-free operator application, we count
 539 effective flops in the element kernel by traversing the intermediate representation of
 540 the generated code, the required data movement assumes a perfect cache model for
 541 any fields (each degree of freedom is only loaded for main memory once), and includes
 542 the cost of moving the indirection maps. The spec sheet memory bandwidth per node
 543 is 119.4 GB/s, and we measure a STREAM triad bandwidth of 74.1 GB/s per node;
 544 the guaranteed not to exceed floating point performance is 518.4 Gflop/s per node (one
 545 AVX multiplication and one AVX addition issued per cycle per core). As evidenced
 546 in **Figure 3**, there is almost no extra performance available for the application of

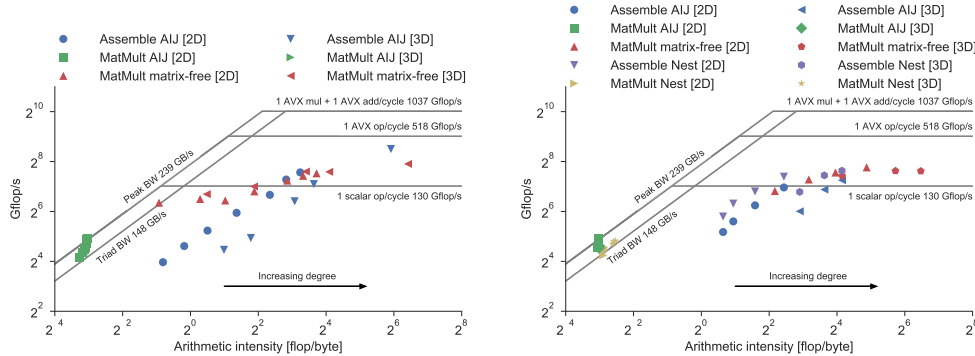


(a) Degrees of freedom per second processed for matrix assembly and matrix-vector products.

(b) Bytes of memory per degree of freedom.

FIG. 2. Performance of matrix-vector products for the Rayleigh-Bénard equation discretised on simplices in two and three dimensions (48 MPI processes).

547 assembled operators: the matrix-vector product achieves close to the machine peak
 548 in all cases. In contrast, the matrix-free actions, with significantly higher arithmetic
 549 intensity, are quite a distance from machine peak: this suggests a direction for future
 optimisation efforts in Firedrake.



(a) Performance of assembly and matrix-vector products for the Poisson operator. The assembled matrix achieves performance close to machine peak, while matrix-free products (and matrix assembly) are a way away.

(b) Performance of assembly and matrix-vector products for the Rayleigh-Bénard operator. The `nest` matrix has higher arithmetic intensity than the `aij` matrix due to using a blocked CSR format for the diagonal velocity block. As with the Poisson operator, assembled matrices achieve almost machine peak, whereas the matrix-free operator has room for improvement.

FIG. 3. Roofline plots for the experiments of Figure 1 and Figure 2.

550

551 **5.2. Runtime solver composition.**

552 **5.2.1. Poisson.** We now consider solving the Poisson problem (2) in three di-
 553 mensions. We choose as domain a regularly meshed unit cube, $\Omega = [0, 1]^d$, and apply

554 homogenous Dirichlet conditions on $\partial\Omega$, along with a constant forcing term. For
 555 low degree discretizations, “black-box” algebraic multigrid methods are robust and
 556 provide high performance. Their performance, however, degrades with increasing ap-
 557 proximation degree. Here we show how we can plug in the additive Schwarz approach
 558 described in [subsection 3.2](#) to provide a preconditioner with mesh and degree inde-
 559 pendent iteration counts. The preconditioner itself is not fully independent of these
 560 parameters because the time to solution increases: we are not using an $\mathcal{O}(n)$ coarse
 561 grid solve, instead using algebraic multigrid V cycles. **TODO: rework better to say**
 562 **only that iteration counts are constant, to be mesh and degree-independent in terms**
 563 **of time we need an $\mathcal{O}(n)$**

564 The main cost of this preconditioner is the application of the (dense) patch in-
 565 verses, the cost of our implementation is therefore quite high. We also comment that
 566 if patch operators are not stored between iterations, the overall memory footprint of
 567 the method is quite small. Developing fast algorithms to build and invert these patch
 568 operators is the subject of ongoing work.

569 In [Table 1](#) we compare the algorithmic and runtime performance of hypre’s
 570 boomerAMG algebraic multigrid solver applied directly to a P_4 discretization with
 571 the additive Schwarz approach. The only changes to the application file were in the
 572 specification of the runtime solver options. The provided solver options are shown
 573 in [Appendix B.1](#) for the hypre preconditioner and [Appendix B.2](#) for the Schwarz
 574 approach.

