

Multiphysics Simulations: Challenges and Opportunities

Mathematics and Computer Science Division

About Argonne National Laboratory

Argonne is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC under contract DE-AC02-06CH11357. The Laboratory's main facility is outside Chicago, at 9700 South Cass Avenue, Argonne, Illinois 60439. For information about Argonne and its pioneering science and technology programs, see www.anl.gov.

Availability of This Report

This report is available, at no cost, at <http://www.osti.gov/bridge>. It is also available on paper to the U.S. Department of Energy and its contractors, for a processing fee, from:

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
phone (865) 576-8401
fax (865) 576-5728
reports@adonis.osti.gov

Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor UChicago Argonne, LLC, nor any of their employees or officers, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of document authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, Argonne National Laboratory, or UChicago Argonne, LLC.

Multiphysics Simulations: Challenges and Opportunities¹

Report from a Workshop Sponsored by the Institute for Computing in Science (ICiS)
July 30 - August 6, 2011
Park City, Utah

Workshop Co-Organizers:

David Keyes, KAUST and Columbia University
Lois Curfman McInnes, Argonne National Laboratory
Carol Woodward, Lawrence Livermore National Laboratory

Theme: Practices and Perils of Multiphysics Applications

Principal Lead: **Eric Myra**, University of Michigan
Co-Lead: **Charbel Farhat**, Stanford University

Theme: Multiphysics Algorithms

Principal Lead: **Michael Pernice**, Idaho National Laboratory
Co-Lead: **Don Estep**, Colorado State University

Theme: Multiphysics Software Design

Principal Lead: **William Gropp**, University of Illinois at Urbana-Champaign
Co-Lead: **John Magerlein**, IBM Research Center

Additional Contributors:

| | |
|---|--|
| John Bell, Lawrence Berkeley National Laboratory | Qiming Lu, Fermi National Accelerator Laboratory |
| Jed Brown, Argonne National Laboratory | Reed Maxwell, Colorado School of Mines |
| Alain Clo, KAUST | Michael McCourt, Cornell University |
| Jeffrey Connors, Lawrence Livermore National Laboratory | Miriam Mehl, Technische Universität München |
| Emil Constantinescu, Argonne National Laboratory | Roger Pawlowski, Sandia National Laboratories |
| Kate Evans, Oak Ridge National Laboratory | Amanda Peters, Harvard University |
| Ammar Hakim, Tech-X Corporation | Daniel Reynolds, Southern Methodist University |
| Glenn Hammond, Pacific Northwest National Laboratory | Beatrice Riviere, Rice University |
| Glen Hansen, Sandia National Laboratories | Ulrich Rüde, University Erlangen-Nuremberg |
| Judith Hill, Oak Ridge National Laboratory | Tim Scheibe, Pacific Northwest National Laboratory |
| Tobin Isaac, University of Texas at Austin | John Shadid, Sandia National Laboratories |
| Xiangmin Jiao, Stonybrook University | Brendan Sheehan, Colorado State University |
| Kirk Jordan, IBM Research Center | Mark Shephard, Rensselaer Polytechnic Institute |
| Dinesh Kaushik, KAUST | Andrew Siegel, Argonne National Laboratory |
| Efthimios Kaxiras, Harvard University | Barry Smith, Argonne National Laboratory |
| Alice Koniges, Lawrence Berkeley National Laboratory | Xianzhu Tang, Los Alamos National Laboratory |
| Kihwan Lee, SLAC National Accelerator Laboratory | Cian Wilson, Columbia University |
| Aaron Lott, Lawrence Livermore National Laboratory | Barbara Wohlmuth, Technische Universität München |

¹Please cite this document as follows: *Multiphysics Simulations: Challenges and Opportunities*, D. E. Keyes, L. C. McInnes, C. Woodward, W. D. Gropp, E. Myra, M. Pernice, J. Bell, J. Brown, A. Clo, J. Connors, E. Constantinescu, D. Estep, K. Evans, C. Farhat, A. Hakim, G. Hammond, G. Hansen, J. Hill, T. Isaac, X. Jiao, K. Jordan, D. Kaushik, E. Kaxiras, A. Koniges, K. Lee, A. Lott, Q. Lu, J. Magerlein, R. Maxwell, M. McCourt, M. Mehl, R. Pawlowski, A. Peters, D. Reynolds, B. Riviere, U. Rüde, T. Scheibe, J. Shadid, B. Sheehan, M. Shephard, A. Siegel, B. Smith, X. Tang, C. Wilson, and B. Wohlmuth, Tech. Rep. ANL/MCS-TM-321, Argonne National Laboratory, Report of workshop sponsored by the Institute for Computing in Science (ICiS), July 30 - Aug. 6, 2011, Park City, Utah (Dec. 2011).

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 1 |
| 1.1 | What Constitutes Multiphysics? | 2 |
| 1.2 | Prototype Algebraic Forms | 3 |
| 1.3 | Structure and Motivation for This Review | 6 |
| 2 | Practices and Perils in Multiphysics Applications | 8 |
| 2.1 | Examples of PDE-Based Multiphysics Applications | 8 |
| 2.1.1 | Interaction of fluids and structures | 8 |
| 2.1.2 | Fission reactor fuel performance | 10 |
| 2.1.3 | Conjugate heat transfer/neutron transport coupling in reactor cores | 12 |
| 2.1.4 | Multiscale methods in hard and soft materials: crack propagation and ultrafast DNA sequencing | 13 |
| 2.1.5 | Simulation of processes in magnetic confinement fusion | 15 |
| 2.1.6 | Subsurface science | 17 |
| 2.1.7 | Surface and subsurface hydrology | 18 |
| 2.1.8 | Climate modeling | 19 |
| 2.1.9 | Radiation hydrodynamics | 21 |
| 2.1.10 | Geodynamics and magma dynamics | 22 |
| 2.1.11 | Particle-accelerator design | 23 |
| 2.2 | Crosscutting Issues in Multiphysics Applications | 24 |
| 2.2.1 | Choices and challenges in coupling algorithms | 25 |
| 2.2.2 | Software engineering | 26 |
| 2.2.3 | Analysis and verification | 26 |
| 3 | Algorithms for Multiphysics Coupling | 27 |
| 3.1 | Solver Methods | 27 |
| 3.1.1 | Methods for systems of linear equations | 27 |
| 3.1.2 | Methods for systems of nonlinear equations | 30 |
| 3.2 | Continuum-Continuum Coupling | 32 |
| 3.2.1 | Methods for coupling multiphysics components in space | 33 |
| 3.2.2 | Methods for coupling multiphysics components in time | 35 |
| 3.3 | Continuum-Discrete Coupling | 39 |
| 3.4 | Error Estimation | 40 |
| 3.4.1 | Operator splitting for reaction-diffusion equations | 42 |
| 3.4.2 | Iterative solution of parabolic problems coupled through a common boundary | 43 |
| 3.4.3 | Solution of systems of elliptic equations with independent discretizations | 44 |
| 3.5 | Uncertainty Quantification | 46 |
| 4 | Multiphysics Software | 47 |
| 4.1 | Status of Software for Multiphysics | 47 |
| 4.1.1 | Current practices | 47 |
| 4.1.2 | Common needs | 48 |
| 4.1.3 | Software successes | 49 |
| 4.2 | Challenges in Multiphysics Software Design | 52 |
| 4.2.1 | Enabling introduction of new models, algorithms, and data structures | 52 |
| 4.2.2 | Sharing methods and codes among application fields | 54 |
| 4.2.3 | Relating multiphysics spatial discretizations | 55 |
| 4.2.4 | Timestep control | 57 |
| 4.2.5 | Achieving performance | 57 |
| 4.2.6 | Software engineering issues for multiphysics integration | 57 |
| 4.3 | Difficulties in Collaborative Multiphysics Software | 58 |

| | | |
|--------------|---|-----------|
| 4.3.1 | Maintenance of a stable and efficient development environment | 58 |
| 4.3.2 | Investment in software tools and techniques | 59 |
| 4.3.3 | Adoption of third-party software | 60 |
| 4.3.4 | Licensing | 60 |
| 5 | Opportunities in Multiphysics Simulation Research | 61 |
| 5.1 | Exascale Issues | 61 |
| 5.2 | Concepts Ripe for Leveraging | 61 |
| 5.2.1 | Architecturally optimized, distributed, hierarchical data structures from solver packages | 62 |
| 5.2.2 | High-accuracy interpolation | 62 |
| 5.2.3 | Multiphysics as a form of adaptation | 62 |
| 5.2.4 | Applications of classical numerical analysis | 62 |
| 5.2.5 | Exploitation of additional concurrency | 63 |
| 5.2.6 | Exploitation of “tricks” for implicit algorithms | 63 |
| 5.2.7 | Provision of special hardware | 63 |
| 5.2.8 | Software maintenance best practices | 63 |
| 5.3 | Required Breakthroughs | 64 |
| 5.3.1 | Formulation of objective functions useful for optimizing simulations | 64 |
| 5.3.2 | Characterization of uncertainty | 64 |
| 5.3.3 | Equidistribution and control of error | 64 |
| 5.3.4 | Dynamic characterization of cross-physics interaction strength | 64 |
| 5.3.5 | Reduction of data transfers | 65 |
| 5.3.6 | Relaxation of synchrony | 65 |
| 5.3.7 | User-specified checkpointing | 65 |
| 5.3.8 | Combining continuum-based and discrete phenomena | 65 |
| 5.3.9 | Multiphysics software design | 66 |
| 5.3.10 | Multiphysics performance models | 66 |
| 6 | Insertion Paths for Algorithms and Software into Multiphysics Applications | 67 |
| 6.1 | Mechanisms | 67 |
| 6.1.1 | Leveraging of interfaces | 67 |
| 6.1.2 | Small, side effort | 67 |
| 6.1.3 | Full rewrites | 67 |
| 6.2 | Factors Affecting Successful Insertion Paths | 68 |
| 6.2.1 | Clear motivation | 68 |
| 6.2.2 | Collaboration | 68 |
| 6.2.3 | Discretization and meshes | 69 |
| 6.2.4 | Software packages | 69 |
| 6.2.5 | Funding | 69 |
| 6.2.6 | Personnel | 69 |
| 6.3 | Observations | 70 |
| 7 | Conclusions | 71 |
| | Acknowledgments | 72 |
| | Appendix A Glossary | 73 |
| | Appendix B Multiphysics Exemplars and Benchmarks | 75 |
| Appendix B.1 | Advection-reaction | 76 |
| Appendix B.2 | Heat transfer coupled across an interface | 78 |
| Appendix B.3 | Conjugate heat transfer | 80 |
| | References | 96 |

Abstract

This report is an outcome of the workshop *Multiphysics Simulations: Challenges and Opportunities*, sponsored by the Institute of Computing in Science (ICiS). Additional information about the workshop, including relevant reading and presentations on multiphysics issues in applications, algorithms, and software, is available via <https://sites.google.com/site/icismultiphysics2011/>.

We consider multiphysics applications from algorithmic and architectural perspectives, where “algorithmic” includes both mathematical analysis and computational complexity and “architectural” includes both software and hardware environments. Many diverse multiphysics applications can be reduced, en route to their computational simulation, to a common algebraic coupling paradigm. Mathematical analysis of multiphysics coupling in this form is not always practical for realistic applications, but model problems representative of applications discussed herein can provide insight. A variety of software frameworks for multiphysics applications have been constructed and refined within disciplinary communities and executed on leading-edge computer systems. We examine several of these, expose some commonalities among them, and attempt to extrapolate best practices to future systems. From our study, we summarize challenges and forecast opportunities. We also initiate a modest suite of test problems encompassing features present in many applications.

Keywords: multiphysics, multimodel, multirate, multiscale, implicit and explicit algorithms, strong and weak coupling, loose and tight coupling

1. Introduction

Simulations that couple multiple physical phenomena are as old as simulations themselves. However, multiphysics simulation deserves fresh assessment, in light of steadily increasing computational capability and greater aspirations for simulation in domains of scientific prediction, engineering design, and policy making. An oft-quoted motivation for extreme computing is to relax assumptions of decoupling; however, it is by no means obvious that the promises claimed for coupled multiphysics simulation will be realized in extreme-scale computational environments in the principal way by which individual codes are coupled today, namely, through divide-and-conquer operator splitting. Coupling individual simulations may introduce limitations of stability, accuracy, or robustness that are more severe than the limitations imposed by the individual components. Furthermore, the data motion and data structure conversions required to iterate between independent simulations for each component may be more costly in latency and electrical power than those of the individually tuned components. Thus, “one plus one” may cost significantly more than “two” and may be less amenable to scalable execution than expected.

The 2008 report by D. Brown et al. on *Applied Mathematics at the U.S. Department of Energy: Past, Present, and a View to the Future* [1] emphasizes these challenges:

Today’s problems, unlike traditional science and engineering, do not involve physical processes covered by a single traditional discipline of physics or the associated mathematics. Complex systems encountered in virtually all applications of interest to DOE involve many distinct physical processes ...

... The issue of coupling models of different events at different scales and governed by different physical laws is largely wide open and represents an enormously challenging area for future research.

In this review, we consider multiphysics applications from algorithmic and architectural perspectives, where “algorithmic” includes both mathematical analysis and computational complexity and “architectural” includes both software and hardware environments. Many multiphysics applications can be reduced, in principle, to an algebraic paradigm whose linearization brings to bear powerful tools of analysis in which individual components (denoted here as “uniphysics”) are represented by diagonal blocks and the multiphysics coupling between them as off-diagonal blocks. Such analyses are not always practical for realistic applications, but model problems representative of applications discussed herein provide insight. Various software frameworks for multiphysics applications have been constructed and refined within disciplinary communities and executed on leading-edge computer systems. We discuss here some commonalities among these frameworks and extrapolate best practices to future systems.

We begin by presenting a spectrum ranging from applications that all would agree are “multiphysics” to others that may involve semantic stretching but that have similar structure when the discrete problem is executed on a digital computer. Our objectives are broadly to characterize the state of the art in multiphysics simulation, to illustrate opportunities for leveraging successful approaches and software across physically disparate applications, and to identify gaps in understanding that are retarding the overarching goal of uncertainty-quantified high-fidelity simulations.

Whereas multiphysics is often approached bottom-up as the assembly of individual components, we consider the complementary perspective—that problems are intrinsically coupled and that the uniphysics applications often run are idealizations made in asymptotic limits. Those limits are not always carefully examined, because their examination may be nontrivial in realistic applications; hence, the importance of gaining insight from idealized situations is heightened. Because we strongly advocate examining coupling strength before pursuing a decoupled or split strategy for solving a multiphysics problem, we propose “coupled until proven decoupled” as a perspective worthy of 21st-century simulation purposes and resources.

1.1. What Constitutes Multiphysics?

Semantically, a multiphysics system consists of more than one component governed by its own principle(s) for evolution or equilibrium, typically conservation or constitutive laws. A major classification in such systems is whether the coupling occurs in the bulk (e.g., through source terms or constitutive relations that are active in the overlapping domains of the individual components) or whether it occurs over an idealized interface that is lower dimensional or a narrow buffer zone (e.g., through boundary conditions that transmit fluxes, pressures, or displacements). Typical examples of bulk-coupled multiphysics systems with their own extensively developed literature include radiation with hydrodynamics in astrophysics (radiation-hydrodynamics, or “rad-hydro”), electricity and magnetism with hydrodynamics in plasma physics (magnetohydrodynamics), and chemical reaction with transport in combustion or subsurface flows (reactive transport). Typical examples of interface-coupled multiphysics systems are ocean-atmosphere dynamics in geophysics, fluid-structure dynamics in aeroelasticity, and core-edge coupling in tokamaks. Beyond these classic multiphysics systems are many others that share important structural features.

Success in simulating forward models leads to ambitions for inverse problems, sensitivity analysis, uncertainty quantification, model-constrained optimization, and reduced-order modeling, which tend to require many forward simulations. In these advances, the physical model is augmented by variables other than the primitive quantities in which the govern equations are defined. These variables may be probability density functions, sensitivity gradients, Lagrange multipliers, or coefficients of system-adaptive bases. Equations that govern the evolution of these auxiliary-dependent variables are often derived and solved together with some of the physical variables. When the visualization is done *in situ* with the simulation, additional derived quantities may be carried along. Error estimation fields in adaptive meshing applications may constitute yet more. Though the auxiliary variables may not be “physical” in the standard sense, they give the overall simulation the structure of multiphysics.

In another important class of systems that might fall under the rubric of multiphysics by virtue of being multiscale, the same component is described by more than one formulation, typically with a computationally defined boundary or transition zone between the domains of applicability. We refer, for example, to field-particle descriptions of N-body systems in celestial mechanics or molecular dynamics, in which the gravitational or electrical forces between particles that are not immediate neighbors are mediated by a field that arises from the particles themselves. Typically, each particle defines a partition of the domain into “near” and “far” for this purpose, and the decomposition is strict. Another example is provided by atomistic-continuum models of solids, such as are used in crack propagation. In this case, the atomistic and continuum models both hold in a zone of finite thickness.

Recent schemes based on projective integration keep both fine and coarse models simultaneously in the picture and pass between the different “physics” for reasons of computational complexity. Specifically, they compute as much as possible with the coarse model, which may have constitutive terms that are difficult to derive from first principles but can be computed locally in time and space by taking suitable windowed averages or moments of the dependent variables of the fine model, which are closer to first principles. In these models, the coarse-to-fine transformation is called “lifting,” and the fine-to-coarse transformation is called “restriction.” Lifting may mean populating an ensemble of particles according to a distribution, and restriction may mean averaging.

Still other systems may have a multiphysics character by virtue of being multirate or multiresolution. A chemical kinetics model may treat some components as being in equilibrium, idealizing a fast relaxation down to a constraint

manifold on which other components vary more slowly. Some phenomena may be partitioned mathematically by projections onto wavelet bases of different frequency or wavenumber properties that are naturally treated differently.

Stretching the semantics of multiphysics still further, we may distinguish only between different mathematical formulations or even just different discretizations of what is essentially the same physical model. An example is grafting a continuum-based boundary-element model for the far field onto a finite-element model for the near field.

Systems of partial differential equations of different types (e.g., elliptic-parabolic, elliptic-hyperbolic, or parabolic-hyperbolic) for the same component may be thought of as multiphysics because each of the classical partial differential equation (PDE) archetypes represents a different physical phenomenon. Even a single equation with terms of different types represents a multiphysics model because each term must often be handled through separate discretization or solver methods.

All our examples so far ultimately define partitions or sequences of partial updates on dependent variables. With a slightly stronger semantic stretch, a system in which independent variable spaces are handled differently or independently may also have an algebraic character similar to a true multiphysics system. Examples include the “alternating direction implicit” method, in which physical space is decoupled, or methods in which physical and phase space or physical and frequency space variables are handled separately.

We enumerate these quotidian examples of multiscale, multirate, multilevel, and multimodel problems because they possess similarities to multiphysics problems, particularly in the time domain, that allow leveraging and insight. They also introduce data structure issues that can be treated in common with data structures for multiphysics simulations. Should any inhomogeneity in a problem be regarded as multiphysics? Not by traditional classifications, perhaps. However, all simulations are ultimately transformed to algebraic problems, which are then transformed to arithmetic operations over discrete data structures. In algebraic form and in low-level operations, inhomogeneity-induced blockings from disparate root causes look structurally similar. We seek both dividends of understanding and dividends of software reuse from these similarities. Here “multiphysics” is synecdoche for a broad class of coarsely partitioned problems.

Many algorithmic approaches to multiphysics simulations are “one-off” and do not leverage algebraic insights from related problems. Many software approaches are unpublished in code form, or even are proprietary. Approximations that go into the decoupling are inaccessible to measurement. These situations have retarded the growth of successful and reliable simulations of complex phenomena. Our efforts begin by establishing a common parlance for describing multiphysics systems that is simple and abstract enough to allow comparisons and reuse of software solutions.

1.2. Prototype Algebraic Forms

The two simplest systems that exhibit the crux of a multiphysics problem are the coupled equilibrium problem

$$F_1(u_1, u_2) = 0 \tag{1a}$$

$$F_2(u_1, u_2) = 0 \tag{1b}$$

and the coupled evolution problem

$$\partial_t u_1 = f_1(u_1, u_2) \tag{2a}$$

$$\partial_t u_2 = f_2(u_1, u_2), \tag{2b}$$

where the same operator notation may be overloaded in continuous or discrete settings, which are distinguished by context. When (2a–2b) is semi-discretized in time, the evolution problem leads to a set of problems that take the form (1a–1b) and are solved sequentially to obtain values of the solution $u(t_n)$ at a set of discrete times. Here u refers generically to a multiphysics solution, which has multiple components (such as those introduced in Section 1.1) indicated by subscripts $u = (u_1, \dots, u_{N_c})$; the simplest case of $N_c = 2$ components is indicated here. Noting this, we will generically use the notation in (1a–1b) to refer to either a coupled equilibrium problem or a single timestep of a coupled evolution problem.

We assume for convenience initially that the Jacobian $J = \frac{\partial(F_1, F_2)}{\partial(u_1, u_2)}$ is diagonally dominant in some sense and that $\frac{\partial F_1}{\partial u_1}$ and $\frac{\partial F_2}{\partial u_2}$ are nonsingular. These assumptions are natural in the case where the system arises from the coupling of

two individually well-posed systems with legacies of being solved separately. In the equilibrium problem, we refer to F_1 and F_2 as the component *residuals*; in the evolution problem, we refer to f_1 and f_2 as the component *tendencies*.

The choice of solution approach for these coupled systems relies on a number of considerations. From a practical standpoint, existing codes for component solutions often motivate operator splitting as an expeditious route to a first multiphysics simulation capability making use of the separate components. This approach, however, may ignore strong couplings between components and give a false sense of completion. Solution approaches ensuring a tight coupling between components require smoothness, or continuity, of the nonlinear, problem-defining functions, F_i , and their derivatives. Any potential discontinuities must be identified and dealt with before putting confidence in these approaches. For now, we illustrate multiphysics solution strategies within the context of the prototypical systems above and defer details until Section 3.