TABLE 1

Krylov iterations, and time to solution for P_4 Poisson problem using hypre and the Schwarz preconditioner described in [subsection 3.2](#) as the problem is weakly scaled. The required number of Krylov iterations grows slowly for the hypre preconditioner, but is constant for Schwarz. However, the overall time to solution is still lower with hypre.

DoFs ($\times 10^6$)	MPI processes	Krylov its		Time to solution (s)	
		hypre	schwarz	hypre	schwarz
2.571	24	19	19	5.62	9.48
5.545	48	20	19	6.45	10.6
10.22	96	20	19	6.17	10.3
20.35	192	21	18	6.53	10.7
43.99	384	22	19	7.53	11.9
81.18	768	22	19	7.52	11.7
161.9	1536	23	19	8.98	13
350.4	3072	24	19	8.56	14
647.2	6144	26	19	9.32	13.9
1291	12288	28	19	10.2	17.3
2797	24576	29	19	13	22.5

575 **5.2.2. FieldSplit examples.** Merely being able to solve the Poisson equation
 576 is a relatively uninteresting proposition. The power in our (and PETSc’s) approach
 577 is the ease of composition, *at runtime*, of scalable building blocks to provide pre-
 578 conditioners for complex problems. To demonstrate this, we consider solving the
 579 Rayleigh-Bénard equations for stationary convection (7).

580 A block preconditioner for this problem was developed in [22], but its perfor-
 581 mance was only studied in two-dimensional systems, and the implementation of the
 582 preconditioner was tightly coupled with the problem. The components of this pre-
 583 conditioner are: an inexact inverse of the Navier-Stokes equations, for which the

584 block preconditioners discussed in [18] provide mesh-independent iteration counts;
 585 an inexact inverse of the scalar (temperature) convection diffusion operator. For
 586 the Navier-Stokes block we approximate the Schur complement with the pressure-
 587 convection-diffusion approach (which requires information about the discretization
 588 inside the preconditioner). The building blocks are an approximate inverse for the ve-
 589 locity convection-diffusion operator, and approximate inverses for pressure mass and
 590 stiffness matrices. For moderate velocities, the velocity convection-diffusion operator
 591 can be treated with algebraic multigrid. Similarly, the pressure mass matrix can be
 592 inverted well with only a few iterations of a splitting-based method (e.g. point Ja-
 593 cobi), while multigrid is again good for the stiffness matrix. Finally, the temperature
 594 convection-diffusion operator can again be treated with algebraic multigrid.

595 Using the notation of (18) and (22), we need approximate inverses \tilde{N}^{-1} and \tilde{K}^{-1} .
 596 Where \tilde{N}^{-1} itself needs approximate inverses \tilde{F}^{-1} , K_p^{-1} , and M_p^{-1} . We can make
 597 different choices for all of these inverses, as well as to apply operators in a matrix-free
 598 manner or not, the matrix format for assembled operators, and convergence tolerance
 599 for all approximate inverses. These options (and others) can all be configured at
 600 runtime, while maintaining a single code base for the specification of the underlying
 601 PDE model, merely by modifying solver options.

602 Explicitly assembling the Jacobian and inverting with a direct solver requires a
 603 relatively short options list: [Appendix B.3](#). Conversely, to implement the precon-
 604 ditioner of (22), with algebraic multigrid for all approximate inverses (except the
 605 pressure mass matrix), and the operator applied matrix-free, we need significantly
 606 more options. These are shown in full in [Appendix B.4](#).

607 **5.2.3. Algorithmic and parallel scalability.** Firedrake and PETSc are de-
 608 signed such that the user of the library need not worry in detail about distributed
 609 memory parallelisation, provided they respect the collective semantics of operations.
 610 Since our implementation of solvers and preconditioners operates at the level of public
 611 APIs, we only need to be careful that we use the correct communicators when con-
 612 structing auxiliary objects. Parallelisation therefore comes “for free”. In this section,
 613 we show that our approach scales to large problem sizes, with scalability limited only
 614 by the performance of the building block components of the solver.

615 We consider the algorithmic performance of the Rayleigh-Bénard problem (7) in
 616 a regularly meshed unit cube, $\Omega = [0, 1]^3$. We choose as boundary conditions:

$$617 \quad (26a) \quad u = 0 \quad \text{on } \partial\Omega$$

$$618 \quad (26b) \quad \nabla p \cdot n = 0 \quad \text{on } \partial\Omega$$

$$619 \quad (26c) \quad T = 1 \quad \text{on the plane } x = 0$$

$$620 \quad (26d) \quad T = 0 \quad \text{on the plane } x = 1$$

$$621 \quad (26e) \quad \nabla T \cdot n = 0 \quad \text{otherwise}$$

623 and take $\text{Ra} = 200$ and $\text{Pr} = 6.18$. The constant pressure nullspace is projected out
 624 in the linear solver. The solution to this problem is shown in [Figure 4](#).

625 We perform a weak scaling experiment (increasing both the number of degrees of
 626 freedom, and computational resource) to study any mesh dependence in our solver.
 627 For the full set of solver options see [Appendix B.4](#). Newton iterations reduce the
 628 residual by 10^8 in three iterations, with only a weak increase in the number of Krylov
 629 iterations, as seen in [Table 2](#). The scalability does not look as good as these results
 630 would suggest, with only 20% parallel efficiency for this weakly scaled problem on
 631 6144 cores. Looking at the inner solves indicates the problem, although the outer




FIG. 4. Solution to the Rayleigh-Bénard problem of (7) with boundary conditions as specified in (26), and g pointing up. Shown are streamlines of the velocity field coloured by the magnitude of the velocity, and isosurfaces of the pressure.

TABLE 2

Newton iteration counts, total Krylov iterations, and time to solution for Rayleigh-Bénard convection as the problem is weakly scaled. The required number of linear iterations grows slowly as the mesh is refined, however the time to solution grows much faster.

DoFs ($\times 10^6$)	MPI processes	Newton its	Krylov its	Time to solution (s)
0.7405	24	3	16	31.7
1.488	48	3	16	36.3
2.973	96	3	17	43.9
5.769	192	3	17	47.3
11.66	384	3	17	56
23.39	768	3	17	64.9
45.54	1536	3	18	85.2
92.28	3072	3	18	120
185.6	6144	3	19	167

632 Krylov solve performs well, our approximate inner preconditioners are not fully mesh
633 independent. Table 3 shows the total number of iterations for both the Navier-Stokes
634 solve and the temperature solve as part of the application of the outer preconditioner.
635 In addition to iteration counts increasing, the time to compute a single iteration also
636 increases. This is observable more clearly in the previous results for the Poisson
637 operator (Table 1). This is due to sub-optimal scalability of the algebraic multigrid
638 that is used for all the building blocks in these solves. Our results for the Poisson

TABLE 3

Total iterations for Navier-Stokes and temperature solves (with average iterations per outer linear solve in brackets) for the nonlinear solution of the Rayleigh-Bénard problem. We see weak mesh dependence in the per-solve iteration counts. When multiplied up by the slight mesh dependence in the outer solve, this results in a noticeable inefficiency.

DoFs ($\times 10^6$)	Navier-Stokes iterations	Temperature iterations
0.7405	329 (20.6)	107 (6.7)
1.488	338 (21.1)	110 (6.9)
2.973	365 (21.5)	132 (7.8)
5.769	358 (21.1)	133 (7.8)
11.66	373 (21.9)	137 (8.1)
23.39	378 (22.2)	139 (8.2)
45.54	403 (22.4)	151 (8.4)
92.28	420 (23.3)	154 (8.6)
185.6	463 (24.4)	174 (9.2)

639 equation using hypre’s boomerAMG appear similar to previously reported results on
 640 weak scalability from the hypre team [3], and so we do not expect to gain much
 641 improvement here without changing the solver. This can, however, be done without
 642 modification to the existing solver: as soon as a better option is available, we can just
 643 drop it in.

644 **6. Conclusions and future outlook.** We have presented our approach in Fire-
 645 drake extending the existing solver interface to support matrix-free operators, and the
 646 necessary preconditioning infrastructure. Our approach is extensible and composable
 647 with existing algebraic solvers supported through PETSc. In particular, it removes
 648 much of the friction in developing block preconditioners requiring auxiliary opera-
 649 tors. The performance of such preconditioners for complex problems still relies on
 650 having good approximate inverses for the blocks, but our composable approach can
 651 seamlessly take advantage of any such advances.

652 **Appendix A. Code availability.** For reproducibility, we cite archives of the
 653 exact software versions that were used to produce the results in this paper. The
 654 experimentation and job submission framework (along with the plotting scripts and
 655 raw results) is available as [46]. The Additive Schwarz preconditioner from [subsec-](#)
 656 [tion 3.2](#) is [53]. For all components of the Firedrake project, we used recent versions:
 657 COFFEE [45], FIAT [47], FInAT [48], Firedrake [49], PETSc [50], petsc4py [51],
 658 PyOP2 [52], TSFC [54], and UFL [55].