Classic multiphysics algorithms preserve the integrity of the two uniphysics problems, namely, solving the first equation for the first unknown, given the second unknown, and the second equation for the second unknown, given the first. This represents the reductionist approach of science, and as discussed in Section 4, software generally exists to do this. Multiphysics coupling is taken into account by iteration over the pair of problems, typically in a Gauss-Seidel manner (see Algorithm 1), linear or nonlinear, according to context. Here we employ superscripts to denote iterates.

Algorithm 1 Gauss-Seidel Multiphysics Coupling

```

Given initial iterate  $\{u_1^0, u_2^0\}$ 
for  $k = 1, 2, \dots$ , (until convergence) do
  Solve for  $v$  in  $F_1(v, u_2^{k-1}) = 0$ ; set  $u_1^k = v$ 
  Solve for  $w$  in  $F_2(u_1^k, w) = 0$ ; set  $u_2^k = w$ 
end for

```

When this iteration converges, the accuracy with which the discrete equations are solved can be improved by continuing the iterations. The largest implicit aggregate is the largest of the uniphysics problems; we refer to this iteration as “loosely coupled.” A Jacobi-like iteration can be similarly defined. This further decoupling exposes more parallelism, albeit possibly at the cost of a slower convergence rate.

The simplest approach to the evolutionary problem likewise employs a field-by-field approach in a way that leaves a first-order-in-time splitting error in the solution. Algorithm 2 gives a high-level description of this process, which produces solution values at time nodes $t_0 < t_1 < \dots < t_N$. Here, we use notation $u(t_0), \dots, u(t_N)$ to denote discrete timesteps. An alternative that staggers solution values in time is also possible.

Algorithm 2 Multiphysics Operator Splitting

```

Given initial values  $\{u_1(t_0), u_2(t_0)\}$ 
for  $n = 1, 2, \dots, N$  do
  Evolve one timestep in  $\partial_t u_1 + f_1(u_1, u_2(t_{n-1})) = 0$  to obtain  $u_1(t_n)$ 
  Evolve one timestep in  $\partial_t u_2 + f_2(u_1(t_n), u_2) = 0$  to obtain  $u_2(t_n)$ 
end for

```

The individual component evolutions in Algorithm 2 can be implicit or explicit and performed with or without subcycles. However, there is no point of their being of high order unless a higher-order coupling scheme than this is used, such as Strang splitting [2] for second-order or more intricate splittings or temporal Richardson extrapolations [3, 4] for higher-order. An inner loop may be placed inside each timestep in which the coupling variables are updated to satisfy implicit consistency downstream while still preserving the legacy software for each component. In the literature, this loosely coupled evolution is depicted with the diagram of Figure 1, where in the context of aeroelasticity, for instance, the first component consists of fluid velocity and pressure and the second component structural displacements.

As described in more detail in Section 2, this is an interface transmission form of coupling in which structural displacements provide boundary conditions for the fluid, and fluid pressures provide boundary conditions for the structure in the context of a dynamic mesh.

If the residuals or tendencies and their derivatives are sufficiently smooth and if one is willing to write a small

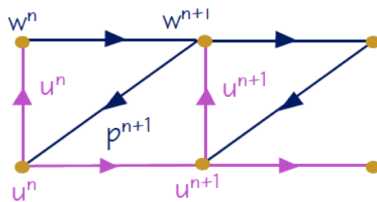


Figure 1: Illustration of loose coupling, with time advancing to the right.

amount of solver code that goes beyond the legacy component codes, a good algorithm for both the equilibrium problem and the implicitly time-discretized evolution problem is Jacobian-free Newton-Krylov (JFNK) [5]. Here, the problem is formulated in terms of a single residual that includes all components in the problems,

$$F(u) \equiv \begin{pmatrix} F_1(u_1, u_2) \\ F_2(u_1, u_2) \end{pmatrix} = 0, \quad (3)$$

where $u = (u_1, u_2)$. The basic form of Newton’s method to solve Equation 3, for either equilibrium or transient problems, is given by Algorithm 3. Because of the inclusion of the off-diagonal blocks in the Jacobian, for example

$$J = \begin{bmatrix} \frac{\partial F_1}{\partial u_1} & \frac{\partial F_1}{\partial u_2} \\ \frac{\partial F_2}{\partial u_1} & \frac{\partial F_2}{\partial u_2} \end{bmatrix},$$

Newton’s method is regarded as being “tightly coupled.” Section 3 provides more details, including the Jacobian-free Newton-Krylov variant, as well as the additive Schwarz preconditioned inexact Newton (ASPIN) [6] method, which offers advantages when the types of nonlinearities in the two components are quite different.

Algorithm 3 Newton’s method

Given initial iterate u^0

for $k = 1, 2, \dots$, (until convergence) **do**

 Solve $J(u^{k-1}) \delta u = -F(u^{k-1})$

 Update $u^k = u^{k-1} + \delta u$

end for

The operator and algebraic framework described here is relevant to many divide-and-conquer strategies in that it does not “care” (except in the critical matter of devising preconditioners and nonlinear component solvers for good convergence) whether the coupled subproblems are from different equations defined over a common domain, the same equations over different subdomains, or different equations over different subdomains. The general approach involves iterative corrections within subspaces of the global problem. All the methods have in common an amenability to exploiting a “black-box” solver philosophy that amortizes existing software for individual physics components. The differences are primarily in the nesting and ordering of loops and the introduction of certain low-cost auxiliary operations that transcend the subspaces.

Not all multiphysics problems can easily or reliably be cast into these equilibrium or evolution frameworks, which are primarily useful for deterministic problems with smooth operators for linearization. In formulating multiphysics problems, modelers first apply asymptotics to triangularize or even to diagonalize the underlying Jacobian as much as possible, pruning provably insignificant dependences, but bearing in mind the conservative rule of coupling: “guilty until proven innocent.” One then applies multiscale analyses to simplify further, eliminating stiffness from irrelevant mechanisms. Because of significant timescale and resolution requirement differences, these activities often reveal justifications for splitting some physics or models from others. In these cases, an operator splitting approach should be applied. However, caution is warranted to ensure that temporal or spatial scales do not overlap between split physics as the simulation progresses. If this happens, a more coupled approach is generally required.

1.3. Structure and Motivation for This Review

The previous section discussed a two-component multiphysics system at the operator and discrete level, primarily in terms of operations that are already available in simulations for the individual components, introducing generic notation that will be useful throughout the paper. Generalization of the algorithms to N_c components is trivial. However, the convergence of the methods is not guaranteed.

Applications are described in terms of this notation in Section 2; model problems are treated by these methods in Section 3, and software that implements data structures and coupling to accommodate these methods is discussed in Section 4. In principle, most of the multiphysics applications can be made to exhibit the partition above, with variations on the functional forms of the cross-component dependences.

Even when all the components entered into a multiphysics simulation are well verified, the peril of operator splitting is that systems coupled by new off-diagonal terms may admit destabilizing modes not present in any component system alone. Even when the coupling is stable in a continuous sense, the accuracy of the coupled system and the stability of the numerical discretization can be compromised by operator splitting, as shown in Section 3.4.

Couplings that are relaxed at the component level may include fluxes at interfaces between physical domains, bulk source terms, complex constitutive laws, or cross-dimensional effects. With the advent of nonlinear implicit solvers and their availability in well-established packages such as PETSc [7, 8], SUNDIALS [9, 10], and Trilinos [11, 12], there exists a robust alternative to operator splitting; however, preconditioning fully temporally implicit discretizations remains challenging, particularly with advanced, high-order spatial discretizations. For efficiency, it is important to understand the interaction of the many forms of errors and mechanisms of instability, as well as the relative costs of reducing them, particularly in the context of emerging architectures, on which traditional flop-based complexity is less relevant than the complexity of transmitting or copying data.

Despite these mathematical perils, computational models of multiphysics phenomena are often approached through operator splitting, which has attractions from a software engineering viewpoint. However, in an era of architectures that penalize data motion but allow nearly “free” flops, splitting usually results in an increase (proportional to the degree of the splitting) in the ratio of loads and stores relative to flops. That is, operator splitting shrinks the arithmetic intensity of the simulation, which is the wrong direction of code evolution. The move toward more tightly coupled approaches looms as a prime opportunity for adaptation to the exascale, with its memory bandwidth stringencies [13]. In Sections 5 and 6 we address opportunities in multiphysics applications, propose transferable best practices, and emphasize applications at extreme scale. The topics to be addressed include physical modeling (fidelity, ranges of applicability), mathematical analysis (formulations, well-posedness in the continuum sense), numerical analysis (accuracy, stability), algorithms (complexity, optimality), software engineering (programmer productivity, reusability), and implementation on emerging architectures (performance, portability, scalability). Section 7 summarizes our conclusions. Appendix A contains a glossary of basic multiphysics terms as used in this report. In Appendix B we propose the start of an illustrative multiphysics problem suite.

We conclude this section by quoting from recent reports of DOE communities:

The dominant computational solution strategy over the past 30 years has been the use of first-order-accurate operator-splitting, semi-implicit and explicit time integration methods, and decoupled nonlinear solution strategies. Such methods have not provided the stability properties needed to perform accurate simulations over the dynamical timescales of interest. Moreover, in most cases, numerical errors and means for controlling such errors are understood heuristically at best. (H. Simon, R. Stevens, T. Zacharia, eds., 2007 [14])

The following priority research direction [was] identified: develop scalable algorithms for non-hydrostatic atmospheric dynamics with quasi-uniform grids, implicit formulations, and adaptive and multiscale and multiphysics coupling ... Improvements in scalability alone will not be sufficient to obtain the needed throughput (the time it takes to complete a climate simulation). Obtaining the needed level of throughput will also require incorporating as much implicitness as possible ... (W. Washington, ed., 2008 [15])

Often, scientists would ideally employ a high-order timestepping scheme and take relatively large timesteps for computational economy. However, if operator-splitting techniques are used, the low-order splitting

error thwarts this objective. Moreover, computational challenges on the immediate horizon – optimization for design or control, inverse problems for parameter identification, multiphysics coupling, etc. – are most naturally tackled with fully implicit formulations. (W. Tang, D. Keyes, eds., 2009 [16])

The great frontier of computational physics and engineering is in the challenge posed by high-fidelity simulations of real-world systems, that is, in truly transforming computational science into a fully predictive science. Real-world systems are typically characterized by multiple, interacting physical processes (multiphysics), interactions that occur on a wide range of both temporal and spatial scales. (Advanced Scientific Computing Advisory Committee (ASCAC) Subcommittee on Exascale Computing, R. Rosner, ed., 2010 [17])

2. Practices and Perils in Multiphysics Applications

In this section, we discuss examples of multiphysics applications, selected from diverse disciplines within science and engineering. By outlining these real-world examples, we illustrate, from a practical standpoint, how multiphysics applications are assembled and used. These examples also illustrate the rich combinatorics of possible couplings among uniphysics components. In taking this approach, we provide a departure point from which to discuss issues that can be encountered by naively coupling two or more robust uniphysics codes to form a multiphysics application.

One of the most important of these issues is the coupling itself—a term that we now define and discuss further in Section 2.2.1. The concept of physical coupling within a multiphysics application is intuitive. However, upon continued discussion, the term can become increasingly vague since coupling exists at least at both a physical level and an algorithmic level. Hence, some definitions are in order. Throughout this report we refer to *strong (versus weak) coupling of physical models* as meaning that physical components of the modeled system are coupled by a strong (versus weak) interaction. Correspondingly, we refer to *tight (versus loose) coupling of numerical models* as meaning that algorithmic components are coupled by a tight (versus loose) interaction. See Appendix A for further discussion.

We note that there is not a direct one-to-one correspondence between strong (weak) coupling of physical models and tight (loose) coupling of numerical models, since a tight or a loose coupling scheme may be used for strongly or weakly coupled physical models. We also note that terminology has often been used interchangeably in the literature; for example, in fluid-structure interactions, tight coupling schemes and strong coupling schemes have been discussed interchangeably.

To illustrate the distinction, the effect of the fluid pressure on the displacement of a deformable structure and of the structural displacement on the fluid velocity are instantaneous; thus, the physical interaction is strong. However, it is possible to update the structural displacement, based on a new fluid pressure distribution, without simultaneously adjusting the fluid velocity to respect its new bounding surface. If such is the case, the numerical coupling is loose, and iteration between the two phases may be required to reduce the error of the tight interaction sufficiently. Conversely, in the tightest numerical coupling, the state variables of the two phases would be updated simultaneously, typically in an algorithm of greater complexity. Either the tight or loose numerical coupling could satisfy a given accuracy, and either may be more computationally efficient depending upon the problem and how they are applied.

In a reacting flow, the concentration of a chemical species may have insignificant influence on the thermodynamic state of the fluid within a broad range, and may also be insensitive to the thermodynamic state; thus, the physical interaction is weak. However, the numerical model may evolve these state variables independently (weak coupling) or simultaneously (strong coupling), though the latter effort produces no benefit in this case.

We further comment that physical interaction may be strong in one direction and weak in the other, or strong in both directions. We describe such a system as exhibiting one-way or two-way coupling, respectively. One-way coupling leads naturally in limit to a numerical system that is triangular.

We consider such coupling issues for a variety of PDE-based multiphysics applications in Section 2.1. In Section 2.2 we identify commonalities and differences across the application spectrum, and we touch on issues in algorithms and software from an applications perspective, as a foundation for more detailed discussion in Sections 3 and 4.

2.1. Examples of PDE-Based Multiphysics Applications

A broad range of multiphysics simulations are under way in the computational science community, as researchers increasingly confront questions about complex physical and engineered systems characterized by multiple, interacting physical processes that have traditionally been considered separately. Sections 2.1.1 through 2.1.11, introduce some large-scale PDE-based multiphysics applications, including fluid-structure interaction, fission reactor fuel performance, reactor core modeling, crack propagation, fusion, subsurface science, hydrology, climate, radiation hydrodynamics, geodynamics and magma dynamics, and accelerator design. Discussion focuses on current practices and challenges in application-specific contexts.

2.1.1. Interaction of fluids and structures

Numerical simulations that model interaction between incompressible laminar flows and elastic structures require coupling a description of the fluid—typically the incompressible Navier-Stokes equations or a weakly compressible lattice-Boltzmann equation—with a description of the structures. Sample application areas for this scenario include

blood flow in arteriae, veins, and heart chambers; low Mach number aerodynamics; marine propellers; and hydroelectric power plants. The unknowns of the involved equations—velocities and pressure for the fluid, displacements for the structure—are associated with different locations in the overall computational domain, resulting in a surface-coupled problem.

Fluid–structure interaction (FSI) can be simulated in at least two ways. One approach is to solve a large system of equations for all fluid and structure unknowns as a single system—typically ill-conditioned. Alternatively, in a partitioned approach, one has separate solvers for the fluid and the structure, together with a suitable coupling method. In the latter case, boundary conditions for both single physics problems at the coupling surface have to be defined. The respective interface values are passed from one solver to the other. This approach requires mapping methods for physical variables between (in general) nonmatching solver grids at coupling surfaces. Important features desired of such mappings are accuracy, consistency, and conservation of energy and momentum. Two main classes of mapping methods can be identified: interpolation methods [18–22], based on geometric relations between the involved grid points, and mortar methods, in which boundary conditions are formulated in weak form by using Lagrange multipliers [18, 23–25].

The coupling itself can be done with different methods, leading to looser or more tightly coupled timestepping methods; see Figure 2 for two variants. The loosest coupling is a one-way coupling, where the flow solver computes a force exerted on the structure using a rigid-structure geometry, and structural movements are computed in a postprocessing-like manner based on these forces. This strategy is obviously applicable only for small and static structure deformations. The most widely used class of iteration schemes is Gauss-Seidel-like coupling iterations (see Algorithm 1), with variants ranging from a single iteration loop per timestep to repeated iteration-to-convergence within each timestep, with or without (Aitken) underrelaxation [26–28], to interface quasi-Newton methods that efficiently compute approximate Newton iterations based on sensitivities resulting from Gauss-Seidel iterations [29]. To account for what are usually moderately different timescales in the fluid and the structure, a subcycling in the flow solver can be used. In the absence of turbulence, spatial scales are essentially the same throughout fluid and structure domains.

The choice of coupling implementation can lead to development of numerical instabilities in multiphysics codes. For incompressible fluids, the so-called added-mass effect induces instabilities in loosely coupled simulations and in Gauss-Seidel-like iterations: the force exerted by the fluid on a moving structure can be interpreted as a virtual added mass (see, e.g., [30]) of the structure. For an incompressible flow, each acceleration or deceleration of the structure causes an immediate change in this added mass (whereas the added mass change for a compressible flow increases continuously over time). If this change is too large, both loosely and tightly coupled Gauss-Seidel-like coupling schemes become unconditionally unstable; in other words, a reduction of the timestep does not cure the instability [31]. In the case of a massless structure, a reduction of the timestep even worsens instabilities [29]. Typically, only a few low-frequency Fourier modes of interface displacements or velocities are unstable [29, 31]. Interface quasi-Newton methods rapidly capture the unstable Fourier modes and thus lead to stable coupling iterations [29]. However, standard multigrid techniques, with Gauss-Seidel-like iterations as a smoother, do not work for incompressible flows. Although they smooth high frequencies, they are unstable for low frequencies. In addition, the convergence rate for low-frequency modes does not improve on spatially coarser grids.

Since Gauss-Seidel-like coupling iterations, including the interface quasi-Newton method, are inherently sequen-

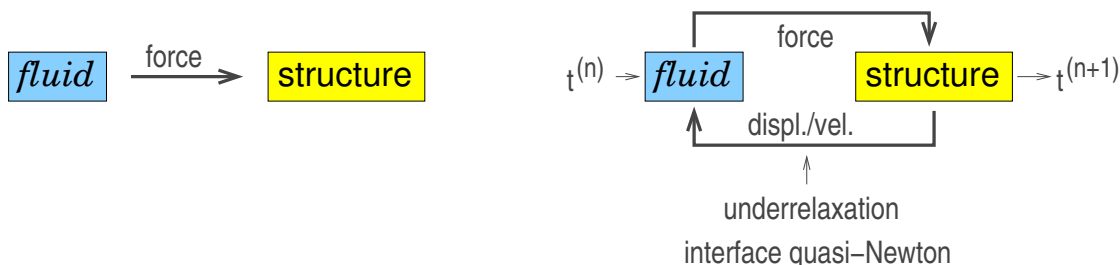


Figure 2: One-way and Gauss-Seidel-like coupling strategies for fluid-structure interaction simulations.

tial in execution order of flow and structure solver, and since the structure solver is usually much less costly than is the flow solver, alternative Jacobi-like methods (see [25] for compressible flows) are desirable on massively parallel systems. These avoid processor idle time during the execution of the structural solver and the coupling numerics.

As a practical application of FSI, we consider examples in high-performance, high-fidelity analysis, which finds application in modeling the flow of airstreams over aircraft and vehicles as well as fluid flow over underwater structures. Airflow problems necessitate the solution of nonlinear compressible FSI problems. Specific examples include the parametric identification of an F-16 Block-40 fighter aircraft in clean wing configuration and subsonic, transonic, and supersonic airstreams [32, 33]; the aeroelastic analysis of an F/A-18 5.6 Hz limit cycle oscillation configuration; the flutter clearance of the laminar flow wing of a supersonic business jet concept; and the aeroelastic tailoring of a Formula 1 car. Examples of incompressible FSI problems include the study of intense implosive collapses of gas-filled underwater structures and their effects on nearby structures [34].

Many of these problems feature multiple spatial and temporal scales. In all of them, the fluid domain can be occupied by one or multiple interacting fluids, the structure can be linear or nonlinear with self-contact, and the fluid-structure interaction occurs at physical interfaces.

The AERO code [32] is an example of a multiphysics FSI software package that considers these issues. The functional modules of the AERO code include some of the most important considerations for producing accurate models: (1) the structural analyzer AERO-S, (2) the compressible turbulent flow solver AERO-F, (3) the auxiliary module MATCHER, which enables the discretization of fluid-structure transmission conditions at nonmatching, discrete fluid-structure interfaces [35], and (4) the auxiliary module, SOWER, which manages all parallel I/O associated with this software. AERO-F can operate on unstructured body-fitted meshes as well as fixed meshes that embed discrete representations of surfaces of obstacles around and/or within which the flow is to be computed. The body-fitted meshes and the embedded discrete surfaces can be fixed, move, and/or deform in a prescribed manner or can be driven by interaction with the structural analyzer AERO-S. In the case of body-fitted meshes, the governing equations of fluid motion are formulated in the arbitrary Lagrangian–Eulerian framework. In this case, large mesh motions are handled by a corotational approach, which separates the rigid and deformational components of the motion of the surface of the obstacle [36] and robust mesh motion algorithms that are based on structural analogies [37]. In the case of embedded surfaces that can have complex shapes and arbitrary thicknesses, the governing equations of fluid motion are formulated in the Eulerian framework, and the wall boundary or transmission conditions are treated by an embedded boundary method [38]. Both AERO-F and AERO-S feature explicit and implicit time integrators with adaptive timestepping. Both modules are coupled by a family of partitioned analysis procedures that are loosely coupled but exhibit excellent numerical stability properties and are provably second-order time accurate [39]. AERO-F and AERO-S communicate via runtime software channels, for example, using MPI. They exchange aerodynamic (pressure) and elastodynamic (displacement and velocity) data across nonmatching, discrete fluid and structure mesh interfaces using a conservative method for the discretization of transmission conditions [18] and the data structures generated for this purpose by MATCHER.

2.1.2. Fission reactor fuel performance

Simulating the performance of light water nuclear reactor fuel involves complex thermomechanical processes between fuel pellets, made of fissile material, and the protective cladding that surrounds the pellets [40]. Figure 3 shows the relationship between the fuel pellet and cladding and the helium-filled gap separating them. Advancing fuel design and analysis capability requires progress both in the modeling of the fuel material behavior and in the development of a multidimensional computational framework that couples these models in a consistent, accurate, and efficient manner to describe how the fuel geometry and behavior change over the lifetime of the fuel.

The base physics begins with heat diffusion that partially drives the mechanical response of the fuel pellet and cladding system [41]. Thermomechanical response problems are often solved in an operator splitting fashion, but this approach can result in convergence and stability issues as the pellet and cladding come into contact as the fuel ages. Newman et al. [42] solve the fuel performance problem in a fully coupled manner on a single mesh. Further, an oxygen diffusion equation is added since, under hyperstoichiometric conditions, both the thermal conductivity and mechanical properties of the fuel pellet are impacted by the diffusion of oxygen within the matrix. Additional material models are used to describe other details of the thermal and mechanical behavior of the fuel and cladding, such as swelling of the fuel due to the accumulation of fission products within it [43]. The expansion of the pellet

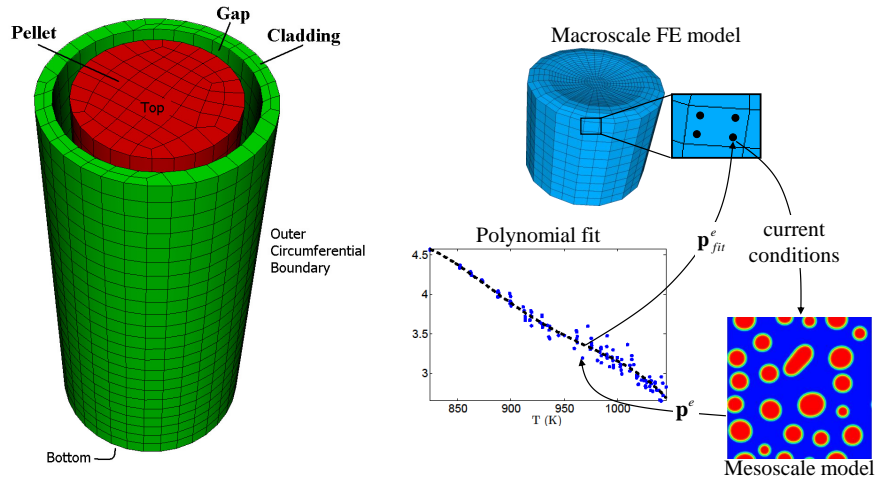


Figure 3: This figure illustrates that the temperature at the integration points is coupled to the phase field simulation, resulting in a thermal conductivity field. The conductivity values are used to form a regularized thermal conductivity surface that is subsequently used to solve the heat diffusion equation, thus completing the multiphysics coupling loop. **Left:** Meshed representation of a short section of the fuel pellet and cladding, showing the gap between the two (exaggerated in the diagram). **Right:** Pictorial showing the multiscale coupling between a macroscale finite element and a mesoscale phase field model.

and, indirectly, the width of the gap separating the pellet and the cladding are functions of the thermal profile within the pellet, the age of the fuel (the degree of fission product driven swelling), and the creepdown of the cladding onto the outside surface of the pellet.

Given the significant radial and axial displacement of the pellet relative to the cladding, it is infeasible to mesh the helium-filled gap between the two, especially in a multiple pellet simulation. Since the gap is long and narrow, Rayleigh numbers are unlikely to be high enough within the gap to support significant convection. However, it is important to accurately model the conductive heat transfer across the gap, since the gap is one of the principal factors that influence performance of the fuel. Thus, a fully coupled thermomechanical contact model was developed to model the evolution of the gap within the fully coupled multiphysics code [44].

Irradiation of the fuel and the cladding affects the thermal and mechanical properties of both. It is difficult to describe these effects on the coarse scale of the finite-element code described above. Modeling the irradiation of the material is more easily done at lower-length scales, such as the mesoscale. One might employ a phase-field simulation approach to calculate the evolution of thermal conductivity of the fuel under irradiation. Use of this model can be challenging, however, since the characteristics of the irradiation damage change with fuel temperature. In [45] and [46], a mesoscale solution of the Cahn-Hilliard equation is coupled to the macroscale approach to form a multiscale simulation. The mesoscale simulation requires coupling a temperature field from the macroscale as input to the simulation. Further, the thermal conductivity calculated at the mesoscale impacts the temperature field calculated on the macroscale, as shown in Figure 3(b). In the approach, a mesoscale calculation is performed at selected integration points of the macroscale finite-element problem, where temperature from the macroscale is used to initialize the mesoscale calculation. The thermal conductivity at the mesoscale is then homogenized to serve as a value that applies to the corresponding integration point.

The equations that describe reactor fuel performance are solved by using a fully coupled, Jacobian-free Newton Krylov approach [42, 47, 48] (see Algorithm 3 and Section 3.1.2). A fully implicit time integration strategy is used to support accurate time integration needed over short transient events, such as the initial warmup of the reactor and power transients during operation, and to allow long timesteps to be taken under the long quasi-static periods of normal operation. All the phenomena within the reactor fuel are tightly interrelated, particularly the mechanical aspects of the fuel. As the fuel warms up, the pellet expands, thereby narrowing the gap and increasing the pressure of the gases between the pellet and cladding. The narrowing gap increases the heat transfer efficiency across it, which cools the pellet, checking its expansion. Attempting to decouple these processes in order to simplify the solution process can result in convergence issues and even “numerical chatter,” where the pellet and cladding chatter on incipient contact

instead of arriving at an equilibrium solution, as described by Hansen [44]. In the gap integration technique presented here, fluxes and tractions are exchanged between the pellet and cladding using the mortar finite-element method. The integration space and Lagrange multiplier implementation employed the Moertel package [49] in Trilinos.

2.1.3. Conjugate heat transfer/neutron transport coupling in reactor cores

Accurately predicting neutron densities in a reactor core is a critical aspect of design and safe operation. Reactor cores are designed to balance power generation with efficient heat transfer and fuel burnup objectives, optimizing energy density relative to design limits for key global properties such as peak temperature and temperature gradient, neutron fluence, and pumping power (pressure drop).

On relatively fast timescales (e.g., compared with the physics of neutronics/fuel coupling) and within certain parameter regimes, methods for computing the accurate coupling between the neutron transport and conjugate heat transfer, for fixed fuel and structural configuration, becomes the key problem of interest for reactor analysts. Historically, reactor design and analysis tools have been based on weak coupling between bulk fluid characterizations and few-group diffusion or nodal neutronics methods. The Simulation-based High-efficiency Advanced Reactor Prototyping (SHARP) [50] project is investigating accurate, high-fidelity fluid/neutronic coupling methods for a range of reactor phenomena.

Assuming that suitable boundary conditions are applied, the coupled reactor core neutronics/heat transfer equations can be written as a single coupled system:

$$\partial_t f = L(f) + N(f),$$

where $L(f)$ denotes the linear part of the operator, $N(f)$ the nonlinear part, and $f = [\psi \ T \ \rho \ w]^T$. Here ψ is the angular neutron flux, T the medium temperature, ρ the medium density, and w the three-dimensional coolant velocity. The temperature and density can be further divided into coolant and fuel regions, denoted by $T = T_c \cup T_f$ and $\rho = \rho_c \cup \rho_f$. The angular flux ψ is then obtained from the linear transport equation:

$$\left[\frac{1}{v} \frac{\partial}{\partial t} + \hat{\Omega} \cdot \nabla + \sigma_{\rho, T}(r, E) \right] \psi(r, \hat{\Omega}, E, t) = q_{ex}(r, \hat{\Omega}, E, t) + \int dE' \int d\Omega' \sigma_s(r, E' \rightarrow E, \hat{\Omega}' \cdot \hat{\Omega}) \psi(r, \hat{\Omega}', E', t) + \frac{\chi(E)}{4\pi} \int dE' \nu_f \sigma(r, E') \int d\Omega' \psi(r, \hat{\Omega}', E', t). \quad (4)$$

In Equation (4), v is the scalar neutron speed, E the neutron energy, $r = (x, y, z)$ the spatial coordinate, t the time, $\hat{\Omega} = (\theta, \phi)$ the direction of neutron travel, q_{ex} an external neutron source, $\sigma_{\rho, T}$ the total interaction cross section, σ_s the scattering cross section, χ the fission spectrum, and ν_f the number of neutrons emitted per fission. The subscripts on $\sigma_{\rho, T}$ are explicitly included to denote the dependence of cross sections on both the fuel and coolant temperature and density, which is the principal source of (nonlinear) coupling between different physics modules.

With a solution to Equation (4), the volumetric heat generation rate q''' can be estimated by

$$q'''(r, t) = \int dE' \sigma(E') W(E') \int d\Omega' \psi(r, \hat{\Omega}', E', t).$$

The heat equation for both fluid and fuel can then be solved with q''' as a source (where $w = 0$ in the fuel region):

$$\rho C_p \left(\frac{\partial}{\partial t} + w \cdot \nabla \right) T(r, t) = \nabla \cdot \kappa_T \nabla T(r, t) + q'''(r, t) + f. \quad (5)$$

The fluid velocity w in Equation (5) is obtained from the Navier-Stokes equation. For simplicity, we write the Boussinesq approximation of the equation where density changes are considered negligible except in the buoyancy term (in general, an equation of state is needed to close the system).

$$\frac{\partial w}{\partial t} + w \cdot \nabla w = -\frac{1}{\rho_0} \nabla p + \nabla \cdot \nu_T \nabla w + \frac{\rho'}{\rho_0} g \hat{k} \quad (6)$$

$$\nabla \cdot \rho_0 w = 0$$

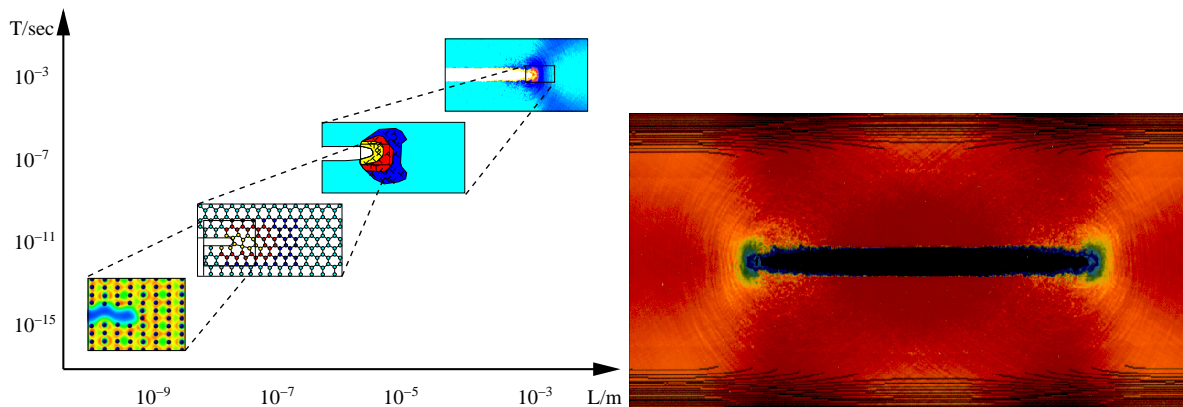


Figure 4: Illustration of the length and timescales in crack propagation. **Left:** Schematic representation of the length and timescales involved. **Right:** Multiscale simulation of crack propagation in silicon, a prototypical brittle solid, which includes all the relevant length-scales; the color maps show the distribution of stress (highest at the tip of the crack), with stress waves passing seamlessly from one region to the next (see [61]).

Equation (6) is nonlinearly coupled to the heat equation (5) through the temperature dependence of the dynamic viscosity ν , denoted explicitly by ν_T . The details of the discrete formulations of the above equations are reported in companion papers [51–53] and are not essential to understanding the intermodule coupling problem.

When the physics is split into separate components, these couplings take the form of specific data values that are interpolated from source to target meshes and sent between the modules through well-defined interfaces. SHARP includes the following physics modules to solve the above equations: (1) Nek [54, 55], a spectral element code that solves the 3D, incompressible (Boussinesq), time-dependent Navier-Stokes equation with conjugate heat transfer on unstructured, higher-order quadrilateral meshes; and (2) UNIC [53], an unstructured, deterministic neutron transport code that incorporates parallel even parity, sweeping, and ray-tracing algorithms.

Separating physics modules into distinct components and implementing coupling as interactions between these components imposes a number of requirements on the overall design of the SHARP framework, discussed in [50]. SHARP couples physics components and other services through the spatial discretization or mesh provided by MOAB [56–58], which is introduced in Section 4.1.3.

New research in the Center for Exascale Simulation of Advanced Reactors (CESAR) [59] is building on this foundation to develop a coupled, next-generation nuclear reactor core simulation tool capable of efficient execution on exascale computing platforms. The approach incorporates tight coupling of extreme-fidelity models of neutron transport and conjugate heat transfer.

2.1.4. Multiscale methods in hard and soft materials: crack propagation and ultrafast DNA sequencing

Most physical phenomena of interest to humankind involve a range of temporal and spatial scales. For example, brittle fracture of solids under external stresses can lead to the failure of large structures such as bridges or ships but originate at the atomic scale with the propagation of microcracks. Likewise, tidal currents in bays extend over many miles, but their behavior is dictated by the water viscosity, determined from the molecular-scale interactions of water molecules. A computational approach that captures the complexity of such physical systems must be able to treat all the relevant scales at the level of accuracy required for a realistic description. We illustrate the issues involved in such multiscale simulations through two examples that require a broad range of scales, from the continuum to the atomistic and even to the electronic level (for a review of similar problems and methods for their simulation, see [60]). The two examples are chosen to demonstrate both the progress that has been made and important challenges that remain in the areas of hard and soft materials.

Our first example involves crack propagation and fracture of solids. Crack propagation is often the result of chemical impurities (corrosion), which can change the mechanical behavior of a solid from ductile (tough, difficult to break solids) to brittle (weak, easily breakable ones). The length scale at which this phenomenon is observed in everyday situations is roughly millimeters and above. At this scale cracks are detectable by the naked eye and, if not

arrested, can lead to the failure (occasionally catastrophic) of large structures such as ships, airplanes, and bridges. The presence of chemical impurities, often in minute concentrations (a fraction of a percent or lower), changes the nature of the bonds between the atoms, which is crucial in the region near the tip of the crack where bonds are being stressed to the point of breaking. These atomic-scale changes are ultimately responsible for the macroscopic-scale change in the mechanical behavior of the solid.

A schematic breakdown of the different length and scale regimes relevant to this phenomenon is illustrated in Figure 4. The macroscopic scale (10^{-3} m or above) on which the crack propagation is of practical interest contains a huge number of atoms, of order 10^{24} , and can be handled only by continuum approaches such as elasticity theory. The response of the solid to external stresses is determined by the motion of defects such as dislocations, which provide the mechanism for deformation. Their distribution and motion at the mesoscopic scale (10^{-5} m) can be described by theories that explicitly take into account these defects and the pattern of stress they create within the solid [62–65]. The origin of dislocations and the nature of the interactions between them can be described only by considering their structure at the microscopic scale (10^{-7} m) in terms of atomic degrees of freedom interacting by springlike forces [66]. Furthermore, the presence of chemical impurities and their effect on atomic bonding, and by extension of the forces that hold atoms together near the tip of a crack, must include a quantum-mechanical description of the ion-electron interactions that can capture the chemical properties of different types of atoms. This scale can be treated in the context of density functional theory using, for example, methodologies such as the Car-Parrinello approach [67]. A first attempt to capture all these scales for the case of fracture in silicon, a prototypical brittle solid, was the multiscale-atomistic-ab initio-dynamics approach, which was able seamlessly to couple within the same method the quantum-mechanical scale to the classical atomistic and to the continuum [61]. A direct link between the atomistic scale and the scale of defects (dislocations, domain boundaries, etc.) dynamics was accomplished by the quasi-continuum method [68], which enabled realistic simulations of large deformation of materials [69, 70]. These examples addressed effectively many of the computational issues in the coupling of length scales for multiscale simulations of materials. For timescale issues, the methods originally developed by Voter [71, 72] for accelerating molecular dynamics can extend the scale of atomistic motion (10^{-15} s) by several orders of magnitude, addressing an important bottleneck in multiscale simulations.

Our second example concerns ultrafast sequencing of DNA through electronic means, a concept being vigorously pursued by several experimental groups [73–78]. The idea is to form a small pore, roughly of the same diameter as the DNA double helix, and detect the DNA sequence of bases by measuring the tunneling current across two electrodes at the edges of the pore during the translocation of the biopolymer from one side of the pore to the other. Estimates of the rate at which this sequencing can be done are on the order of 10 kbp/s, which translates to sequencing the entire human genome in a couple of days. While this goal has not been achieved yet, the successful demonstration of this concept has the potential to produce long-lasting changes in the way medicine is practiced.

This system is also complex, involving several scales that are shown schematically in Figure 5. At the coarsest scale, the biopolymer in solution needs to be directed toward the nanopore. The scale of the biopolymer in solution is on the order of micrometers, and the timescale involved in finding the pore is on the order of fractions of a second. At these scales, the system can be reasonably modeled as consisting of a continuous polymeric chain in a uniform solvent under the influence of an external nonuniform but continuous field that drives it toward the pore. Understanding the translocation process involves modeling the polymer at scales set by the persistence length (~ 50 nm for DNA), while including the effect of the solvent. The coupling of the molecular motion to the fluid motion has been successfully demonstrated, with the solvent described by the lattice-Boltzmann approach to modeling fluid dynamics [79, 80] and the polymer motion described by Newtonian dynamics under the influence of local friction from the fluid motion [81]. This lattice-Boltzmann plus molecular-dynamics (LBMD) description has successfully reproduced details of the observed translocation process in a realistic manner, as observed in experiments for both single-file and multifile translocation. An extension of the LBMD method to blood flow in the main arteries of the heart, including the explicit dynamics of red blood cells, was able to capture both the effects of non-Newtonian blood flow on the arterial walls (the “endothelial shear stress,” linked to heart disease) and the formation of correlated cell motion [82]. The coupling to the atomistic and electronic scales, required to capture the whole process of electronic sequencing during translocation, remains a significant challenge. The electron current can be calculated from quantum-transport theories, once the relevant electronic states are known. To this end, detailed calculations of the electronic properties of DNA bases need to be carried out, at the level of quantum chemistry, including all the details of the atomistic structure dictated by the morphology during the translocation process. Assembling the entire picture in a coherent model is a

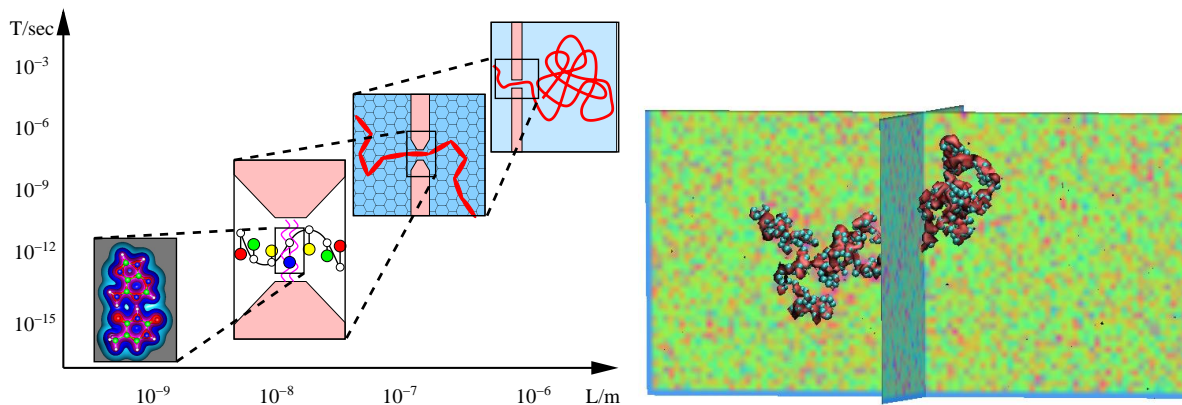


Figure 5: Illustration of the length and timescales in ultrafast DNA sequencing by electronic signals during translocation through nanopores. **Left:** Schematic representation of the process and the various scales involved. **Right:** multiscale simulation coupling the molecular scale of DNA, represented by the blue spheres at the molecular persistence-length scale, to the motion of the fluid solvent whose local velocity is given by the colored contours.

serious challenge to current computational capabilities and is likely to require significant conceptual and algorithmic innovations in the way we handle the phase-space of complex systems.

These two examples illustrate several of the issues that arise in the context of multiscale simulations for hard and soft materials. The physics coupled in both cases is the motion of atoms and electrons, at the finest scale, as dictated by boundary conditions on the system at macroscopic scales, where continuum theories can be applied. The two ends of the spatial scale are intimately coupled, with the atomic motion constrained by the imposed boundary conditions and, in turn, determining the response of the system to the externally applied stress. The information passed between the scales is essentially the forces on the key components at each scale, which are derivatives of the appropriate energy functionals. These functionals have explicit dependence on degrees of freedom that are coupled across the scales.

Two key sources of errors and instabilities in these multiscale simulations are the differences in the energy functionals across two scales and the loss of information in passing from the finer to the coarser scale. The differences in energy functionals lead to the appearance of “ghost forces,” since the physical quantities that underlie each functional may have a range much larger than the boundary region between the two scales. A case in point is the range of electronic wavefunctions, which can extend over much larger distances than what is practical to capture with a multiscale method involving a quantum-mechanical description at the finest length scale (for metallic solids, these wavefunctions may extend over the entire range of the solid). The second limitation, the loss of information due to coarsening of degrees of freedom, is intrinsic to the multiscale problem and can be addressed to some extent by performing finite-temperature simulations at the atomistic scale. These, however, are both extremely computationally demanding and restricted to a single temperature value. A more general framework is required to overcome this limitation.

The most severe limitation of multiscale methods that include atomistic or electronic scales is that the timestep is dictated by the dynamics at the smallest scale, which for atomic motion is on the order of 10^{-15} s. Interesting methodologies for extending the timescale have been proposed (see, e.g., [71, 72]), but they have their own limitations in the size of system that can be treated or the range by which the timescale can be extended. This issue remains one of the most challenging theoretical problems in performing multiscale simulations for materials problems that can be of direct relevance to real applications.