659 **Appendix B. Full solver parameters.**

660 **B.1. Poisson: hypre.** We use hypre’s boomerAMG algebraic multigrid im-
 661 plementation, and select more aggressive coarsening strategies to obtain a lower-
 662 complexity coarse grid operator than the default.

```

663 -ksp_type cg -ksp_rtol 1e-8 -mat_type aij
664 -pc_type hypre -pc_hypre_type boomeramg
665 -pc_hypre_boomeramg_P_max 4
666 -pc_hypre_boomeramg_no_CF
667 -pc_hypre_boomeramg_agg_n1 1
668 -pc_hypre_boomeramg_agg_num_paths 2
669 -pc_hypre_boomeramg_coarsen_type HMIS
670 -pc_hypre_boomeramg_interp_type ext+i
671

```

673 **B.2. Poisson: schwarz.** We use exact inverses for the patch problems, and
 674 PETSc's GAMG algebraic multigrid for the P_1 inverse. The telescoping preconditioner [34]
 675 for the low-order P_1 operator is used to reduce the number of active MPI
 676 processes, since it has many fewer degrees of freedom than the P_4 operator.

```
677 -ksp_type cg -ksp_rtol 1e-8 -mat_type matfree
678 -pc_type python -pc_python_type ssc.SSC
679 -ssc_pc_composite_type additive
680 -ssc_sub_0_pc_patch_save_operators True
681 -ssc_sub_0_pc_patch_sub_mat_type seqaij
682 -ssc_sub_0_sub_ksp_type preonly
683 -ssc_sub_0_sub_pc_type lu
684 -ssc_sub_1_lo_pc_type telescope
685 -ssc_sub_1_lo_pc_telescope_reduction_factor 6
686 -ssc_sub_1_lo_telescope_ksp_max_it 4
687 -ssc_sub_1_lo_telescope_ksp_type richardson
688 -ssc_sub_1_lo_telescope_pc_type gamg
689
```

691 **B.3. Rayleigh-Bénard: direct.** To invert the full linearised Jacobian with a
 692 direct solver (here we use MUMPS [2]), we use the options:

```
693 -mat_type aij
694 -ksp_type preonly
695 -pc_type lu
696 -pc_factor_mat_solver_package mumps
697
```

699 **B.4. Rayleigh-Bénard: iterative.** To configure the nonlinear iteration, and
 700 then also split the Navier-Stokes block from the temperature block, we use:

```
701 -snes_type newtonls -snes_rtol 1e-8 -snes_linesearch_type basic
702 -ksp_type fgmr -ksp_gmres_modifiedgramschmidt
703 -mat_type matfree
704 -pc_type fieldsplit
705 -pc_fieldsplit_type multiplicative
706 -pc_fieldsplit_0_fields 0,1
707 -pc_fieldsplit_1_fields 2
708
```

710 now we configure the temperature solve to use GMRES and algebraic multigrid.

```
711 -prefix_push fieldsplit_1_
712 -ksp_type gmres
713 -ksp_rtol 1e-4,
714 -pc_type python
715 -pc_python_type firedrake.AssembledPC
716 -assembled_mat_type aij
717 -assembled_pc_type telescope
718 -assembled_pc_telescope_reduction_factor 6
719 -assembled_telescope_pc_type hypre
720 -assembled_telescope_pc_hypre_boomeramg_P_max 4
721 -assembled_telescope_pc_hypre_boomeramg_agg_n1 1
722 -assembled_telescope_pc_hypre_boomeramg_agg_num_paths 2
723 -assembled_telescope_pc_hypre_boomeramg_coarsen_type HMIS
724 -assembled_telescope_pc_hypre_boomeramg_interp_type ext+i
725 -assembled_telescope_pc_hypre_boomeramg_no_CF True
726 -prefix_pop
727
```

729 Finally we configure the Navier-Stokes solve to use GMRES with a lower Schur complement
 730 factorization as a preconditioner, and the pressure-convection-diffusion approximation
 731 for the schur complement.

```
732 -prefix_push fieldsplit_0_
733 -ksp_type gmres
734 -ksp_gmres_modifiedgramschmidt
735 -ksp_rtol 1e-2
736 -pc_type fieldsplit
737 -pc_fieldsplit_type schur
738 -pc_fieldsplit_schur_fact_type lower
739
```

```

740 -prefix_push fieldsplit_0_
741 -ksp_type preonly
742 -pc_type python
743 -pc_python_type firedrake.AssembledPC
744 -assembled_mat_type aij
745 -assembled_pc_type hypre
746 -assembled_pc_hypre_boomeramg_P_max 4
747 -assembled_pc_hypre_boomeramg_agg_nl 1
748 -assembled_pc_hypre_boomeramg_agg_num_paths 2
749 -assembled_pc_hypre_boomeramg_coarsen_type HMIS
750 -assembled_pc_hypre_boomeramg_interp_type ext+i
751 -assembled_pc_hypre_boomeramg_no_CF
752 -prefix_pop
753
754
755 -prefix_push fieldsplit_1_
756 -ksp_type preonly
757 -pc_type python
758 -pc_python_type firedrake.PCDPC
759 -pcd_Fp_mat_type matfree
760 -pcd_Kp_ksp_type preonly
761 -pcd_Kp_mat_type aij
762 -pcd_Kp_pc_type telescope
763 -pcd_Kp_pc_telescope_reduction_factor 6
764 -pcd_Kp_telescope_pc_type ksp
765 -pcd_Kp_telescope_ksp_ksp_max_it 3
766 -pcd_Kp_telescope_ksp_ksp_type richardson
767 -pcd_Kp_telescope_ksp_pc_type hypre
768 -pcd_Kp_telescope_ksp_pc_hypre_boomeramg_P_max 4
769 -pcd_Kp_telescope_ksp_pc_hypre_boomeramg_agg_nl 1
770 -pcd_Kp_telescope_ksp_pc_hypre_boomeramg_agg_num_paths 2
771 -pcd_Kp_telescope_ksp_pc_hypre_boomeramg_coarsen_type HMIS
772 -pcd_Kp_telescope_ksp_pc_hypre_boomeramg_interp_type ext+i
773 -pcd_Kp_telescope_ksp_pc_hypre_boomeramg_no_CF
774
775 -pcd_Mp_mat_type aij
776 -pcd_Mp_ksp_type richardson
777 -pcd_Mp_pc_type sor
778 -pcd_Mp_ksp_max_it 2
779 -prefix_pop
780 -prefix_pop