2.1.5. Simulation of processes in magnetic confinement fusion

Magnetic fusion is a long-term solution for producing electrical power for the world, and the large international device (ITER) [83] being constructed will produce net energy and a path to fusion energy provided the computer modeling is accurate. Fusion computation is enormously difficult because of the wide range of temporal and spatial scales in fusion devices [16, 84]. Consequently, the fusion community has, over decades, developed a suite of models for tokamaks (and other confinement devices).

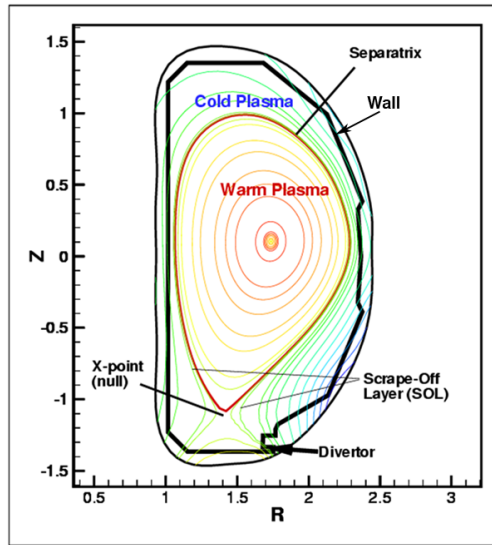


Figure 6: Schematic diagram of a tokamak cross section. The thick black line shows the material wall. The region inside the separatrix (closed brown contour) is the core, while outside is the edge. The physics in the core, edge, and material walls occur at different timescales and include different physical phenomena. Coupled simulations are needed to self-consistently understand the physics of the tokamak in order to optimize the fusion gain.

A model that quantitatively accounts for all the interacting physical processes at work in a tokamak discharge has come to be known as a whole device model (WDM). Such a model must account for nonlinearly interacting phenomena, including sources of energy and particles into the plasma; transport of energy, particles, and momentum through the plasma; interface between the hot central plasma and the cool edge plasma and the material walls; and interaction of the plasma with external magnetic coils, circuits, and control systems. While work of this kind has been going on in fusion research for a very long time, limitations of theory and computer power traditionally have restricted such modeling to much-simplified models for the physical processes. Including this physics in a single, integrated, multiphysics, multiscale simulation presents many challenges of physics, applied mathematics, and computer science. Beyond the mathematics and computer science required to realize the physics components in an integrated simulation, there are important mathematical issues related to model coupling and stable, accurate timestepping algorithms, as well as computer science issues of data communication and efficient management of very large-scale computer resources.

As an example of one significant WDM project, FACETS, the Framework Application for Core-Edge Transport Simulations [85–87], has the goal of providing modeling of a fusion device from the core to the wall. Figure 6 shows the cross section of a tokamak plasma, which has a toroidal (donut) shape. The separatrix (so labeled) separates hot plasma in the center from the cooler, surrounding plasma. The central plasma is hot because it can move primarily only along the magnetic field lines, and those always remain in the core of the system. Thus, the plasma is never in direct contact with material walls, which would cool the plasma. In the edge region of the plasma, just outside the separatrix, the magnetic field lines (projections on a poloidal plane shown in green) wrap around the core but then terminate on the walls. This contact with the walls leads to much cooler plasma.

In the two regions, the plasma is analyzed with different approximations. In the core, the velocity distributions of the particles are nearly Maxwellian; and bulk parameters, such as temperature and density, are nearly constant along field lines and, hence, on flux surfaces, the surfaces on which the field lines remain. This constancy in flux surfaces implies that the plasma has only one-dimensional variation, on the dimension that crosses flux surfaces. In the edge, the plasma varies strongly along a field line, from the region in the midplane near the core to the ends, which are in contact with the walls. Hence, in this region the plasma has two-dimensional variation. Being cooler, the edge region can have a high density of atomic species, whereas in the core the plasma is for the most part fully ionized.

Just inside the separatrix is a transition region, where the cool, two-dimensional plasma of the edge morphs into

the hot, one-dimensional plasma of the core. The atomic physics processes become less important; and the character of the turbulence, which controls the plasma confinement, changes. In this region a pedestal may form. At the pedestal there can be nearly a transport barrier, such that the temperature of the plasma can be very different across it. Good pedestals are needed for the beneficial H-mode of operation.

Important to the dynamics of this system is how the edge plasma interacts with the wall. The fuel for fusion is hydrogen (actually, its isotopes, deuterium and tritium). The nuclei of these atoms can diffuse into the wall and be released at a later time. These processes are important to fusion confinement devices because they act as boundary conditions on the edge plasma, determining, in part, the influx of particles.

FACETS uses a component-based approach to couple the various physics. Currently, the core, edge, and wall equations are evolved separately and coupled with an explicit coupling scheme; preliminary work on implicit coupling is promising. The coupling between the core and the edge involves a transformation of the 2D edge data to a boundary condition for the 1D core equations, while a reverse transformation is required to transfer fluxes from the core to the edge. The edge-wall coupling is a pointwise coupling in which each cell on the edge boundary is coupled to a wall component.

Coupled whole-device simulations of tokamaks are required in order to understand and optimize fusion gain. The current approach to coupling is essentially code-to-code coupling between existing physics components that were not originally designed for this purpose. Hence, the challenge is not only to retrofit these components to work with one another but also to design robust and accurate algorithms that evolve the coupled physics. Further, developing new WDM components that take into account emerging high-performance architectures is a significant challenge that needs to be addressed in order to satisfy the modeling requirements as ITER and other tokamaks become operational.

2.1.6. Subsurface science

Problems in hydrogeology encompass a broad range of multiphysics and multiscale issues. In this section, we discuss several solution and coupling approaches in the area of subsurface fluid flow, solute transport, and reactions.

First, we consider continuum-scale subsurface reactive transport. For several decades researchers have studied the coupling of fluid flow, multicomponent solute transport, and geochemical reactions. In the 1980s, researchers [88–90] developed formulations for coupling multispecies chemistry with multicomponent transport. Yeh and Tripathi [91] and Steefel and MacQuarrie [92] provide overviews of these techniques, which can be divided into three general categories: (1) the global implicit approach (GIA), where solute transport and geochemical reaction are tightly coupled and solved simultaneously in a single system of nonlinear equations; (2) the sequential iteration approach (SIA), where solutes are transported independently, then coupled through geochemical reaction, with the process iterating until the global residual equations converge to a prescribed tolerance; and (3) the sequential non-iterative approach (SNIA), where the solutes are transported independently and coupled through geochemical reaction, but with the process not iterated to convergence. The SNIA approach is more commonly referred to as “operator splitting.” Steefel and MacQuarrie [92] present a concise overview of these three approaches and provide examples illustrating the tradeoffs of each.

Through the years, researchers have argued in favor of SNIA or SIA over the fully coupled GIA because of GIA’s excessive computing requirements [91]. However, both SIA and SNIA, with less coupling, present their own challenges regarding accuracy and performance [93–97]. SNIA is known to introduce operator-splitting error in the solution, and SIA may result in prohibitive iteration counts when geochemical species are tightly coupled to the immobile phase (e.g., strong sorption, mineral precipitation-dissolution with fast reaction kinetics). An interesting modification can be made to SIA by selectively coupling a subset of strongly coupled reacting species and solving the resulting subsystem of equations with GIA within the SIA framework [98]. This selective coupling can reduce SIA iteration counts by two orders of magnitude and model runtimes by more than an order of magnitude.

Recent advances in high-performance computing have enabled the use of GIA to solve three-dimensional problems comprising large numbers of chemical species. The feasibility of GIA has been demonstrated on a groundwater-flow and reactive-transport problem involving tens of millions of degrees of freedom simulated on thousands of processor cores [99]. With the use of increasing numbers of processor cores (e.g., 10K or more), simulating these larger nonlinear systems of equations requires the development of new solver algorithms to preserve nonlinear and linear iterative solver iteration counts and to minimize communication requirements. Physics-based preconditioners have been investigated for reactive transport [100], where the GIA Jacobian operator is split into two separate operators for global transport and local reaction. These two transport and reaction operators are then applied sequentially in the

preconditioning step, much like operator splitting is applied as a solution technique. This physics-based preconditioner performs well for subsurface problems composed of primarily aqueous reactions, but more poorly for scenarios where one would expect larger operator splitting error (e.g., fast reaction kinetics and solute retardation through sorption and mineral precipitation).

Next, we go beyond the continuum scale and consider multiscale approaches that couple pore-scale to continuum-scale models. At the pore scale, the geometry of solid grains and pore spaces is explicitly defined, and the physics represented are Navier-Stokes equations for single-phase fluid flow, advection and diffusion as transport processes, and various aqueous and heterogeneous reactions (both equilibrium and kinetic). Pore-scale models take various forms, including traditional mesh-based CFD methods, lattice-Boltzmann methods, pore network models, and fully Lagrangian smoothed particle hydrodynamics (SPH) particle methods.

In this multiscale approach, at the continuum scale, the system is treated as a “porous medium” without consideration of explicit grains and pores. Continuum-scale physics includes conservation of mass combined with Darcy’s law for single-phase fluid flow, mean advection and dispersion (including diffusion and effects of unresolved velocity variations) as transport processes, and a similar set of chemical reactions. Most continuum-scale models use traditional numerical methods (finite difference, finite element, finite volume) to discretize and solve the system of PDEs, as discussed above.

Since efforts in this treatment of multiscale-multiphysics coupling in subsurface reactive-transport problems have been initiated only recently, the literature contains only a small number of examples. These use both concurrent and hierarchical approaches. In concurrent methods, different physics are solved on separate subdomains with exchange of information and boundary condition matching at the boundaries between the subdomains (or in some cases on an overlapping region). Balhoff et al. [101] used a concurrent approach to couple a pore network model of flow within a sand-filled fracture to a continuum model of flow in the surrounding porous matrix. Tartakovsky et al. [102] have developed a coupled pore-continuum model of a diffusion/reaction system in which SPH is used at both scales; this allows straightforward model coupling at the domain boundary but is difficult to adapt to situations in which advective flow exists. Mortar methods [103] also provide a means of coupling subdomains with different physics by solving a boundary-matching problem on the mortar grid. Battiato et al. [104] used an iterative boundary-matching scheme to couple pore- and continuum-scale models of reactive transport including advective flow. Alternatively, hierarchical methods solve both models simultaneously over common domains and pass information between scales using various approaches. Tartakovsky and Scheibe [105] have successfully applied a hierarchical method called dimension reduction with numerical closure to simulation of mixing-controlled mineral precipitation. In this approach, conceptually similar to the heterogeneous multiscale method of E et al. [106], macroscopic properties and states are defined in terms of pore-scale variables, and short bursts of pore-scale simulation are used to provide numerical closure for a given set of local conditions. Then, a larger timestep is taken with the macroscopic model (eliminating a significant amount of pore-scale simulation) before it is necessary to update the numerical closure to reflect new macroscale conditions.

2.1.7. *Surface and subsurface hydrology*

A blueprint for modeling fully integrated surface, subsurface, and land-surface processes, originally proposed forty years ago [107], is now becoming a reality. Although truly coupled models have appeared only recently in the literature [108–113], a growing library of models and community of modelers now contribute considerably to our understanding of the coupled terrestrial hydrologic and energy cycles. These coupled and integrated hydrologic models are being used on a range of important issues. The models have common features in that they solve some form of the nonlinear diffusion equation (e.g., Richards) for subsurface water flow and some form of the shallow water equations (e.g., kinematic or diffusive wave) for surface flow. Additionally, these hydrologic models are being coupled to land-surface energy, biogeochemistry/ecology models and to atmospheric models [108, 109, 111, 113–116]. Updated numerical and computational technologies have, in part, enabled new approaches for modeling these coupled interactions. While these models all take different numerical, discretization, and coupling approaches, they share the common goal of modeling—rigorously and mathematically—the terrestrial hydrologic and energy cycle as an integrated system. Research addressing these issues encompasses a range of scales and includes a variety of processes [117]. Hydrologic simulation models break into two broad solution categories: a globally implicit formulation and operator-splitting or Gauss-Seidel methods. Among the globally implicit hydrologic models, there

are two distinct formulations for solution of surface and subsurface flow [117]: first-order exchange [108, 109] and continuity of pressure [111]. Integrated hydrologic models solve equations for surface flow (generally the shallow wave equations) and subsurface flow (generally Richards equation) simultaneously [108–113, 118–120] and represent both coupling approaches and a range of numerical implementation of solution (finite-element methods and finite-difference/finite-volume methods). While not widespread, integrated hydrologic models have been designed to run on large-scale computing infrastructure and, in one case, demonstrate very good weak scaling out to large numbers of processors [121].

ParFlow solves the variably saturated Richards equation in three dimensions. Overland flow and groundwater surface water interactions are represented through a free-surface overland flow boundary condition, which routes ponded water as via the kinematic wave equation [111]. The coupled subsurface-overland equations are solved simultaneously using a highly efficient and robust Newton-Krylov method with multigrid preconditioning, which exhibits excellent parallel scaling and allows simulation of large-scale, highly heterogeneous problems [122, 123]. Simultaneous solution ensures full interaction between surface and subsurface flows; this is particularly important in regions with spring-fed streams, large baseflows, and rivers that alternate between gaining and losing.

As one example of multiscale hydrological modeling, ParFlow solves the coupled water and energy balance at the land surface by incorporating a modified version of the Common Land Model (CLM3.0) [124], which calculates evaporation from vegetation canopy and the ground surface; transpiration and plant water uptake over the soil column; snow accumulation and melt; and latent, sensible, and ground heat fluxes. Vegetation and land-energy processes are calculated at each timestep by using an operator-splitting approach that ensures full mass and energy conservation [113, 114]. A fully coupled hydrologic and climate model PF.WRF was developed by combining ParFlow with WRF [116]. The general solution procedure begins with explicit advancement of the atmospheric code. Using the operator-splitting approach, surface fluxes that are relevant to ParFlow, such as infiltration and evaporation rates, are integrated within the atmospheric model over the ParFlow timestep. These accumulated surface fluxes are applied to ParFlow at the start of its implicit time advancement. The subsurface moisture field calculated by ParFlow is passed directly to the Noah land-surface model [125] within WRF for use in the next timestep. Noah is advanced for each WRF timestep to provide the necessary land-surface fluxes, but soil moisture values are now specified by ParFlow. ParFlow is incorporated as a subroutine in Noah, with communication over all Noah soil layers; subsurface hydrology in Noah is entirely replaced by ParFlow. The combined model is fully mass and energy conservative.

2.1.8. Climate modeling

Earth system modeling, specifically coupled global climate modeling (GCM), is one of the truly definitive multiphysics applications in several respects. GCMs encompass physical behavior ranging in spatial scales from molecular dynamics to global atmospheric wave patterns. Moreover, GCMs capture climate variations covering months to ten thousand years over regional to global areas of the Earth's surface. A number of GCMs are under active development around the world, and the Community Earth System Model (CESM), which has been developed by hundreds of scientists and contributes to the Intergovernmental Panel of Climate Change Assessment Report [126], is a good example to illustrate the multiphysics nature of global climate modeling. In most cases, the multiphysics characteristics of the CESM apply to other GCMs.

The CESM comprises five Earth system components—atmosphere, land surface, ocean, sea ice, and ice sheets—and has provision for future components, as shown in Figure 7. These components periodically exchange data as fluxes at component boundaries through a flux coupler, which coordinates the transfer of coupled physics across model components—except for ice sheets, which currently receive one-way surface-mass balance values from the land model. Typically, coupling occurs every simulated day, although recent high-resolution coupled model development has indicated the necessity of exchanging this information more frequently [127]. The source-model flux data is interpolated to the target component by using second-order or conserved mapping weights within the coupler. Coastal “masks,” which account for the boundary between land surface and ocean, are created as a data file to determine whether atmosphere model grid cells couple to the land surface or ocean below and vice versa. These masks need to be created for each combination of component grids used in a coupled-model run.

Within each model component, resolved physics are represented by discretized equations referred to as the dynamical cores, or dycores. Currently, all the physics and dynamics variables exist on the scale of the three-dimensional discretized grid, although the land model performs computations using statistical representations of subgrid-scale

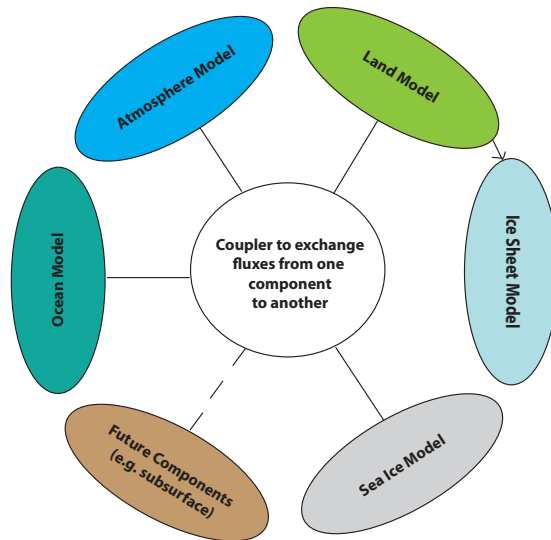


Figure 7: A schematic diagram illustrating the multi-fidelity structure of the earth-system model. Except for the ice sheet model, all components in the model mutually couple through the flux coupler, which is shown as the hub of the diagram. The ice sheet model does not couple in the same way. It is coupled only to the land model, receiving surface-mass balance values. As is shown, this coupling is one way—the land model receives no feedback from the ice sheet model.

land-use characteristics. The atmospheric component in CESM, the Community Atmosphere Model (CAM), has several dycore options that scientists can choose, based on desired attributes such as computational efficiency and local conservation properties. Each model component uses a different strategy for solving the model equations. For details, the reader is referred to the model documentation [128]. Unlike CAM, which has several choices for spatial discretization, the ice sheets, sea-ice, and ocean components in CESM use a finite-difference scheme on regularly spaced grids. Their time integration methods are diverse, arising from the different equations and physical scales being treated in each component. These methods include subcycled fully explicit (sea-ice), semi-implicit to “step-over” barotropic dynamics (ocean), and fully implicit velocity, with subcycled advection (ice-sheets) methods. Several of these components have developmental dycores using an unstructured grid to allow local refinement [129] and a testbed of fully implicit algorithms to overcome stability limitations [130].

As mentioned above, CAM has several dycore options for scientists to select when configuring the coupled climate model for experiments. These options include a spectral-element-discretized cubed-sphere configuration, which uses a subcycled fully explicit numerical scheme. Other dycore options include an Eulerian spectral and spherical finite-volume discretization, which both use an operator-splitting, semi-implicit time integration method to separate waves of different scales. In this case, a common splitting of advection and gravity waves adapted from the geophysical fluid dynamics community is employed because it takes advantage of a natural separation at coarse resolution and favorable numerical properties of the resulting equations. The atmospheric dynamical core resolves horizontal waves ranging from vertically propagating gravity waves (~ 200 m/s) to low-frequency global modes within the background flow ($\sim 2\text{--}20$ m/s). Vertical motions are smaller in scale and are represented by the hydrostatic approximation, which assumes that the vertical pressure gradient is largely balanced by gravity and is valid to scales down to individual cloud systems of about 100 km in extent.

Climate is uniquely challenging because the most important drivers of global climate change—chemistry, moisture, and radiation—vary over microscopic spatial scales but impact the global-scale net energy balance of the Earth system. The global-scale atmospheric flow is coupled to local chemical, cloud, and radiative forcings; and the time scale for this coupling ranges from the diurnal cycle to decades. Subgrid-scale effects are treated either as prescribed values from datasets or as separate, vertical-column, horizontally independent models that are updated as forcings to the dynamical equations. These independent models are themselves multiphysical. For example, the chemistry model in CESM [131] uses a split explicit-implicit solution method for the slow and fast species reactions, respectively.

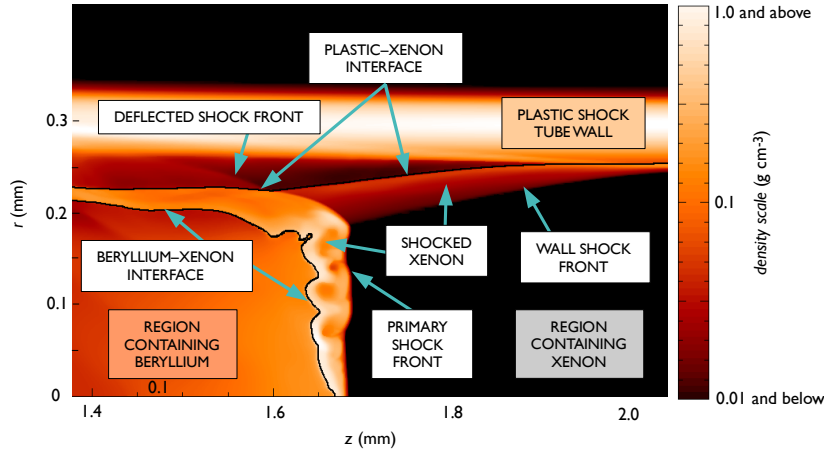


Figure 8: Illustration, adapted from Myra et al. [136], showing interacting shock waves propagating in a simulation of a high-energy-density physics shock-tube experiment. The figure shows a slice representing a half cross-section of the cylindrical shock tube with the central axis at the bottom of the figure. In addition to the shock fronts, a number of material interfaces are shown, which represent boundaries between the multiple materials used in the experiment (beryllium and xenon within a polyimide-plastic shock tube). The horizontal extent of this schematic represents roughly 0.6 mm of a shock tube that is about 4.0 mm in total length.

2.1.9. Radiation hydrodynamics

Physical processes that involve radiation or hydrodynamics individually are part of many phenomena of everyday life and everyday scientific research. However, phenomena where radiation and hydrodynamics couple and stand together on an equal footing—where strong coupling between the two components affects the evolution of each—rarely occur terrestrially but are common astrophysically. The energetics involved explains this apparent dichotomy. In astrophysics, radiation hydrodynamics plays a part in highly energetic phenomena, such as stellar core collapse, supernova shock propagation, supernova remnants, stellar pulsations, and accretion flow. Where radiation-hydrodynamic phenomena occur terrestrially, they are restricted to high-energy density environments: inertial-confinement fusion, laboratory astrophysics, and high-temperature shock-wave and explosive phenomena.

In addition to radiation hydrodynamics (RHD) being a multiphysics field, in that it strongly couples the evolution of distinct physical components (matter and radiation), many problems in RHD introduce other multiphysics aspects. For example, a single problem can involve any number of the four fundamental interactions across multiphysics domains and can mix classical- and modern-physics paradigms (e.g., classical or relativistic space-time and gravity, quantum statistics, electroweak, and strong nuclear interactions).

Numerical RHD models are typified by a number of common features [132–134]: Conservation PDEs play a fundamental role (e.g., mass, species, energy, momentum). As a result of the energies involved, shock waves are frequently present in such systems, leading to discontinuities in physical quantities. Additionally, material interfaces, which are part of many models of terrestrial setups, introduce another form of discontinuity into the problem. Some RHD problems can be classified as reactive-flow problems, as evidenced by complex and strong coupling terms between components. Figure 8 shows an example of the complexity of these systems: Interacting shock waves are propagating in a simulation of a terrestrial shock-tube experiment as part of the CRASH project [135]. A beryllium foil is driven into a xenon-filled shock tube by a drive laser to speeds in excess of 100 km/s. This process generates a radiative primary shock wave, which in turn generates a wall shock as the plastic walls of the shock tube are heated and ablated.

Other complexities also are common in RHD models. Frequently, the equations for the radiation are nonlinear, requiring solution by nested iterative procedures, such as preconditioned Newton–Krylov methods. (These nonlinearities can be caused by strong coupling between matter and radiation as well as by quantum-mechanical degeneracy effects.) Coupling between physical components can be a sensitive issue because, in complex phenomena, coupling varies in strength as a function of time, position, and/or energy and thus can require different solution approaches in different situations. High dimensionality is also an issue. Proper treatment of the radiation requires solution of the

Boltzmann transport equation over six-dimensional phase space as a function of time. Almost always, finite computer resources force a compromise, leading to reduced dimensionality in some fashion [137]. Such compromises can affect self-consistency of coupling (e.g., radiation–matter momentum exchange in flux-limited diffusion approximations).

Since numerical solution of RHD problems predates the computer age, solution techniques are many and varied. Today, however, solutions typically involve the solution of Euler’s equations for the matter component. (Reynolds numbers are typically very high in such systems.) Standard methods include finite-difference and finite-volume methods. Lagrangian, Eulerian, and hybrid approaches (e.g., arbitrary Lagrangian–Eulerian [ALE] methods) are used. Such techniques also translate to solution of the radiation [138–143]. Here, however, finite-element methods also find use. [144].

When a similar discretization is used for both matter and radiation, coupling between components reduces to a problem of obtaining consistency and accuracy of paired coupling terms. However, when different discretizations are used for matter and radiation (e.g., finite volume for matter and finite element for radiation), then accuracy of the transfer operator can become problematic. Unlike frequently encountered multiphysics couplings discussed in this report, matter–radiation coupling occurs, in some strength, over the entire physical domain of interest; it is a volumetric coupling. Methods for improving the accuracy of this coupling are an area of active research.

RHD models commonly use a mixture of explicit and implicit approaches within the context of operator splitting. Choices stem from the timescales involved. Typically, hydrodynamic timescales are much greater than radiative timescales. This situation leads naturally to use of explicit methods for the matter and implicit methods for the radiation (usually sparse iterative approaches for the latter). However, for problems with long evolutionary timescales (e.g., neutrino emission for cooling neutron stars [145]), implicit approaches increasingly are being used for both components.

The difficulties inherent in adequate treatment of multiphysics coupling in RHD models are illustrated by a sample list of the myriad of relevant timescales involved: acoustic, radiation-transport and radiation-source, electron–ion equilibration, electron diffusion, and viscous (if present). Arrayed in relative size, these may span a range of 10^{15} . A challenge for the application modeler is to determine which of the associated processes are to be followed in detail and which can be assumed to occur so quickly that an equilibrium approximation is adequate.

2.1.10. *Geodynamics and magma dynamics*

Geodynamics is the study of dynamic processes in the Earth from crust to core. Processes range from earthquakes to plate motions, large-scale convection to microscale melt migration. Physical coupling occurs over a wide range of simulations in geodynamics. Short-term crustal behavior is generally modeled with some combination of a porous, viscous, elastic, or plastic material model [146]. Beneath the crust, incompressible mantle flow calculations couple the advection-diffusion of energy/temperature to the Stoke’s equations [147]. The complexity may be increased by introducing the advection of chemical components with low or zero diffusivity, which is often achieved by using particle-based methods [148]. Compressible flow, which normally uses an anelastic approximation to filter out seismic waves and timescales, may also introduce reference equations of state derived from models of seismic wave propagation [149]. These simulations are often multiscale, including narrow plume conduits in large domains [150, 151], along with physical discontinuities in material properties at a variety of depths [152]. At greater depths, in the outer core, models couple turbulent flow from the Navier-Stokes equations to thermal and magnetic models of rapidly convecting liquid iron [153, 154].

The change of behavior and physical timescales with depth generally leads to different models being used for different shells of the Earth. Boundary conditions are then used to parameterize the fluxes between models. As an example, owing to high convective vigor in the outer core, the lower thermal boundary condition on the mantle may be treated as a Dirichlet condition [156], while the motions of crustal plates may be treated as a rigid shell, with prescribed shear at the upper boundary of the mantle [157]. Depending on the application, this outer boundary may also be treated as a free surface, supporting dynamic topography, which involves coupling the solution of a kinematic boundary condition to the flow solution in the global domain. This approach potentially introduces instabilities in incompressible flows [155], as shown in Figure 9. Smaller-scale, higher-resolution studies may occur on overlapping domains. As an example, one could focus on the production of melt in subcrustal regions, such as mid-ocean ridges or subduction zones. The differing temporal and spatial scales of pore-filling fluid and melt lead to calculations that are often multiphase, coupling Darcy’s flow with Stokes flow in the mantle or poro-visco-elastic flow in the crust [158–160]. Experimental, theoretical, and numerical work attempts to link these high-resolution studies to larger domains




Figure 9: The “drunken sailor” instability. In a simple case of a compositionally (grey shading) driven Rayleigh-Taylor instability, an explicit coupling between the incompressible Stokes equations and the kinematic free-surface on the top boundary leads to a sloshing instability in the velocity field (vectors) [155].

through parameterizations of rheology and melting [161].

2.1.11. Particle-accelerator design

A key component technology in particle accelerators is the accelerator cavity, which is used to generate electromagnetic fields that control acceleration of the subatomic particles in the accelerator beam. To achieve high efficiency, a modern accelerator cavity uses superconducting devices to generate the fields that drive the cavity at its resonant frequency.

In radio-frequency (RF) cavity design, finding a sufficiently accurate resonant frequency for real geometry has become possible with the help of a higher-order, high-fidelity finite-element method [162–165]. One of the main sources of operational failures of RF cavities is strongly related to the excessive thermal heating arising from high power or high currents in operation [166, 167]. As the power and current increase for future experiments, optimizing the performance and cost effectiveness of RF cavities becomes an increasingly important technological challenge for designing next-generation accelerators. TEM3P, a multiphysics analysis tool using finite-element methods, has been developed to capture thermal, structural, and electromagnetic effects such as cavity wall heating [168], structural deformations [169], and Lorentz force detuning [170]. Parallel implementation of TEM3P utilizes supercomputers for large-scale computations [171–173] with high geometric fidelity and high numerical accuracy. TEM3P solves the thermoelastic equation, in which the heat flux and radiation pressure are interpolated from the electromagnetic solution file from Omega3P (a parallel finite-element code for frequency domain analysis of accelerator cavities) [174].

The design process starts with accurate electromagnetic solutions obtained from Omega3P analysis. Omega3P solves Maxwell’s equations with higher-order curved elements for highly accurate solutions. TEM3P reads the electromagnetic solution file (in NCDF format), which is interpolated to produce accurate physical quantities of interest for the given multiphysics application. Heat flux is calculated for the thermal solver in TEM3P [168], while radiation pressure is calculated for the structural solver in TEM3P [169]. Interpolation of the electromagnetic fields accounts for the general case, in which the nodal points in the TEM3P mesh do not coincide with the nodal points in the Omega3P mesh. Once the desired heat flux or radiation pressure is obtained, it is applied as a boundary condition to the existing thermal model for an RF heating problem or to the existing structural model for a Lorentz Force detuning problem.

Internally, TEM3P couples its thermal solution to its structural solver. When a temperature distribution is known, thermal stresses due to the high-temperature gradient can be readily calculated from the structural solver. As shown in Figure 10(b), the high-temperature gradient in the RF gun is strong enough to cause geometric distortions so that the resulting frequency shifts can be larger than the operating frequency bandwidth [169].

Commonly, in multiphysics analysis, the primary source of heat flux is RF heating. In this case, an electro-thermo-structural coupling, as shown in the figure, is used to provide viable solutions that address frequency shift problems.




Figure 10: Multiphysics coupling analysis of LCLS gun.

TEM3P has also been applied to thermal instability problems in superconducting cavities [168]. Thermal instability can lead to serious problems such as thermal quenching, which terminates the superconducting state of a cavity [175]. In the face of this instability, operation is usually not possible (or only at a significantly reduced efficiency if it is possible) [176]. Superconductor materials exhibit highly nonlinear thermal characteristics, and nonlinear boundary conditions need to be applied in order to obtain an accurate solution. A robust nonlinear solver is implemented in TEM3P to address these problems arising from the nonlinear material properties.

The current version of TEM3P provides steady-state solutions for both thermal and structural problems. Multiphysics problems with multiple spatial and temporal scales are not yet addressed in TEM3P. As designs become more concerned with fully time-dependent multiphysics coupling, further research and development become increasingly important.

Shape optimization studies have been performed with adjoint gradient information (inverse eigenvalue problems) [177, 178]. Because adjoint variables enable gradient calculation at a given design point regardless of the number of design variables, achieving a near-optimal shape in a reasonable turnaround time has become feasible [179–181]. With the successful implementations of multiphysics coupling, applying its solution to optimization problems could open up new research areas. The solutions from the multiphysics coupling could provide insights in the complex physical interactions not resolvable from single-physics analysis. Thus, combining an adjoint-based gradient optimization with multiphysics coupling could potentially play an important role in achieving near-optimal solutions in multiphysics environments.

2.2. Crosscutting Issues in Multiphysics Applications

The previous subsection contains a diverse collection of multiphysics applications, including brief descriptions of solution approaches, treatment of multiphysics coupling, and various application-specific issues. Here, we extract commonalities and differences from those descriptions. Our goal is twofold: (1) to identify issues in algorithms and software that need to be addressed through, perhaps, different approaches and more rigorous analysis, and (2) to provide pointers to related discussion in the remainder of the document. Many of the implementation issues in multiphysics applications can be classified as relating to either splitting or coupling of physical processes among or between solution modules. We consider both spatial and temporal coupling issues. We also introduce software engineering questions that touch on best practices issues as well as the cultural environment for software collaboration and the hard choices that must be made when mapping an application to specific hardware.

Many multiphysics applications owe their coupling to historical accident. Multiphysics applications have rarely been designed as a first creation. Instead, such applications typically represent the product of collaboration between two scientific groups of diverse fields who find a common interest that merits joint investigation. The result is a marriage of two diverse codes (or algorithms); common practices from a software perspective are discussed in Section 4.1.1. The driving force is the new capability that this union makes available; the drive toward an answer trumps the drive toward an *accurate* answer or a computationally efficient simulation. In spite of valuable scientific advances that may be made, time is rarely made for mathematical rigor. Thus, it is common that, for example, two well-constructed second-order accurate codes are coupled by a quick *ad hoc* scheme that may be, at best, first order. In this

report, our goal is to suggest approaches for more formal analysis of spatial and temporal coupling methods in such areas as accuracy and order of convergence. Of special importance are interfaces between subdomains of a problem, volumetric interfaces, and temporal interfaces.

2.2.1. Choices and challenges in coupling algorithms

Spatial interfaces remain a primary challenge, especially between subdomains using different computational approaches and discretizations (e.g., an interface between two subdomains—one discretized by finite elements and the other by finite volumes). Many commonly used interpolation approaches are noncommutative and thus can lead to nonconservation of supposedly conserved quantities. A number of applications we have described use mortar methods to bridge this difficulty (see Sections 2.1.1, 2.1.2, and 2.1.6). Methods for coupling multiphysics components in space are discussed in Section 3.2.1.

A survey of applications suggests that explicit methods are still prevalent across many fields, especially for “production” applications. This situation is understandable: compared with implicit methods, explicit methods are easier to implement, maintain, and enhance from a software-engineering perspective, as discussed in Section 4.1.1. Maintaining the stability of an explicit method mandates that the programmer understand at least some the important timescales involved in the problem. Implicit methods are free of splitting errors that are introduced by operator splitting, and may allow the use of much larger timesteps than an explicit method. However, when using an implicit method, time steps must be carefully selected to control at least local time integration error in order to avoid “silently” producing inaccurate answers by potentially taking large timesteps over important processes having fast timescales. This is of special concern when, as occurs frequently, the equations being solved are nonlinear. (Nonlinearities can occur in nearly all the application examples above, but notably those described in Sections 2.1.1, 2.1.3, 2.1.5, 2.1.6, 2.1.7, 2.1.8, 2.1.9, and 2.1.11.) In summary, the explicit-implicit debate becomes one of accuracy versus efficiency. Given the increasing use of implicit methods, it is necessary to ask how we should approach judging the accuracy of a method. For example, it is easy to integrate $y'(t) = y(t)^2$ from $y(0) = 1$ in a stable way that completely misses the singularity at $t = 1$ and continues smoothly with unbounded error into $t > 1$. How do we know we are not skipping over physics? What is the collection of metrics we can use to provide confidence in our answers? Which integrated metrics are of greatest value? Which pointwise metrics can be relied on?

A number of the applications discussed in this report have timescales for important physical processes that may vary by many orders of magnitude. (See Sections 2.1.4, 2.1.8, 2.1.9, and 2.1.10.) The challenge for the model builder is to have a clear picture of what quantities need to be tracked in the problem being solved. Such insight helps to determine which of the associated processes need to be followed in detail and which can be assumed to occur too quickly for their evolutionary details to be important. An additional challenge is to write that intelligence into the code itself. Section 3.2.2 discusses methods for coupling multiphysics components in time, including implicit-explicit (IMEX) and adaptive techniques, which can help address such multiphysics challenges in temporal coupling. Additional opportunities are considered in Section 5.

Since diverse timescales are common in multiphysics applications, it is worth asking whether an approach using uncertainty quantification (UQ) methods may yield insight into the relation between timestep selection and results obtained; UQ issues are discussed in Section 3.5.

A related issue involves quantifying certain cultural knowledge that allows one to “cheat” a little on constraints when quantities of interest are not affected. An example is the careful violation of the Courant-Friedrichs-Levy (CFL) condition in explicit codes.

Another interesting challenge is spatial coupling of the discrete to the continuum. This issue arises in problems in which the macroscopic evolution of the system relies on detailed knowledge of phenomena occurring on a microscopic level. Examples include the modeling of crack propagation and the ultrafast sequencing of DNA discussed in Section 2.1.4. An additional class of problems in this vein is nonequilibrium fluid flow: adequate description of state variables in the fluid may rely upon, for example, the result of a simulation at the molecular, atomic, or even electronic or nuclear scale. Problems in radiation hydrodynamics (see Section 2.1.9) can fall into this class.

Discrete-to-continuum coupling may also arise at interfaces between diverse materials where, in one material, a microscopic model is required for a satisfactory physical description, while in the other material, a classical continuum description is adequate. Whenever these sorts of couplings arise, a one or more additional mesoscale descriptions that link the two extremes has found use. (Section 2.1.4 describes applications that benefit from this approach.) We discuss discrete-continuum coupling more fully in Section 3.3.

Any need for adaptivity in algorithms throughout a given multiphysics simulation also presents challenges. As an example, consider a multiphysics application in which a key piece of science changes in the middle of the model's evolution (e.g., a rate may become much larger or smaller as a result of some change in a linked quantity). This situation often leads to the component equations having a changed character. For example, what might have previously been a parabolic problem might now take on hyperbolic behavior. This circumstance may require a change of solution methods (introduced in Section 3.1), anywhere from the preconditioner to the whole solver approach. Equally extreme are phase transitions: solution methods are likely to be different for a subdomain now containing a solid phase instead of a gaseous phase. These issues, as well as related challenges in supporting changes in modeling and algorithms over the lifetimes of multiphysics applications, are considered from a software-design perspective in Section 4.

2.2.2. *Software engineering*

The current status of multiphysics software is discussed in Section 4.1, including current practices for multiphysics coupling in production codes, common needs, and examples of libraries, frameworks, and tools that support multiphysics simulations in high-performance computing environments. The development of large-scale multiphysics applications raises many of the same issues as does development of any large-scale application. One of the most fundamental issues is the “buy” versus “build” choice. In the multiphysics world, however, this question takes on increased importance because of the added application complexity. The “build” approach frequently requires large collaborations of diverse expertise. However, the sources where one may “buy” are often similarly scattered among disciplines and institutions. Neither approach alone guarantees robust, accurate code. Risk mitigation also plays a role in this choice. On the one hand, third-party code is not designed with the end user's specific purpose in mind. Learning the code (and perhaps learning that the code is not suitable) can be time consuming. Even if such a third-party code is suitable, additional investment is required to maintain multiple, interacting software packages, perhaps across multiple platforms. A decision to “build” can be equally risky. In addition to the obvious issues of funding, assembling, and directing a collaborative team, and then having the team deliver on time, there is an ever-present danger that the drive for science leads to compromises in code accuracy, especially where the subtle coupling of multiphysics models occurs. Challenges, near-term needs, and risks in multiphysics software are discussed in Sections 4.2 and 4.3, while longer-term opportunities and insertion paths are considered in Sections 5 and 6.

2.2.3. *Analysis and verification*

Analysis and verification of coupled models involve issues that span throughout both algorithms and software and have been only partially addressed by the community overall. Section 3.4 discusses error estimation and uncertainty quantification for multiphysics problems from an applied mathematics perspective, while Section 4.2 considers issues in software testing. Verification of coupled codes is challenging. Unit tests do not test the synergistic effects of coupling—the very aspect of simulation that probably requires the greatest testing. In addition, large multiphysics codes have too many coupling combinations to allow practical coverage of the code space by any conceivable test suite. Furthermore, standard reference problems are lacking. Given this situation, verification tests tend to devolve into regression tests, whose value further devolves when existing physics is changed or new physics introduced. Appendix B is a first step at presenting a few multiphysics exemplars in a unified format.

3. Algorithms for Multiphysics Coupling

In this section, we discuss methods for formulating and solving multiphysics problems. The formulation of a multiphysics problem is highly dependent on the nature of the problem and how the component physics fields are coupled in space and time. In some circumstances, coupling between continuum and discrete models is required. Examples include heat conduction problems that require calculation of temperature-dependent thermal conductivity from first principles and plasticity problems where empirical stress-strain relationships are not available.

This section also addresses error estimation and uncertainty quantification for multiphysics problems. The complexity of multiphysics problems makes assessment of errors and uncertainties both problematic and essential. We note that in practical terms, it is often difficult or impossible to obtain solutions of multiscale, multiphysics models that are asymptotically accurate uniformly in space and time. Hence, the fruits of classical numerical analysis may be unobtainable in the study of multiphysics models. Further, while uncertainty quantification is a relatively young field with many foundational research problems remaining to be tackled, uncertainty quantification for multiphysics models is simply nascent.

We begin with a general discussion of solver methods in Section 3.1. These algorithms form the building blocks of many multiphysics codes that run on the largest currently available high-performance computing systems. This is followed by a discussion of strategies for coupling physics components in space in Section 3.2.1 and in time in Section 3.2.2. The special case of coupling continuum and discrete models is discussed in Section 3.3. We then discuss error estimation in Section 3.4 and uncertainty quantification in Section 3.5.

3.1. Solver Methods

While multiphysics problems (1a–1b) and (2a–2b) are inevitably nonlinear, most solution methods require solving a series of linear subproblems. Here we briefly summarize currently available techniques for solving systems of linear equations, followed by a review of methods for nonlinear problems. Related issues in software are discussed in Section 4.

3.1.1. Methods for systems of linear equations

Consider the system of linear equations

$$Au = f, \quad A \in \mathbb{R}^{N_d \times N_d}, \quad u, f \in \mathbb{R}^{N_d}, \quad (7)$$

where N_d is the number of degrees of freedom. Equation (7) can represent an entire multiphysics problem or some subproblem encountered in a multiphysics solution algorithm. Methods for solving (7) fall into two categories: direct and iterative methods. Direct methods generally rely on some canonical decomposition of A and scale poorly (in both storage and operations) when applied to large problems, even when sparsity in A is effectively exploited. For many years, researchers relied on *stationary* iterative methods, such as classical Jacobi, Gauss-Seidel, or successive overrelaxation methods. While effectively alleviating the large storage requirements of direct methods and being easy to implement, stationary methods lack robustness and generally do not converge fast enough to be considered practical for realistic multiphysics problems. Current best practices for efficient solution of (7) are based almost exclusively on multigrid methods and preconditioned Krylov subspace methods.

Multigrid methods. Systems of linear equations (7) usually must be solved many times in a multiphysics application. Consequently, linear solvers are often bottlenecks in multiphysics applications. This situation places a premium on algorithms that have optimal arithmetic complexity; ideally, solutions to (7) should be computed with $\mathcal{O}(N_d)$ floating-point operations. Multigrid methods [182–184] achieve this objective for a large class of problems.