```

782

REFERENCES

- 783 [1] M. S. ALNÆS, A. LOGG, K. B. ØLGAARD, M. E. ROGNES, AND G. N. WELLS, *Unified Form Language: a domain-specific language for weak formulations of partial differential equations*,
784 ACM Transactions on Mathematical Software, 40 (2014), <https://doi.org/10.1145/2566630>,
785 <https://arxiv.org/abs/1211.4047>.
786 [2] P. R. AMESTOY, I. S. DUFF, AND J.-Y. L'EXCELLENT, *Multifrontal parallel distributed symmetric and unsymmetric solvers*, Computer methods in applied mechanics and engineering,
787 184 (2000), pp. 501–520, [https://doi.org/10.1016/S0045-7825\(99\)00242-X](https://doi.org/10.1016/S0045-7825(99)00242-X).
788 [3] A. H. BAKER, R. D. FALGOUT, T. GAMBLIN, T. V. KOLEV, M. SCHULZ, AND U. M. YANG, *Scaling algebraic multigrid solvers: on the road to exascale*, in Competence in High Performance Computing 2010: Proceedings of an International Conference on Competence
789 in High Performance Computing, June 2010, Schloss Schwetzingen, Germany, C. Bischof,
790 H.-G. Hegering, W. E. Nagel, and G. Wittum, eds., Berlin, Heidelberg, 2012, Springer
791 Berlin Heidelberg, pp. 215–226, https://doi.org/10.1007/978-3-642-24025-6_18.
792 [4] S. BALAY, S. ABHYANKAR, M. F. ADAMS, J. BROWN, P. BRUNE, K. BUSCHELMAN, L. DALCIN,
793 V. ELJKHOUT, W. D. GROPP, D. KAUSHIK, M. G. KNEPLEY, L. C. MCINNES, K. RUPP,
794 B. F. SMITH, S. ZAMPINI, H. ZHANG, AND H. ZHANG, *PETSc Users Manual*, Tech. Report
795 ANL-95/11 - Revision 3.7, Argonne National Laboratory, 2016.
796 [5] S. BALAY, W. D. GROPP, L. C. MCINNES, AND B. F. SMITH, *Efficient management of parallelism in object oriented numerical software libraries*, in Modern Software Tools in Scientific
797 Computing, E. Arge, A. M. Bruaset, and H. P. Langtangen, eds., Birkhäuser Press, 1997,
798 pp. 163–202, https://doi.org/10.1007/978-1-4612-1986-6_8.
799 [6] M. BENZI, G. GOLUB, AND J. LIESEN, *Numerical solution of saddle point problems*, Acta
800 Numerica, 14 (2005), pp. 1–137, <https://doi.org/10.1017/S0962492904000212>.
801
802
803
804
805