Multigrid methods are based on the observation that different components of error can be more efficiently represented and eliminated on grids of different resolution. This idea can then be applied recursively to develop solution algorithms that utilize a grid hierarchy to achieve high efficiency. The key ingredients are an inexpensive smoothing method to apply on each level in the grid hierarchy (typically a stationary method), transfer operators to move data between successive grid levels, and a schedule for visiting each grid level in the hierarchy, resulting in the well-known V-cycle, W-cycle, or full multigrid cycle.

Developed in the context of elliptic and parabolic PDEs that typically enjoy smooth solutions, multigrid methods have demonstrated near-optimal $\mathcal{O}(N_d \log N_d)$ scaling on such problems on up to hundreds of thousands of processors

on modern parallel architectures [185]. Variations for anisotropic problems and problems with discontinuities have also been developed. While the early development of multigrid methods focused on constructing coarser problems through reference to the underlying grid, the fundamental properties of smoothing and grid transfer operations were soon abstracted and combined with coarsening strategies that are adaptively based on analysis of the matrix entries. This research led to the class of *algebraic* multigrid methods, which offer the potential for near-optimal black-box solvers for large-scale problems without any reference to an underlying mesh structure. Current research in multigrid methods strives to improve the scalability of coarsening strategies [186, 187]. High-performance implementations of both geometric and algebraic multigrid methods are available in numerous software libraries [11, 188].

Krylov subspace methods. Let $u_0 \in \mathbb{R}^{N_d}$ be an initial approximation to the solution of (7), $r_0 = f - Au_0$ the initial residual, and for each integer $k \geq 1$ let

$$\mathcal{K}_k(r_0, A) = \text{span}\{r_0, Ar_0, \dots, A^{k-1}r_0\}$$

be the *Krylov subspace* generated by r_0 and A . Consider first the case where A in (7) is symmetric and positive definite. In this case, A induces an inner product $\langle u, v \rangle_A = u^\top Av$ and corresponding proper norm $\|u\|_A = \langle u, u \rangle_A$. For symmetric positive definite A , the *conjugate gradient* (CG) method [189] generates a sequence $u_k \in u_0 + \mathcal{K}_k(r, A)$ of approximations to the solution u^* of (7) such that $\|u_k - u^*\|_A$ is minimized over $u_0 + \mathcal{K}_k(r, A)$ and the residuals $r_k = f - Au_k$ are orthogonal to $\mathcal{K}_k(r, A)$. Because of these optimality properties, CG is guaranteed to converge in at most N_d iterations in exact arithmetic. Remarkably, these optimality properties can be achieved by using simple, three-term recurrence relations and a small number of auxiliary vectors. Hence, CG was initially seen as a direct method with low storage that could compete with factorization-based direct solution methods. It was soon realized, however, that the key orthogonality relationships could not be maintained with finite precision, and interest in CG waned until the early 1970s. At that time, it was pointed out [190] that CG could still be used effectively as an iterative method. Ideas for accelerating convergence using preconditioning appeared soon afterward [191, 192]. See [193] for an early history and [194, 195] for more complete recent overviews of Krylov subspace methods.

When paired with a suitable preconditioner, CG is a robust and efficient method for solving (7) when A is symmetric positive definite. In multiphysics problems, however, few circumstances are encountered where A is symmetric positive definite. Numerous generalizations of CG to nonsymmetric problems that seek to exploit the optimality properties of CG have been proposed. Two principal approaches have been taken. One is to maintain $r_k \perp \mathcal{K}_k(r_0, A)$ through an explicit Gram-Schmidt orthogonalization process. This approach constructs an orthonormal basis of $\mathcal{K}_k(r_0, A)$ and leads to the Arnoldi family of iterative methods, of which the generalized minimum residual (GMRES) method [196] is the best-known variant. The main disadvantages of this approach are the storage requirements for the basis of $\mathcal{K}_k(r_0, A)$ and growth in the operation count for incrementally constructing that basis. Practical approaches to dealing with this difficulty amount to setting a maximum dimension of the Krylov subspace. Alternatively, a three-term recurrence relationship can be used. This approach is unable to maintain orthogonality, but it is inexpensive and can often be effective when paired with a suitable preconditioner. This strategy leads to the Lanczos-based family of methods. BiCGStab [197] is perhaps the most widely used method in this class. For general nonsymmetric A , maintaining orthogonality while using short recurrence relationships is not possible [198]. High-performance implementations of Krylov subspace methods are widely available [8, 11, 188].

Preconditioning. The key observations for Krylov subspace methods are that $\{\mathcal{K}_k(r_0, A)\}_{k=1}^{N_d}$ is a nested sequence of subspaces that approximates the range of A and that the efficiency of a Krylov subspace method is determined by how well (7) is represented by its restriction to a low-dimensional Krylov subspace. *Preconditioning* is a technique for transforming (7) to an equivalent system where a lower-dimensional Krylov subspace provides a better approximation. This amounts to the introduction of a preconditioner $P \approx A^{-1}$; and typical formulations include left, right, and split preconditioning:

$$\begin{aligned} PAu &= -Pf, && \text{(left),} \\ AP^{-1}u &= f, && \text{(right),} \\ P_1AP_2P_2^{-1}u &= -P_1f, && \text{(split).} \end{aligned}$$

Observe, in particular, that for left preconditioning, the sequence of approximations $\{u_k\}$ now resides in the Krylov subspaces $\{\mathcal{K}_k(Pr_0, PA)\}_{k=1}^{N_d}$. To illustrate with an extreme case, if $P = A^{-1}$, then $Pr_0 = u^* - u_0$ (i.e., the initial error), and the Krylov subspace method converges in one iteration. Thus, the art in selecting a preconditioner is managing the tradeoff between the cost of applying the preconditioner and the reduction in the number of required iterations.

Many strategies are available for constructing a preconditioner. Early efforts were based on the notion of *incomplete factorizations* [199], and much research into black-box, algebraic preconditioning techniques is built around this idea. Using a stationary iterative method as a preconditioner was first explored in [191], and preconditioning with a multigrid method was first explored in [200].

Parallel computing introduces new requirements for preconditioners that have a high ratio of computation to communication. A popular preconditioning approach is to use sparse direct solvers or incomplete factorizations in serial on each parallel process, in which each process solves only the portion of the overall problem that resides on its subdomain. Since these methods follow from the natural parallel decomposition of a problem into tasks, they are referred to as *domain decomposition* methods [201, 202]. Such solvers make good reuse of existing serial solver infrastructure, can attain excellent flop performance, work well for problems posed on unstructured meshes, and naturally couple multiphysics processes tightly within each solve.

Consider the full computational domain Ω as being given by a set of p nonoverlapping subdomains $\Omega_1, \dots, \Omega_p$; that is, $\Omega = \cup_{i=1}^p \Omega_i$. Then we can partition the unknowns of (7) into two subsets: the first consists of unknowns belonging to a single subdomain, u_I , and the second consists of unknowns that lie on inter-subdomain boundaries and couple subdomains together, u_Γ . Applying this partition to the matrix A as well results in a block linear system of the form

$$\begin{bmatrix} A_{I,I} & A_{I,\Gamma} \\ A_{\Gamma,I} & A_{\Gamma,\Gamma} \end{bmatrix} \begin{bmatrix} u_I \\ u_\Gamma \end{bmatrix} = \begin{bmatrix} f_I \\ f_\Gamma \end{bmatrix}, \quad (8)$$

where the matrix $A_{I,I}$ is block-diagonal. We further define restriction operators R_i that take a full set of unknowns, u , and return only the set of unknowns on the interior of Ω_i for each $i = 1, \dots, p$.

The simplest domain decomposition method is *block Jacobi*, where all the interprocessor couplings are ignored—that is, Γ is empty—and (8) is solved by using only the decoupled subdomain solves,

$$u^{k+1} = u^k + \left(\sum_{i=1}^p R_i^T A_i^{-1} R_i \right) (f - Au^k), \quad (9)$$

where A_i is block of $A_{I,I}$ belonging to Ω_i . A block Gauss-Seidel strategy, which sacrifices a lower degree of parallelism for faster convergence, can also be formulated, provided some coloring of the subdomains $\{\Omega_i\}_{i=1}^p$ is available. For problems with indefinite or singular A_i , artificial Dirichlet or Neumann boundary conditions are typically employed at interprocessor interfaces. Because of the resulting decoupling between each subdomain, these methods do not perform well at large scales, especially for elliptic and parabolic PDE problems. As a result, various methods have been proposed to increase the coupling between subdomains and attain improved scaling performance.

Additive Schwarz methods [201, 203–205] increase coupling by extending each subdomain to overlap with neighboring subdomains, increasing the size of each A_i matrix and redefining R_i appropriately, then using (9). Extensions of this method include modification of the prolongation matrix R_i^T [206] and inclusion of coarse-grid solves [207]. A multiplicative variant, using the same components of the block Jacobi strategy above, is also possible.

Other classes of domain decomposition methods (e.g., *FETI* [208], *BDD* [209], *BDDC* [210]) instead increase subdomain coupling by using a nonempty interface set Γ and following a Schur complement–like formulation to solve the full problem

$$u^{k+1} = u^k + \begin{bmatrix} I & -A_{I,I}^{-1}A_{I,\Gamma} \\ 0 & I \end{bmatrix} \begin{bmatrix} A_{I,I}^{-1} & 0 \\ 0 & S^{-1} \end{bmatrix} \begin{bmatrix} I & 0 \\ -A_{\Gamma,I}A_{I,I}^{-1} & I \end{bmatrix} (f - Au^k),$$

where $S = A_{\Gamma,\Gamma} - A_{\Gamma,I}A_{I,I}^{-1}A_{I,\Gamma}$ is the Schur complement matrix. Primary differences in methods following this approach lie in their choice of artificial boundary conditions at interior subdomain interfaces and how much of Γ they include, typically choosing only a small subset of the interface problem within S^{-1} at each iteration. Extensive support for constructing domain-decomposed preconditioners is available in high-performance libraries [8, 11].

3.1.2. Methods for systems of nonlinear equations

Perhaps the simplest and most robust approach for solving systems of nonlinear equations (3) is the fixed-point iteration, also known as the Picard or nonlinear Richardson iteration. These methods are performed through reformulation of the root-finding problem (3) into a fixed-point problem $u = G(u)$, through definition of a fixed-point iteration function, for example,

$$G(u) := u - \alpha F(u), \quad (10)$$

where $\alpha > 0$ is a fixed-point damping parameter that is typically chosen to be less than 1. Fixed-point methods then proceed through the iteration

$$u^{k+1} = G(u^k), \quad (11)$$

with the goal that $\|u^{k+1} - u^k\| < \epsilon$. A particular attraction of the fixed-point iteration is that if the iteration function is a contraction, that is, if there exists some $\gamma \in (0, 1)$ such that

$$\|G(u) - G(v)\| \leq \gamma \|u - v\|,$$

for all vectors u and v in a closed set containing the fixed-point solution u^* , the fixed-point iteration is guaranteed to converge based on the Banach fixed-point theorem (also called the contraction mapping theorem [211, 212]). This convergence is typically linear, however, and can result in slow convergence even when starting with a good initial guess at the solution.

Fixed-point iterations can converge slowly, in particular when the contraction constant γ is close to one. Newton's method (Algorithm 3) offers faster (up to quadratic) convergence. Direct computation of δu in Algorithm 3 may be prohibitively expensive for large-scale problems. *Inexact* Newton methods [213] generalize Algorithm 3 by allowing computation of δu with an iterative method, requiring only that $\|J(u^{k-1})\delta u + F(u^{k-1})\| < \epsilon_k$ for some carefully selected set of tolerances $\{\epsilon_k\}$; see [214, 215] for theoretical details. Newton-Krylov methods are important variants of inexact Newton methods in which δu is computed with a Krylov subspace method. This choice is advantageous because the only information required by the Krylov subspace method is a method for computing Jacobian-vector products $J(u^k)v$. A key consequence of this reliance on only matrix-vector products is that if these directional derivatives may be approximated, the Jacobian matrix $J(u^k)$ is never itself needed. The Jacobian-free Newton-Krylov (JFNK) method exploits this approach, using a finite-difference approximation to these products

$$J(u^k)v \approx \frac{F(u^k + \sigma v) - F(u^k)}{\sigma}, \quad (12)$$

where σ is a carefully chosen differencing parameter [216, 217] and F is sufficiently smooth. This important specialization facilitates use of an inexact Newton method by eliminating the need to identify and implement the Jacobian. Of course, efficiency of JFNK depends critically on preconditioning the inner Krylov subspace method, and the need to track changes in the Jacobian places a premium on preconditioners with low setup cost; see [5, 218–225]. Implementations of inexact Newton methods are available in several high-performance software libraries [8, 9, 11].

Fixed-point and inexact Newton methods can be used to solve multiphysics problems in a fully coupled manner, but they can also be used to implement multiphysics coupling strategies such as Algorithms 1 and 2. With an implicit method for one of the components, one can directly eliminate it in a linear or nonlinear Schur complement formulation. The direct elimination process for u_1 in the first equation of the equilibrium system (1a), given u_2 , can be symbolically denoted

$$u_1 = G(u_2),$$

with which the second equation (1b) is well defined as

$$F_2(G(u_2), u_2) = 0,$$

and each iteration thereof requires subiterations to solve (in principle to a high tolerance) the first equation. Unless the first system is much smaller or easier than the second (a special case of merit), this is not likely to be an efficient algorithm, but it may have robustness advantages.

If the problem is linear,

$$\begin{aligned} F_1(u_1, u_2) &= f_1 - A_{11}u_1 - A_{12}u_2 \\ F_2(u_1, u_2) &= f_2 - A_{21}u_1 - A_{22}u_2, \end{aligned} \quad (13)$$

then $F_2(G(u_2), u_2) = 0$ involves the traditional Schur complement

$$S = A_{22} - A_{21}A_{11}^{-1}A_{12}.$$

If the problem is nonlinear and if Newton's method (see, e.g., [217]) is used in the outer iteration, the Jacobian

$$\frac{dF_2}{du_2} = \frac{\partial F_2}{\partial u_1} \frac{\partial G}{\partial u_2} + \frac{\partial F_2}{\partial u_2}$$

is, to within a sign, the same Schur complement.

Similar procedures can be defined for the evolution problem, which, when each phase is implicit, becomes a modified rootfinding problem on each timestep with the Jacobian augmented with an identity matrix.

If the types of nonlinearities in the two components are different, a better method may be nonlinear Schwarz or ASPIN [6]. In nonlinear Schwarz, one solves component subproblems (by Newton or any other means) for componentwise corrections,

$$\begin{aligned} F_1(u_1^{k-1} + \delta u_1, u_2^{k-1}) &= 0 \\ F_2(u_1^{k-1}, u_2^{k-1} + \delta u_2) &= 0, \end{aligned}$$

and uses these legacy componentwise procedures implicitly to define modified residual functions of the two fields:

$$\begin{aligned} G_1(u_1, u_2) &\equiv \delta u_1 \\ G_2(u_1, u_2) &\equiv \delta u_2. \end{aligned}$$

One then solves the modified rootfinding problem

$$\begin{cases} G_1(u_1, u_2) = 0 \\ G_2(u_1, u_2) = 0 \end{cases}$$

by any means. If one uses Algorithm 3 to solve the modified problem, the Jacobian is

$$\begin{bmatrix} \frac{\partial G_1}{\partial u_1} & \frac{\partial G_1}{\partial u_2} \\ \frac{\partial G_2}{\partial u_1} & \frac{\partial G_2}{\partial u_2} \end{bmatrix} \approx \begin{bmatrix} I & \left(\frac{\partial F_1}{\partial u_1}\right)^{-1} \frac{\partial F_1}{\partial u_2} \\ \left(\frac{\partial F_2}{\partial u_2}\right)^{-1} \frac{\partial F_2}{\partial u_1} & I \end{bmatrix}, \quad (14)$$

which clearly shows the impact of the nondimensionalized cross-coupling in the off-diagonals, since the diagonal blocks constitute the identity. (This makes sense since the componentwise rootfinding problems are already solved at this outer level.) In practice, the outer Newton method must converge in a few steps for ASPIN to be worthwhile, since the inner iterations can be expensive. In nonlinear Schwarz, the partitioning of the global unknowns into u_1 and u_2 need not be along purely physical lines, as they would be if they came from a pair of legacy codes. The global variables can be partitioned into overlapping subsets, and the decomposition can be applied recursively to obtain a large number of small Newton problems. The key to the convenience of nonlinear Schwarz is that the outer iteration is Jacobian-free and the inner iterations may involve legacy solvers. (This is also true for JFNK, where the inner preconditioner may involve legacy solvers.)

Unless linear preconditioning of the global system is required to complement what is in effect nonlinear preconditioning through the inner solves, only distributed vector operations are needed to migrate legacy codes requiring coupling to this framework.

Error bounds on linear and nonlinear block Gauss-Seidel solutions of coupled multiphysics problems are derived in [226], wherein it is shown how the norm of the product of the off-diagonal blocks in (14) enters the analysis. Consider the linear case (13) for which the Jacobian of the physics-blocked preconditioned system is

$$\begin{bmatrix} I & A_{11}^{-1}A_{12} \\ A_{22}^{-1}A_{21} & I \end{bmatrix}.$$

Let the product of the coupling blocks be defined by the square matrix $C \equiv A_{11}^{-1}A_{12}A_{22}^{-1}A_{21}$, the norm of which may be bounded as $\|C\| \leq \|A_{11}^{-1}\| \|A_{12}\| \|A_{22}^{-1}\| \|A_{21}\|$. It is shown in [226] that, provided $\|C\| \leq 1$, any linear

functional of the solution (u_1, u_2) of (13) solved by physics-blocked Gauss-Seidel (Algorithm 1) satisfies conveniently computable bounds in terms of residuals of the individual physics blocks of (13). The cost of evaluating the bounds involves the action of the inverse uniphysics operators A_{11}^{-1} and A_{22}^{-1} on vectors coming from the uniphysics residuals and the dual vectors defining the linear functionals of interest. The extra work required to evaluate the bounds provides confidence that block Gauss-Seidel has been iterated enough to produce sufficiently accurate outputs of interest, be they point values, averages, fluxes, etc. The required actions of the inverse uniphysics operators are a by-product of an implicit method for each phase. The nonlinear case is similar, except that the Jacobian matrices from each physics phase that comprise C may change on each block Gauss-Seidel iteration.

Design patterns for multiphysics preconditioning. Preconditioning is essential for efficiency when a Krylov subspace method is used. While there has been considerable success in developing black-box, algebraic strategies for preconditioning linear systems (7), preconditioners for multiphysics problems generally need to be designed by hand, using building blocks described above for implementation. A few strategies for doing so are described next. A more complete discussion can be found in [5].

A simple strategy for constructing a multiphysics preconditioner is to use a Picard linearization to construct an approximate Jacobian for use in a JFNK solution strategy. Another strategy is to mimic operator splitting strategies to construct an approximate inverse through sequential application of single physics solvers. This can be especially useful when the individual physics components have no knowledge of each other but can solve appropriately posed linear subproblems.

Another approach is to use the multiphysics application itself as an approximate solution strategy. In this context, we can think of the outer Krylov or Newton iteration as acting to accelerate and/or stabilize the existing solution method. This idea first appeared in [227]. The multiphysics application has to satisfy certain properties for this idea to be used; in particular, the ability to calculate a clean, full residual and calculating corrections rather than full updates to solution variables are needed.

A widely used approach is to use a block diagonal approximation of the Jacobian of the system, which can be effective in many circumstances. This approach ignores the off-diagonal coupling, however, and analysis of the problem can help to identify which couplings are essential and must be captured, leaving the other couplings to be resolved by the outer Krylov/Newton iterations. Such approaches generally lead to identification of a Schur complement that embodies the important coupling, and their success relies on judicious approximation of the Schur complement [228, 229].

3.2. Continuum-Continuum Coupling

Multiphysics problems require the computation of interacting fields that may feature different scales in time and/or space and may additionally be defined over different spatial regions. Mathematical modeling of such complex systems can lead to large systems of PDEs, possibly with different structures on different domains. Roughly speaking, two approaches are used. In a single-domain approach, the coupled system is considered globally, and the state variables are defined over the entire spatial domain of the problem; a simple example is a thermomechanical system whose solution consists of the temperature and displacements. Alternatively, different sets of state variables and/or models may be considered on different spatial or physics domains. A simple example for different spatial domains is fluid-structure interaction, where the fluid velocity and pressure are defined in the fluid domain and displacements are defined in the solid domain.

In general, single-domain approaches can be conceptually simple and straightforward. However single-domain models may lack flexibility; solution of the single domain problem may be expensive; and, more important, generalization to more complicated models may not be possible. For example, no single-domain approaches exist for coupling multiphase, multicomponent flow in porous media with multicomponent free flow, whereas transfer conditions based on mechanical, chemical, and thermal equilibrium can be derived for a model posed on two domains [230]. On the other hand, the divide-and-conquer approach of multidomain models allows the use of different numerical strategies tuned to the physics of each component. The price of such flexibility is the need to formulate suitable transfer conditions, in both model formulation and model discretization, which guarantee equilibrium and well-posedness of the coupled system, and to devise efficient solution techniques that are robust, stable, and accurate.

Multidomain approaches require coupling of continuum models in space and, for time-dependent problems, in time. This coupling must respect the mathematical formulation of the coupled problem and must account for different

meshes and discretizations used in each domain. Coupling components that are allowed to proceed at their own timescales can be particularly problematic. The following subsections describe techniques currently being used in several selected multiphysics applications. We note that, currently, highly scalable algorithms for these operations are not readily available.