- 806 [7] A. BRANDT, *Multi-level adaptive solutions to boundary-value problems*, Mathematics of Computation, 31 (1977), pp. 333–390, <https://doi.org/10.1090/S0025-5718-1977-0431719-X>.
- 807
- 808 [8] A. BRANDT AND O. LIVNE, *Multigrid Techniques*, Society for Industrial and Applied Mathematics, 2011, <https://doi.org/10.1137/1.9781611970753>.
- 809
- 810 [9] S. C. BRENNER AND L. R. SCOTT, *The mathematical theory of finite element methods*, vol. 15 of Texts in Applied Mathematics, Springer, New York, third edition ed., 2008.
- 811
- 812 [10] J. BROWN, M. G. KNEPLEY, D. A. MAY, L. C. MCINNES, AND B. F. SMITH, *Composable linear solvers for multiphysics*, in Proceedings of the 2012 11th International Symposium on Parallel and Distributed Computing, ISPD '12, Washington, DC, USA, 2012, IEEE Computer Society, pp. 55–62, <https://doi.org/10.1109/ISPD.2012.16>.
- 813
- 814
- 815 [11] G. F. CAREY AND T. J. ODEN, *Finite Elements: Fluid Mechanics Vol., VI*, Prentice-Hall, Inc., 1986.
- 816
- 817
- 818 [12] E. C. CYR, J. N. SHADID, AND R. S. TUMINARO, *Teko: A block preconditioning capability with concrete example applications in Navier-Stokes and MHD*, SIAM Journal on Scientific Computing, 38 (2016), pp. S307–S331, <https://doi.org/10.1137/15M1017946>.
- 819
- 820
- 821 [13] L. D. DALCIN, R. R. PAZ, P. A. KLER, AND A. COSIMO, *Parallel distributed computing using Python*, Advances in Water Resources, 34 (2011), pp. 1124–1139, <https://doi.org/10.1016/j.advwatres.2011.04.013>. New Computational Methods and Software Tools.
- 822
- 823
- 824 [14] T. A. DAVIS, *Algorithm 832: UMFPACK V4.3—an unsymmetric-pattern multifrontal method*, ACM Transactions on Mathematical Software, 30 (2004), pp. 196–199, <https://doi.org/10.1145/992200.992206>.
- 825
- 826
- 827 [15] H. ELMAN, V. E. HOWLE, J. SHADID, R. SHUTTLEWORTH, AND R. TUMINARO, *Block preconditioners based on approximate commutators*, SIAM Journal on Scientific Computing, 27 (2006), pp. 1651–1668, <https://doi.org/10.1137/040608817>.
- 828
- 829
- 830 [16] H. ELMAN, V. E. HOWLE, J. SHADID, R. SHUTTLEWORTH, AND R. TUMINARO, *A taxonomy and comparison of parallel block multi-level preconditioners for the incompressible Navier–Stokes equations*, Journal of Computational Physics, 227 (2008), pp. 1790–1808, <https://doi.org/10.1016/j.jcp.2007.09.026>.
- 831
- 832
- 833
- 834 [17] H. ELMAN AND D. SILVESTER, *Fast nonsymmetric iterations and preconditioning for Navier-Stokes equations*, SIAM Journal on Scientific Computing, 17 (1996), pp. 33–46, <https://doi.org/10.1137/0917004>.
- 835
- 836
- 837 [18] H. ELMAN, D. SILVESTER, AND A. WATHEN, *Finite elements and fast iterative solvers*, Oxford University Press, second edition ed., 2014.
- 838
- 839 [19] P. E. FARRELL AND J. W. PEARSON, *A preconditioner for the Ohta-Kawasaki equation*, 2016, <https://arxiv.org/abs/1603.04570>.
- 840
- 841 [20] W. D. GROPP, D. K. KAUSHIK, D. E. KEYES, AND B. F. SMITH, *Towards realistic performance bounds for implicit CFD codes*, in Parallel CFD 1999, D. Keyes, A. Ecer, J. Peiriaux, and N. Satofuka, eds., North-Holland, 2000, pp. 241–248, <https://doi.org/10.1016/B978-044482851-4.50030-X>.
- 842
- 843
- 844
- 845 [21] M. A. HEROUX, R. A. BARTLETT, V. E. HOWLE, R. J. HOEKSTRA, J. J. HU, T. G. KOLDA, R. B. LEHOUCQ, K. R. LONG, R. P. PAWLOWSKI, E. T. PHIPPS, A. G. SALINGER, H. K. THORNQUIST, R. S. TUMINARO, J. M. WILLENBRING, A. WILLIAMS, AND K. S. STANLEY, *An overview of the Trilinos project*, ACM Transactions on Mathematical Software, 31 (2005), pp. 397–423, <https://doi.org/10.1145/1089014.1089021>.
- 846
- 847
- 848
- 849
- 850 [22] V. E. HOWLE AND R. C. KIRBY, *Block preconditioners for finite element discretization of incompressible flow with thermal convection*, Numerical Linear Algebra with Applications, 19 (2012), pp. 427–440, <https://doi.org/10.1002/nla.1814>.
- 851
- 852
- 853 [23] V. E. HOWLE, R. C. KIRBY, K. LONG, B. BRENNAN, AND K. KENNEDY, *Playa: high-performance programmable linear algebra*, Scientific Programming, 20 (2012), pp. 257–273, <https://doi.org/10.1155/2012/606215>.
- 854
- 855
- 856 [24] B. JESSIC, D. KAY, M. STOLL, AND A. J. WATHEN, *Fast solvers for Cahn-Hilliard inpainting*, SIAM Journal on Imaging Sciences, 7 (2014), pp. 67–97, <https://doi.org/10.1137/130921842>.
- 857
- 858
- 859 [25] D. KAY, D. LOGHIN, AND A. WATHEN, *A preconditioner for the steady-state Navier-Stokes equations*, SIAM Journal on Scientific Computing, 24 (2002), pp. 237–256, <https://doi.org/10.1137/S106482759935808X>.
- 860
- 861
- 862 [26] R. C. KIRBY, *Algorithm 839: FIAT, a new paradigm for computing finite element basis functions*, ACM Transactions on Mathematical Software, 30 (2004), pp. 502–516, <https://doi.org/10.1145/1039813.1039820>.
- 863
- 864
- 865 [27] M. G. KNEPLEY AND D. A. KARPEEV, *Mesh Algorithms for PDE with Sieve I: Mesh Distribution*, Scientific Programming, 17 (2009), pp. 215–230, <https://doi.org/10.1155/2009/948613>.
- 866
- 867

- 868 [28] A. LOGG, K.-A. MARDAL, AND G. N. WELLS, eds., *Automated solution of differential equations*
869 *by the finite element method: the FEniCS book*, vol. 84, Springer, 2012, [https://doi.org/](https://doi.org/10.1007/978-3-642-23099-8)
870 [10.1007/978-3-642-23099-8](https://doi.org/10.1007/978-3-642-23099-8).
- 871 [29] K. LONG, R. C. KIRBY, AND B. VAN BLOEMEN WAANDERS, *Unified embedded parallel finite ele-*
872 *ment computations via software-based Fréchet differentiation*, SIAM Journal on Scientific
873 Computing, 32 (2010), pp. 3323–3351, <https://doi.org/10.1137/09076920X>.
- 874 [30] K. R. LONG, *Sundance rapid prototyping tool for parallel PDE optimization*, in Large-Scale
875 PDE-Constrained Optimization, L. T. Biegler, M. Heinkenschloss, O. Ghattas, and B. van
876 Bloemen Waanders, eds., Springer Berlin Heidelberg, Berlin, Heidelberg, 2003, pp. 331–
877 341, https://doi.org/10.1007/978-3-642-55508-4_20.
- 878 [31] K.-A. MARDAL AND J. B. HAGA, *Automated solution of differential equations by the finite*
879 *element method: the FEniCS book*, in Logg et al. [28], ch. Block preconditioning of systems
880 of PDEs, <https://doi.org/10.1007/978-3-642-23099-8>.
- 881 [32] K.-A. MARDAL AND R. WINTHER, *Preconditioning discretizations of systems of partial dif-*
882 *ferential equations*, Numerical Linear Algebra with Applications, 18 (2011), pp. 1–40,
883 <https://doi.org/10.1002/nla.716>.
- 884 [33] D. A. MAY, J. BROWN, AND L. LE POURHIET, *pTatin3D: High-performance methods for long-*
885 *term lithospheric dynamics*, in Proceedings of the International Conference for High Per-
886 formance Computing, Networking, Storage and Analysis, SC '14, Piscataway, NJ, USA,
887 2014, IEEE Press, pp. 274–284, <https://doi.org/10.1109/SC.2014.28>.
- 888 [34] D. A. MAY, P. SANAN, K. RUPP, M. G. KNEPLEY, AND B. F. SMITH, *Extreme-scale multi-*
889 *grid components within PETSc*, in Proceedings of the Platform for Advanced Scientific
890 Computing Conference, PASC '16, New York, NY, USA, 2016, ACM, pp. 5:1–5:12,
891 <https://doi.org/10.1145/2929908.2929913>, <https://arxiv.org/abs/1604.07163>.
- 892 [35] J. D. MCCALPIN, *Memory Bandwidth and Machine Balance in Current High Performance*
893 *Computers*, IEEE Computer Society Technical Committee on Computer Architecture
894 Newsletter, (1995), pp. 19–25.
- 895 [36] M. F. MURPHY, G. H. GOLUB, AND A. J. WATHEN, *A note on preconditioning for indefinite*
896 *linear systems*, SIAM Journal on Scientific Computing, 21 (2000), pp. 1969–1972, <https://doi.org/10.1137/S1064827599355153>.
- 897 [37] L. F. PAVARINO, *Additive Schwarz methods for the p-version finite element method*, Numerische
898 Mathematik, 66 (1993), pp. 493–515, <https://doi.org/10.1007/BF01385709>.
- 899 [38] L. F. PAVARINO, *Schwarz methods with local refinement for the p-version finite element method*,
900 Numerische Mathematik, 69 (1994), pp. 185–211, <https://doi.org/10.1007/s002110050087>.
- 901 [39] L. F. PAVARINO AND T. WARBURTON, *Overlapping Schwarz methods for unstructured spectral*
902 *elements*, Journal of Computational Physics, 160 (2000), pp. 298–317, [https://doi.org/10.](https://doi.org/10.1006/jcph.2000.6463)
903 [1006/jcph.2000.6463](https://doi.org/10.1006/jcph.2000.6463).
- 904 [40] F. RATHGEBER, D. A. HAM, L. MITCHELL, M. LANGE, F. LUPORINI, A. T. T. MCRAE, G.-T.
905 BERCEA, G. R. MARKALL, AND P. H. J. KELLY, *Firedrake: automating the finite ele-*
906 *ment method by composing abstractions*, ACM Transactions on Mathematical Software,
907 43 (2016), pp. 24:1–24:27, <https://doi.org/10.1145/2998441>, [https://arxiv.org/abs/1501.](https://arxiv.org/abs/1501.01809)
908 [01809](https://arxiv.org/abs/1501.01809).
- 909 [41] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, Society for Industrial and Applied
910 Mathematics, second edition ed., 2003, <https://doi.org/10.1137/1.9780898718003>.
- 911 [42] J. SCHÖBERL, J. M. MELENK, C. PECHSTEIN, AND S. ZAGLMAYR, *Additive Schwarz precondition-*
912 *ing for p-version triangular and tetrahedral finite elements*, IMA Journal of Numerical
913 Analysis, 28 (2008), pp. 1–24, <https://doi.org/10.1093/imanum/drl046>.
- 914 [43] S. WILLIAMS, A. WATERMAN, AND D. PATTERSON, *Roofline: An insightful visual performance*
915 *model for multicore architectures*, Communications of the ACM, 52 (2009), pp. 65–76,
916 <https://doi.org/10.1145/1498765.1498785>.
- 917 [44] J. XU, *Iterative methods by space decomposition and subspace correction*, SIAM review, 34
918 (1992), pp. 581–613, <https://doi.org/10.1137/1034116>.
- 919 [45] ZENODO/COFFEE, *COFFEE: A Compiler for Fast Expression Evaluation*, Dec. 2016, <https://doi.org/10.5281/zenodo.208989>.
- 920 [46] ZENODO/COMPOSABLE-SOLVERS, *composable-solvers: experimentation framework for compos-*
921 *able block preconditioners*, June 2017, <https://doi.org/10.5281/zenodo.802277>.
- 922 [47] ZENODO/FIAT, *FIAT: The Finite Element Automated Tabulator*, June 2017, [https://doi.org/](https://doi.org/10.5281/zenodo.802269)
923 [10.5281/zenodo.802269](https://doi.org/10.5281/zenodo.802269).
- 924 [48] ZENODO/FINAT, *FINAT: a smarter library of finite elements*, June 2017, [https://doi.org/10.](https://doi.org/10.5281/zenodo.802275)
925 [5281/zenodo.802275](https://doi.org/10.5281/zenodo.802275).
- 926 [49] ZENODO/FIREDRAKE, *Firedrake: an automated finite element system*, June 2017, [https://doi.](https://doi.org/10.5281/zenodo.802271)
927 [org/10.5281/zenodo.802271](https://doi.org/10.5281/zenodo.802271).