3.2.1. *Methods for coupling multiphysics components in space*

Multiphysics applications require some means to exchange information between physics components. Preserving important properties, such as conservation or order of accuracy, is a principal concern. In the simplest case, all the physics components are defined on the same physical domain Ω and use the same mesh Ω^h . In this situation, the coupling is defined by the physical model, and no error is incurred in transferring data between meshes. Numerically, some additional interpolations may be needed to align variables with different centerings or to achieve some consistency between different discretization schemes. In addition, the resolution must be fine enough to meet the requirements of all the physics components that share the mesh. More often, each physics component uses a mesh that is tailored to best represent the physics. These meshes can either fully² or partially overlap and can also introduce different discrete representations of the geometry of Ω . Moreover, different physics components can be defined on distinct physical domains Ω_i (where i indexes the physics components) and are coupled across lower-dimensional interfaces $\Gamma_{ij} = \partial\Omega_i \cap \partial\Omega_j$, $i \neq j$, between the domains, which may also evolve in time. In this case, the meshes $\Omega_i^{h_i}$ and $\Omega_j^{h_j}$ may or may not match along the interface. When the restrictions $\Gamma_{ij}^{h_i}$ and $\Gamma_{ij}^{h_j}$ of the meshes to the common interface Γ_{ij} do not match, information transfer across the interface must account for these differences in the discrete representation of Γ_{ij} . In some cases the picture can be further complicated when Γ_{ij} is determined as part of the calculation. The rest of this section provides a brief overview of methods currently used by multiphysics applications, and Section 4.2.3 considers related software issues.

At its core, transferring discrete fields between different meshes is a form of interpolation that is determined by how the field is represented on the source and target meshes. Search algorithms that determine the support of the interpolation on the source mesh are critical for constructing efficient transfer operations. Tree structures such as octrees, alternating digital trees, and k-d trees are used when the points to be evaluated are randomly distributed. Mesh topological adjacencies are used when the points are ordered such that the next point is from a nearby element. When the geometries of the discrete domains do not match, a destination point may lie outside the discrete source domain. In this case extrapolation can be used; ideally, information about the physical domain and boundary conditions will inform the extrapolation process.

In more complex cases, the transfer process must satisfy various kinds of local and global constraints. Examples include matching the integral of the solution and guaranteeing a divergence-free property [231]. In many circumstances, local conservation constraints are imposed. For a mesh transfer \mathcal{T} to be conservative in a region $\mathcal{R} \subset \Omega$, it must satisfy

$$\int_{\mathcal{R}} \mathcal{T}f = \int_{\mathcal{R}} f,$$

where f is the source grid data. Enforcement in a discrete sense can be represented as

$$\int_{\mathcal{R}} \phi_i \mathcal{T}f = \int_{\mathcal{R}} \phi_i f,$$

where $\{\phi_i\}$ is a set of test functions. For interpolation, they are simply delta functions at points on the destination mesh. If f is represented in a Galerkin formulation, $\{\phi_i\}$ is the basis that is used on the destination mesh. Because f and ϕ_i are only piecewise smooth on source or target meshes, the product $\phi_i f$ is in general smooth only within the intersection of source and target elements. Therefore, using quadrature rules over source or target elements to compute $\int \phi_i f dx$ may violate the regularity assumption of quadrature rules, and large errors and even instabilities can result from computing these integrals using quadrature rules based on only the source or destination mesh [232, 233]. Knowledge of the complete intersection of the source and target meshes can then be used together with appropriate quadrature rules to construct \mathcal{T} [232, 234–238].

²A similar need is encountered for arbitrary Lagrangian-Eulerian (ALE) methods, where periodic mesh remapping is needed.

Transferring information across a shared interface between two domains depends on the smoothness of solutions on either side of the interface, discretization schemes, and alignment of the two meshes on either side of the interface. When the meshes align at the interface, direct transfer can be used effectively. For an illustration, consider coupling Navier-Stokes flow on the surface and Darcy flow in subsurface porous media (Figure 11) [239]. The free-flow region is denoted Ω_1 , and the porous medium is represented as several layers (regions A, B, C) with different permeability and faults (region D). The interface conditions enforce continuity of the normal component of the velocity across the interface, a force balance at the interface, and the Beavers-Joseph-Saffman model relating the tangential free flow velocity to the shear stress at the interface [240]. Let a_{NS} represent the bilinear form for the Navier-Stokes flow discretized with a continuous finite-element method, and a_D represent the bilinear form for the Darcy equations discretized with a discontinuous Galerkin method to accurately capture irregularities in the subsurface region. A grid is used that covers the entire domain but captures the interface between the two models. This means that the bilinear form a_Γ for the interface coupling conditions across the interface can easily be constructed by simply using values of the solution obtained at the interface. The fully coupled problem is then given by

$$a_{NS}(u, v) + a_D(u, v) + a_\Gamma(u, v) = L(v), \tag{15}$$

where v is drawn from suitable test spaces and L is a forcing term.

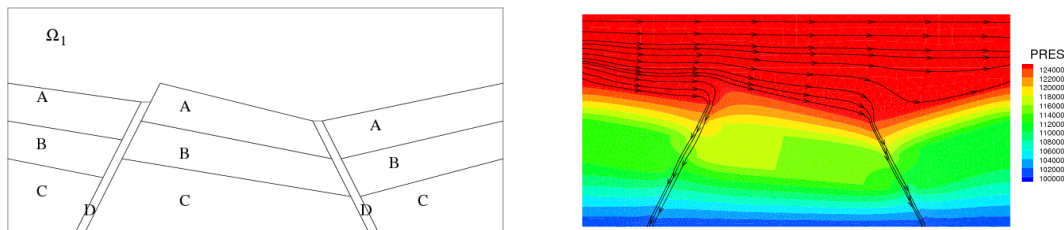


Figure 11: Computational domain (left) and pressure and velocity contours (right) for coupling of Navier-Stokes flow on the surface and Darcy flow in subsurface porous media.

A fully coupled solution method is the most stable and accurate approach to solve (15). However, this leads to a large system of nonlinear equations that may be difficult to precondition and expensive to solve. One way to lower the cost of solving (15) is to use a two-grid technique [241]. In the first step, (15) is solved on a coarse mesh with resolution H ; in the second step, decoupled “single physics” problems are solved on a finer mesh, with resolution $h = H^2$. Boundary data along the interface for the decoupled problems is obtained from the coupled coarse mesh solution. This approach is nearly an order of magnitude faster and produces solutions with accuracy comparable to the fully coupled approach, but it is less stable and accurate. The solution obtained with this algorithm (Figure 11) shows smooth velocity contours, even into the fault regions. For details, see [241].

The mortar finite-element method is another approach for treating interface conditions [242, 243]. It is especially effective when the meshes on the two sides of the interface do not match and when some quantity (such as flux or stress) is needed at the interface. The mortar finite-element method yields a variationally consistent formulation and provides an abstract framework for analyzing the weakly coupled discrete model. The algebraic problem can then be stated as a saddle point or, equivalently, as a constrained system where static condensation of the Lagrange multiplier is realized and the coupling conditions are imposed on the spaces. Mortar formulations are widely used in the context of nonlinear contact problems in structural mechanics and replace simple node-to-node or node-to-segment penalty approaches; see, for example, [244]. The surface traction plays the role of the Lagrange multiplier, which has to satisfy a sign condition in the normal direction, which is due to the nonpenetration condition, and the friction law, which combines normal and tangential surface forces. The inequality conditions can be reformulated as a weakly consistent saddle point system in terms of a nonlinear complementarity function. Then semi-smooth Newton methods provide a superlinear solver, and an efficient algorithmic implementation can be based on the choice of biorthogonal basis functions, resulting in a diagonal mass matrix.

Figure 12 illustrates the acoustic-structure interaction for a 3D loudspeaker geometry; see [245]. The problem has been formulated in a weakly consistent mortar setting, and highly nonmatching meshes are applied for the numerical

simulation.

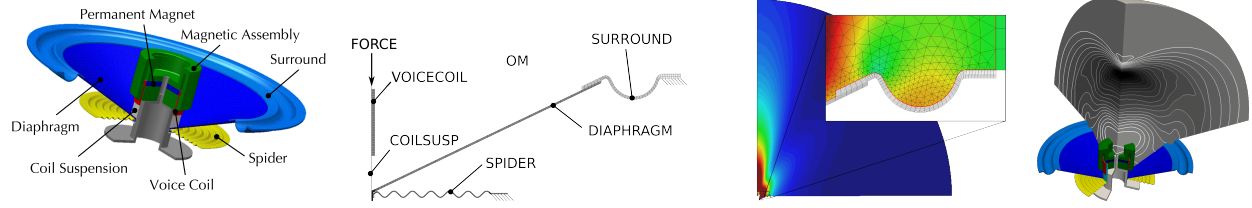


Figure 12: Schematic view of the geometric and acoustic pressure field in acoustic-structure interaction for a 3D loudspeaker geometry.

Of crucial importance for the stability of the method is a uniform inf-sup condition for the pairing between dual and primal variables. In order to obtain optimal *a priori* error results, the discrete Lagrange multiplier space must additionally satisfy a best approximation property. With these two ingredients, one can easily construct locally defined higher-order methods. Moreover, the influence of curvilinear interfaces can be analyzed within the framework of variational crimes, and the weak coupling concept can be easily combined with blending or isogeometric finite elements.

Although mortar techniques were originally designed for nonoverlapping domain decomposition methods, the idea of weakly coupled subdomain formulations can be used to design two-scale overlapping subdomain approaches. Globally, a simple, coarse-scale discrete model is used, whereas a more sophisticated model is applied locally on a small patch. Different equations, mesh sizes, and geometries can be considered. Here, one has to distinguish between one- and bidirectional couplings. One-directional coupling can be regarded as a postprocess and allows locally for a higher resolution of the numerical solution. The possibly nonlinear system has the structure of a lower triangular block system and thus can easily be solved in terms of a Gauss-Seidel-type solver, possibly in combination with massive parallel algebraic multigrid or domain decomposition techniques. For bidirectional coupling, iteration between the small patch and coarse-scale models is required to obtain a physically correct solution.

3.2.2. Methods for coupling multiphysics components in time

Coupling time-dependent continuum models depends on the temporal resolution requirements of the physics components. Typically, each component of the solution is allowed to evolve governed by its own physics. However, there can be wide disparity in the timescales of the constituent physics components, which may in turn be quite different from the timescale of the problem being solved. Replacing instantaneous interactions between different physics with discrete interactions can introduce subtle stability and accuracy issues. Likewise, the multiphysics problem may conserve some global quantity (such as energy), but this property may not be preserved by the simulation.

In evolutionary problems (2a-2b), the spatial discretization leads to a system of ordinary differential equations (ODEs) or differential algebraic equations (DAEs) in general, that is, ordinary differential equations possibly with algebraic constraints. A timestepping approach is then used to advance the solution in time. This strategy is called the *method of lines*. For exposition brevity, we initially focus on ordinary differential equations, with initial conditions $u(t_0)$ and computing values of the solution $\{u(t_n)\}_{n=1}^{N_t}$ at discrete times $t_1 < \dots < t_{N_t}$.

In an *explicit method*, future states are given by a formula involving known computed quantities that can be evaluated directly. The simplest example is forward Euler:

$$u(t_{n+1}) = u(t_n) + \Delta t f(u(t_n)),$$

where $\Delta t = t_{n+1} - t_n$. Explicit methods are inexpensive per iteration because they require a fixed number of function evaluations to complete the step; however, they typically have severe stability restrictions and, therefore, are suited only for nonstiff problems.

A *fully implicit method* advances the solution in time by using current information and an inversion process with no explicit steps. The simplest example is backward Euler:

$$u(t_{n+1}) = u(t_n) + \Delta t f(u(t_{n+1})).$$

Fully implicit methods are relatively expensive because of the need to solve (non)linear systems at each step; however, they have favorable stability properties [246, 247]. Inclusion of explicit elements gives rise to semi-implicit methods [248, 249].

Spatial discretization of multiphysics problems ((2a)–(2b)) may result in problems in which the tendency $f(u)$ has components with widely different dynamics, often classified as stiff and nonstiff [250–255]. This case, called *additive splitting*, generally corresponds to problems in which the physics components act in the same spatial domain, such as reactive transport or rad-hydro [256–261]. When a partitioned time integration method is used with physics components defined on the same spatial region, some work must be done to synchronize the timestepping in each physics component. When one component has a much faster timescale relative to another, a simple strategy is to use an integer multiple of the faster timescale for the slow timescale. In *component partitioning* some solution components correspond to a stiff subproblem and the others to a nonstiff one; examples are fluid-structure interaction and flow with different speeds [262–264]. Component partitioning can be used, for instance, if physics components are defined on different spatial domains; the same subcycling idea can work, but some protocol for communicating information across the interface is needed to ensure accuracy and stability. This leads to the following two conceptual types of partitioning:

$$f(u) = f_1(u) + f_2(u) \quad \text{additive partitioning,} \quad (16)$$

$$f(u) = \begin{bmatrix} f_1(u_1, u_2) \\ f_2(u_1, u_2) \end{bmatrix} \quad \text{component partitioning.} \quad (17)$$

Simple transformations can be used to cast one type of splitting into the other [250, 254]. We will henceforth use form (16) in which f_1 is relatively nonstiff and its presence points one to use an explicit integrator for efficiency, whereas f_2 is relatively stiff and requires an implicit integrator for stability reasons. For instance, in (16) f_1 can incorporate all slow transport components and f_2 stiff kinetic terms, and with (17) u_1 may represent coarse-grid points and u_2 finely refined states or structure and fluid states, respectively. Partitioned integrators were introduced to solve such problems efficiently.

Partitioned methods use time integrators with different properties or timesteps for each (physics) component. They allow black-box implementation of the physics components and different time steps in each physics component. Nevertheless, this formalism allows one to analyze rigorously the stability and accuracy of the compound scheme. Partitioned methods can be tightly coupled (in time); one simple example is the forward-backward Euler method:

$$u(t_{n+1}) = u(t_n) + \Delta t f_1(u(t_n)) + \Delta t f_2(u(t_{n+1})).$$

Strang splitting methods [2] are also particular cases of (loosely coupled) partitioned schemes:

$$\begin{aligned} u_\star &= u(t_n) + \frac{\Delta t}{2} f_2(u_\star) \\ u_{\star\star} &= u_\star + \Delta t f_1(u_\star) \\ u(t_{n+1}) &= u_{\star\star} + \frac{\Delta t}{2} f_2(u_{\star\star}) \end{aligned} .$$

The caveats of such a strategy include order reduction and low accuracy as illustrated in [265]. The temporal splitting error can be fairly easily avoided in tightly coupled schemes [254, 266]. Other tightly-coupled strategies involve embedding explicit evaluations inside an implicit solver [267].

We differentiate two cases that have a more practical flavor: multirate and implicit-explicit (IMEX). *Multirate* methods are partitioned methods that use the same scheme in each partition but with different timesteps [262–264, 268–274]; this goes back to subcycling. Multirate schemes are also used to address accuracy considerations. One instance is applications that employ temporal and spatial adaptive grids. Another instance is shown in Figure 13, which illustrates a multirate scheme that uses smaller timesteps for the reaction component in a reaction-diffusion problem.

An *IMEX* method is a partitioned method that uses an implicit integrator on one partition and an explicit integrator on the other. Therefore, an optimal tradeoff between efficiency and stability becomes possible [254, 266, 275, 276]. Moreover, it is generally desirable to use an explicit integrator when nonlinear or nondifferentiable spatial discretizations are used, such as shock resolving schemes [277]. In this case IMEX alleviates the difficulties posed by such

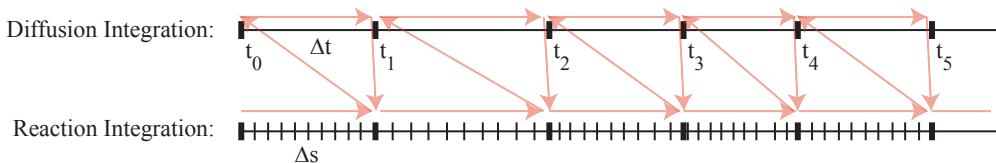


Figure 13: Illustration of multirate operator-splitting time discretization.

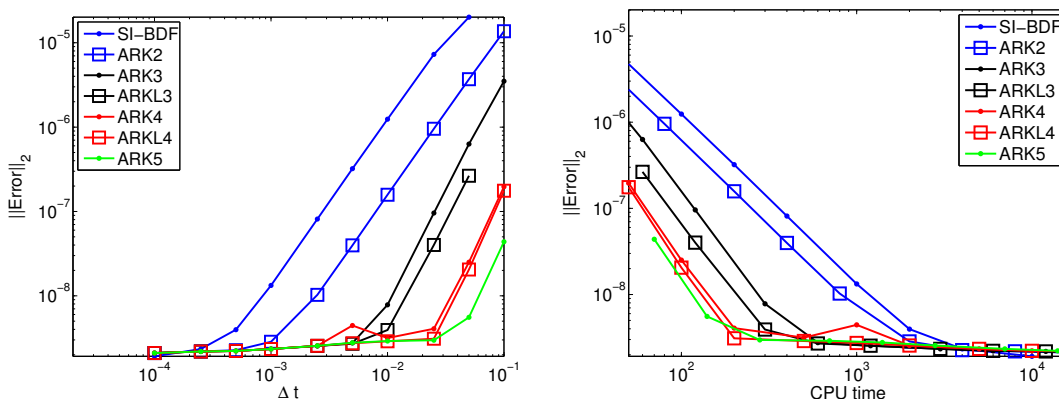


Figure 14: Global convergence and timing of implicit-explicit schemes for a 2D Euler rising bubble test problem. Method SI-BDF is second order, and the rest have the order suffixed to their name.

a problem on the nonlinear solver, which is applied only on the stiff component. Furthermore, IMEX methods generalize fully implicit ($f_1 \equiv 0, f \rightarrow f_2$) and fully explicit ($f_2 \equiv 0, f \rightarrow f_1$) methods and can provide the best balance between accuracy and efficiency. In both cases the temporal splitting error can be made of the same order as the schemes used separately; however, stiffness leakage can be a concern [266]. Partitioned methods based on the classical linear multistep [253, 278], Runge-Kutta [250, 251, 262, 266], and extrapolation methods [254, 279] have been proposed and successfully applied in practice. In Figure 14 we illustrate the temporal convergence rate for a 2D Euler rising-bubble problem by using high accuracy in space [258] and different IMEX time integration schemes. The splitting is based on linearizing the PDE, extracting the stiff wave components, L , and then integrating (2a–2b) as

$$\partial_t u = (f(u) - L u) + L u,$$

where the first term is treated explicitly and the last linear term implicitly. We note that IMEX schemes do not suffer from splitting errors and that the high-order methods are the most efficient.

Stability considerations for partitioned methods. Partitioned methods use multiple time integrators with well-known stability properties when applied to the compound problem (2a–2b). When applied on partitions, however, the stability properties of individual methods do not transfer directly to the stability of the compound problem. Consider two integrators that are stable when applied separately on each partition but where the resulting compound method may lead to unstable overall computations [254, 280–282]. Such a situation is illustrated in Figure 15 for extrapolated IMEX schemes, where the dependence between the implicit and explicit stability regions is clearly visible. Note that larger stability regions for the implicit method correspond to smaller stability regions for the extrapolated explicit methods. This behavior must be taken into account when selecting an extrapolated IMEX scheme for multiphysics problems. If the implicit term is used to integrate symmetric discretizations of hyperbolic operators that lead to pure imaginary eigenvalues, an A -stable implicit method is needed, and one may expect a penalty on the explicit step-size because the explicit stability region is smaller (Method A). Discretizations of parabolic operators that lead to negative real eigenvalues can use an $A(\alpha)$ -stable implicit method for $\alpha < 90^\circ$, so the larger stability region for the explicit

operator allows larger stable explicit steps (Method B). These examples tacitly assume that the implicit and explicit operators are simultaneously diagonalizable, which happens almost exclusively only for scalar problems; however, in practice, this analysis typically holds for vector-valued problems, especially when stiff components are present [280].

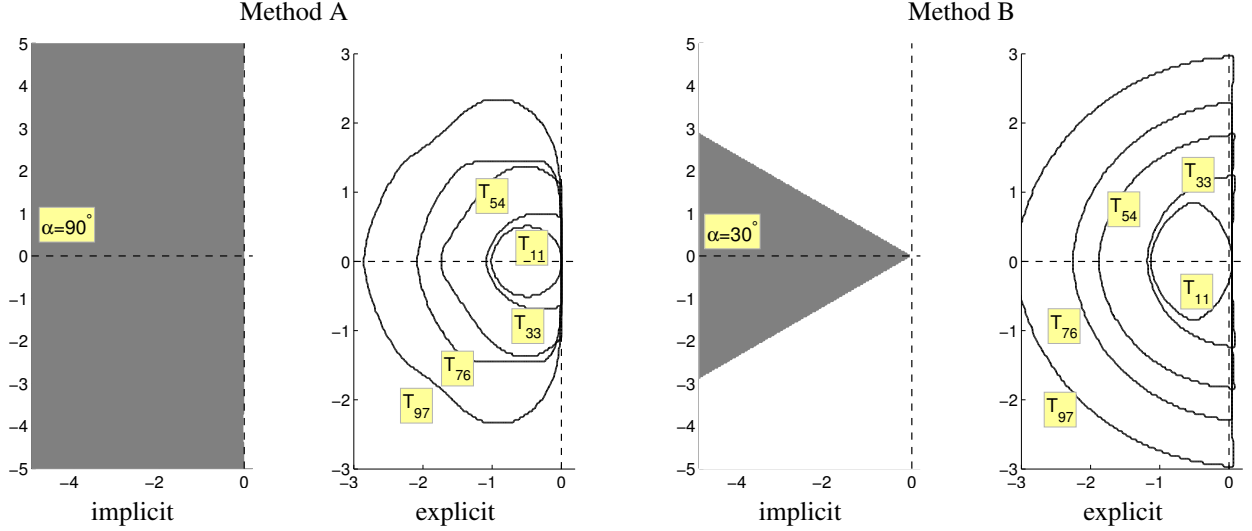


Figure 15: Stability regions of implicit-explicit pairs for extrapolated IMEX schemes. The implicit methods are $A(\alpha)$ -stable for different values of α . The explicit stability regions indicate different orders of extrapolation. The stability region of the IMEX scheme contains the cross product of the regions for these pairs. This figure is adapted from [283].