- 930 [50] ZENODO/PETSC, *PETSc: Portable, Extensible Toolkit for Scientific Computation*, June 2017,
931 <https://doi.org/10.5281/zenodo.802273>.
- 932 [51] ZENODO/PETSC4PY, *petsc4py: The Python interface to PETSc*, June 2017, [https://doi.org/10.](https://doi.org/10.5281/zenodo.802274)
933 [5281/zenodo.802274](https://doi.org/10.5281/zenodo.802274).
- 934 [52] ZENODO/PYOP2, *PyOP2: Framework for performance-portable parallel computations on un-*
935 *structured meshes*, June 2017, <https://doi.org/10.5281/zenodo.802272>.
- 936 [53] ZENODO/SSC, *SSC: subspace corrections in Firedrake & PETSc*, June 2017, [https://doi.org/](https://doi.org/10.5281/zenodo.802279)
937 [10.5281/zenodo.802279](https://doi.org/10.5281/zenodo.802279).
- 938 [54] ZENODO/TSFC, *TSFC: The Two Stage Form Compiler*, June 2017, [https://doi.org/10.5281/](https://doi.org/10.5281/zenodo.802268)
939 [zenodo.802268](https://doi.org/10.5281/zenodo.802268).
- 940 [55] ZENODO/UFL, *UFL: The Unified Form Language*, June 2017, [https://doi.org/10.5281/zenodo.](https://doi.org/10.5281/zenodo.802270)
941 [802270](https://doi.org/10.5281/zenodo.802270).