Stability properties of partitioned timestepping methods are also affected by treatment of interfaces. This can be illustrated by a simple model of heat conduction through two adjacent materials in subdomains Ω_1 and Ω_2 that are coupled across their shared and rigid interface Γ through a jump condition:

$$\begin{cases} \frac{\partial u_i}{\partial t} - \nu_i \Delta u_i = f_i, & x \in \Omega_i, \\ -\nu_i \nabla u_i \cdot \mathbf{n}_i = \kappa(u_i - u_j), & x \in \Gamma, \\ u_i(x, 0) = u_i^0(x), & x \in \Omega_i, \\ u_i = g_i, & x \in \partial\Omega_i \setminus \Gamma \end{cases} \quad (18)$$

for $i, j = 1, 2, i \neq j$. Three time integration strategies can readily be described. The fully coupled approach solves Equation (18) with u_1 and u_2 and all terms evaluated at the same time level. This method is unconditionally stable, conserves energy, and converges optimally under standard smoothness assumptions. A second scheme is to again keep u_1 and u_2 at the same time level but to lag in time the transmission conditions on Γ . This IMEX scheme decouples the subdomain solves but is stable only when κ is small compared with ν_1 and ν_2 , together with a CFL-like stability restriction on the timestep. Here the tradeoff is between the need for smaller timesteps and the reduced cost of solving easier problems in each subdomain. A third, partitioned scheme uses interface data at the previous timestep from the other subdomain. Since each subdomain requires data from the other subdomain, this data-passing approach is less flexible but is stable when κ is large compared to ν_1 and ν_2 . The data-passing approach also does not have a stability constraint on the timestep; see [284] for details and analysis.

Partitioned time integration methods are often used in a sequential fashion as in Algorithm 2, first advancing a single physics component in time, followed by updating information on an interface, advancing a second physics component in time, and finally updating interface information from the second component. The strategy is often referred to as staggered computation; and the techniques used to update data on the interface play a crucial role in the stability of this approach, irrespective of the time integration methods used for the component physics. In particular, there can be a time lag in the interface data that leads to an imbalance in the interface transmission conditions, even for

strategies that are formally second-order accurate in time. For example, in fluid-structure interaction, a small error in the structural displacements may result in large errors in the fluid pressure. The stability and accuracy can be improved by using the combined interface boundary condition (CIBC) [233], which solves the transformed conditions for the interface quantities based on their space derivatives and time derivatives. The CIBC technique allows the construction of a staggered coupling procedure with similar order of accuracy and stability as tightly coupled computations for some classes of fluid-structure coupling problems.

Despite the issues associated with partitioned and mixed approaches to multiphysics time integration, these elements can be combined to produce stable, accurate, and efficient simulations. Such an approach is illustrated by hierarchical adaptive mesh refinement (AMR), which we discuss in the context of lean, premixed turbulent combustion [285]. This is a demanding multiscale simulation involving fluid flow, chemical reactions, and heat transfer that resolves spatial scales that span two orders of magnitude and timescales that span three orders of magnitude. A finite-volume formulation is used to enforce crucial conservation properties. The AMR strategy is tailored to the mathematical structure of the problem to exploit the relationship between the scales and physical processes in the problem. The fluid equations are advanced in time by using a fractional-step projection formulation, whose elements include explicit advection (via a Godunov scheme), implicit diffusion (via a Crank-Nicolson predictor-corrector scheme), Strang splitting [2] for the chemical reactions, and two discrete projection steps to enforce a modified, divergence-free constraint.

AMR introduces two complexities. The first is the need to compute discrete differential operators where the resolution changes; this computation is done through a simple direct interpolation of data at a coarse resolution to a finer resolution. Local piecewise quadratic interpolation of coarse data is used to align the coarse data with fine data and to calculate fine-scale fluxes at the interface [286].³ In addition, a re-fluxing procedure [289] is employed to enforce conservation and free-stream preservation at these changes in resolution. The second complexity is local subcycling in time: a global time step on the entire locally refined grid can be inefficient and will lead to a wide range of CFL numbers across the grid, compromising the accuracy of the Godunov advection scheme. In order to handle this, timesteps are chosen in each refinement region to maintain a CFL number close to 1. Practically speaking, this means that a grid with resolution h requires twice as many timesteps as the next-coarser region with resolution $2h$ (assuming the refinement ratio is 2). Overall, this AMR scheme is second-order accurate in time and space; see [290] and [286] for details.

We note that it is difficult to devise partitioned time integration schemes that are better than second-order accurate in time. *Spectral deferred correction* (SDC) [291, 292] is one potential way to improve time accuracy. Starting with a provisional solution obtained with a simple, low-order method, SDC generates a sequence of corrections that can lead to higher-order accuracy in the solution. Specifically, one writes the solution to the problem in terms of a Picard integral equation and evaluates the integral using interpolated values of the provisional solution. Spectral accuracy of the integral can be obtained by using Gaussian quadrature nodes in the quadrature. In principle, the order of accuracy can be increased up to the order of accuracy of the quadrature used to evaluate the Picard integral. Variations of SDC have also been developed for semi-implicit time integration methods [293, 294].

3.3. Continuum-Discrete Coupling

The development of methods to couple models across multiple physical scales is a fairly new area of investigation that has grown rapidly in the past decade. Much of the initial multiscale methods development has focused on the linking of distinctly different spatial scales, in which the coarse scale is represented by a PDE model and the fine scale by either a discrete model (e.g., atomistic) or a fine-scale heterogeneous continuum model (e.g., composite materials in which the microstructural constituents are modeled). A key feature of these problems is the existence of a strong scale separation that strongly influences the method developed. The majority of the scale linking methods fit into one of two classes: hierarchic (also called information passing) or concurrent methods [295].

In hierarchic methods information is “upscaled” from the fine-scale model to the coarse scale to provide model information (e.g., constitutive relations) at a point in the coarse scale. Figure 16 graphically depicts this process for the case where a finite-element mesh is used to discretize the macroscale continuum equations. As part of the construction

³Local piecewise linear interpolation can also be used [287]. Failure to adequately address this issue can lead to nonphysical artifacts in the solution [288].

of the element stiffness matrix the “force/deformation” relationship at the numerical integration points in the elements needs to be evaluated. In this two-scale model, the “force/deformation” relationship is determined through the analysis of a model defined on a representative volume element composed of one-dimensional fiber elements [296]. There is a long history of development of homogenization methods [297, 298] to link continuum models between two scales. Recently these methods have been generalized to link from atomistic to continuum [299].

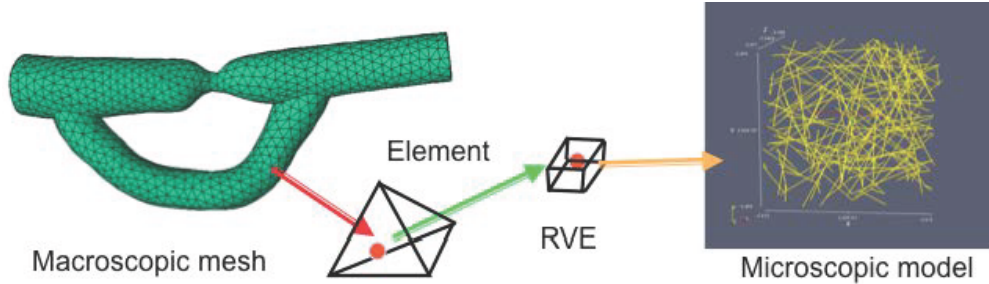


Figure 16: Two-scale model for soft, fibrous material. This diagram indicates the transition from the element level of a macroscopic mesh to a representative volume element (RVE) and then to a microscopic model [296].

In concurrent methods the different models exist concurrently over different portions of the domain. Concurrent methods are frequently used for problems where an atomistic representation is needed in some small fraction of the domain while the majority of the model can use a much coarser representation. Among the most mature of these methods is the quasicontinuum method [68, 300, 301] in which atomistic mechanics are applied in critical regions. In noncritical regions entire groups of atoms are approximated in terms of a small number of representative atoms with specific assumptions applied on the variation of fields over the group of atoms in much the same way as a field is approximated in terms of shape functions over a mesh of elements. Other concurrent methods employ a continuum partial differential equation for the coarsest scale [302–305]. The key component of these methods is the linkage of the atomistic to the continuum. Methods that attempt to link through just the interface between the portions of the domain with different models have not produced good results. Therefore, methods are needed that employ some form of overlap region in which both the continuum and atomistic models exist and are blended from one model to the other. Figure 17 shows a simple graphic of this idea, with a strictly atomistic region to the right; an overlap region in the middle where both atomistic and continuum models interact; and a continuum region to the left that, in this example, includes some dummy atoms that follow the continuum and support finalization of the transition from atomistic to continuum. A key function of the overlap region is supporting the transition of atomistic quantities to continuum quantities, which must be done with care in the zero temperature case and has proved difficult in the finite-temperature case since a method is needed to determine what portion of the energy content goes into thermal loads and what goes into mechanical loads.

In many problems the areas in which the fine-scale models are needed are not known *a priori*. In order to deal with such situations, adaptive multiscale methods are being developed [306–308]. Figure 18 shows the application of an adaptive concurrent atomistic/continuum method to study the growth and interactions of dislocations around nanovoids [303, 307].

An alternative scale-linking approach is the heterogeneous multiscale method [309], which defines compression operators to relate discrete to continuum scales and reconstruction operators to relate continuum to discrete scales. In addition to dealing with the complexities of linking the 10 orders of magnitude of spatial scales when linking atomistic to continuum, the timescales that must be bridged can be as high as 15 orders of magnitude. The same methods used in multiphysics time coupling are considered in multiscale time coupling. Moreover, new methods such as the equation-free multiscale method [310] have been developed that link statistically averaged fine time- and space-scale realizations to the coarse time and space scales.

3.4. Error Estimation

Reliable error estimation for multiphysics problems is difficult. Because *a priori* bounds must account for the largest possible error across a large class of solutions, they do not yield sharp estimates and are useful primarily for

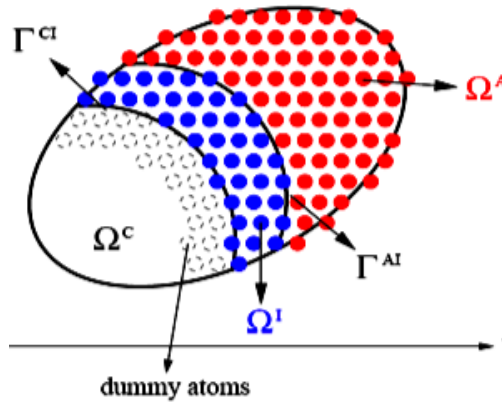


Figure 17: Atomistic-to-continuum concurrent method with overlap region.

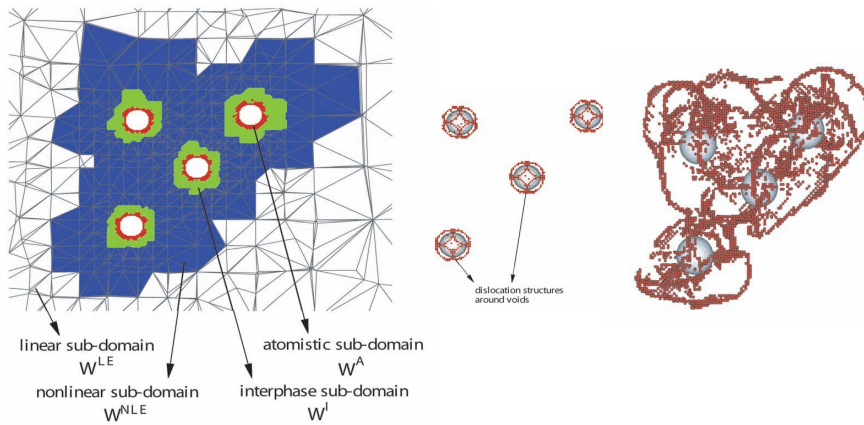


Figure 18: Four-void multiscale example. **Left:** Models adaptively selected at step 20. **Center:** Dislocation structure around voids at step 20. **Right:** Evolved dislocation structure at step 28.

characterizing the general behavior of the error, such as order of accuracy. In addition, most *a priori* estimates are asymptotic and are useful primarily for relatively well-resolved problems having smooth solutions. In this regime, *a priori* bounds can be useful for constructing error estimates based on extrapolation. For example, the AMR algorithm described in Section 3.2.2 uses Richardson extrapolation to determine where additional spatial resolution is needed. Multiphysics applications that use finite-difference or finite-volume discretization can use similar extrapolation-based techniques to estimate error.

In contrast, *a posteriori* analysis focuses on the error in a particular computed solution. Methods for directly estimating the error in a finite-element solution include residual-based estimators that construct estimates based on element residuals, hierarchical estimators that project the residual onto some larger finite-element space, and recovery-type estimators that focus on ways to reconstruct the solution gradient [311]. Generally, these estimators are more reliable for smooth, well-resolved solutions. Extending these ideas to develop robust and accurate error estimators for coarse meshes and nonsmooth solutions is an active area of research [312–314].

Alternatively, the error in some scalar model response quantity Q may be of interest. In this case, duality and adjoint operators can be used to obtain very accurate estimates of the error in Q [315–317]. This approach requires solving a related problem involving the adjoint operator together with a projection of the adjoint solution in the direction of the residual. For time-dependent problems, solving a problem involving the adjoint operator entails integrating backwards in time to the initial state at every timestep. When the problem is nonlinear, some application

state must be saved during the forward time integration in order to construct the adjoint operator during the backward time integration. Storage and data management requirements can make duality-based *a posteriori* error estimation impractical. Also, because of the way that multiphysics applications are constructed from components, global adjoints and residuals are often not available. Consequently these duality-based *a posteriori* techniques are not widely used.

Even if full adjoints and residuals were available, the divide-and-conquer approach commonly encountered in multiphysics applications presents special difficulties for error estimation. Significant extensions and alterations of the standard approaches are required that are well developed for single physics problems. The key features required to obtain accurate error estimates for multiphysics problems include the following.

- Estimating the effects of errors passed between physical components
- Estimating the effects of processing transferred information, for example, upscaling, downscaling, and changing discretizations
- Quantifying the effects of multiphysics solution algorithms on stability by identifying perturbations to adjoint operators
- Estimating the effects of using a finite number of iterations in the solution when an iterative solver is used

The following subsections illustrate some of these issues. In this discussion we collectively refer to these strategies as multiscale, multiphysics, multidiscretization operator decomposition (M³OD), which encompasses multiphysics methods already discussed to this point:

- Operator splitting, or partitioned methods
- Decoupling “loosely coupled” systems that are related through parameter passing
- Coupled systems in which upscaling and downscaling operators are introduced
- Incomplete iteration in block or nonlinear solvers
- Different discretizations for different components, or in different regions or at different times in the same equation

While these methods or issues are very different, they all share the common feature that they involve some substantial alteration of the system of differential operators representing the components of the underlying system. These modifications may arise from formulation and/or discretization and must be accounted for in order to achieve robust and reliable error estimation for multiphysics problems.

This common feature to all M³OD approaches results in a common issue:

M³OD discretizes the instantaneous interactions between the component physics in the multiphysics model.

One potential consequence is the possibility of new forms of instability, some of which may be subtle. Another consequence is that as information is exchanged between physical components, numerical error is also exchanged, and the error in one component may well pollute the solution of another component. We illustrate the range of methods and effects using three common examples.

3.4.1. Operator splitting for reaction-diffusion equations

A classic example of M³OD is operator splitting for a reaction-diffusion equation. Accuracy considerations dictate the use of relatively small steps to integrate a fast reaction component. On the other hand, stability considerations over moderate to long time intervals suggest the use of implicit, dissipative numerical methods for integrating diffusion problems. Such methods are expensive to use per step, but relatively large steps can be used on a purely dissipative problem. If the reaction and diffusion components are integrated together, then the small steps required for accurate resolution of the reaction lead to an expensive computation.

In multirate operator splitting, the reaction and diffusion components are integrated independently inside each time interval of a discretization of time and “synchronized” in some fashion only at the nodes of the interval. The reaction component is often integrated by using significantly smaller substeps (e.g., 10⁻⁵ smaller is not uncommon) than those used to integrate the diffusion component, which can lead to a tremendous computational savings.

We illustrate some possible effects for multirate operator splitting using the well-known Brusselator problem

$$\begin{cases} \frac{\partial u_1}{\partial t} - \nu_1 \Delta u_1 = 0.6 - 2u_1 + u_1^2 u_2, \\ \frac{\partial u_2}{\partial t} - \nu_2 \Delta u_2 = 2u_1 - u_1^2 u_2, \\ \text{suitable initial and boundary conditions,} \end{cases}$$

with $\nu_1 = \nu_2 = 0.025$, over the one-dimensional spatial domain $[0, 1]$. We discretize using:

- a second-order, piecewise linear finite-element method in space with 500 elements;
- a standard first-order multirate splitting scheme;
- the implicit second-order trapezoidal rule with a timestep of 0.2 for the diffusion; and
- first-order implicit Euler with a timestep of 0.004 for the reaction.

In particular, the resolution is sufficient to ensure that spatial discretization errors are negligible and that the reaction and diffusion components are each solved with reasonable accuracy. Overall, the method is first-order accurate in time. See [257] for details.

We show the results in Figure 19. On the left, we plot the error versus timestep for different final times. We see that on sufficiently long time intervals, there is a critical timestep above which convergence fails. On the right, we plot a typical long time unstable solution. Note that the oscillations are *not* on the scale of the spatial discretization.

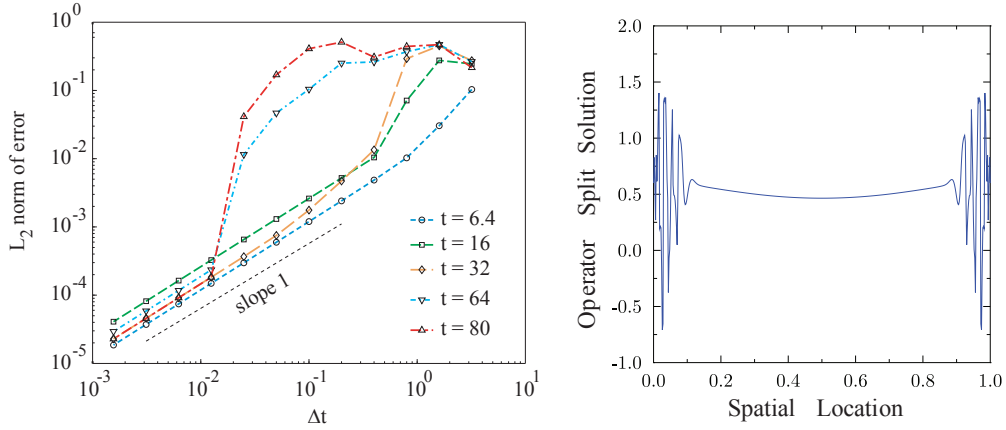


Figure 19: Left: On moderate to long time intervals, there is a critical time step above which convergence fails. Right: A typical unstable solution.

This example well illustrates the fundamental issue arising in solutions using M^3OD . When divorced from the stabilizing effect of diffusion, the solution of an unstable reaction component is free to grow too rapidly. In this case, the reaction component is not constrained by boundary conditions imposed on the solution of the diffusion component, and consequently its values near the boundaries grow far too rapidly.

3.4.2. Iterative solution of parabolic problems coupled through a common boundary

Operator splitting is an extreme form of M^3OD because it is a blunt factorization of the diffusion-reaction differential operator whose effects cannot be repaired, for example, by iteration. We next discuss a less-severe example of M^3OD that nonetheless presents some difficulties and puzzling behavior.

We consider a system of parabolic problems, where each equation is posed on a different domain but the domains share a common boundary Γ . In contrast to the example Eq. (18), we impose continuity of state and normal flux across Γ :

$$\begin{cases} \frac{\partial u_1}{\partial t} - \nabla \cdot (\nu_1 \nabla u_1) = f_1, & x \in \Omega_1, \\ \frac{\partial u_2}{\partial t} - \nabla \cdot (\nu_2 \nabla u_2) = f_2, & x \in \Omega_2, \\ u_1 = u_2, & x \in \Gamma, \\ \nu_1 \nabla u_1 \cdot \mathbf{n}_1 = \nu_2 \nabla u_2 \cdot \mathbf{n}_2, & x \in \Gamma, \\ \text{suitable boundary conditions on the other boundaries,} \\ \text{suitable initial conditions,} \end{cases}$$

where \mathbf{n}_i is the outward unit normal to Ω_i along Γ for $i = 1, 2$. This example arises in conjugate heat transfer, representing the exchange of heat between two abutting solids with different temperatures. It is a subsystem in the

exchange of heat between a solid and an encompassing fluid. This system also arises in other contexts, for example, the core-edge coupling in models of a fusion reactor.

Such a system may be posed as a large, fully implicit system of equations. However, it is common practice to use some form of iterative timestep method. The different components often use different spatial and temporal discretizations and may even be solved with different codes.

A common approach combines Algorithm 1 as an inner iteration in Algorithm 2 and is specified in Algorithm 4. The solves that are performed on each subdomain may use subcycling in time; recall that we use $u_1(t_n)$ and $u_2(t_n)$ to denote the nodal values of the numerical solutions at N_t time nodes $t_0 = 0 < t_1 < t_2 < \dots < t_{N_t}$.

Algorithm 4

```

Given initial values  $u_i(t_0)$ ,  $i = 1, 2$ :
for  $n = 1, \dots, N_t$  do
  Set  $u_i^0(t_n) = u_i(t_{n-1})$ ,  $i = 1, 2$ 
  for  $k = 1, 2, \dots$ , (until convergence) do
    Solve for  $u_1^k(t_n)$  with  $u_1^k(t_n) = u_2^{k-1}(t_n)$  on  $\Gamma$ 
    Solve for  $u_2^k(t_n)$  with  $\nu_2 \nabla u_2^k(t_n) \cdot \mathbf{n}_2 = \nu_1 \nabla u_1^k(t_n) \cdot \mathbf{n}_1$  on  $\Gamma$ 
  end for
  Set  $u_1(t_n) = u_1^k(t_n)$  and  $u_2(t_n) = u_2^k(t_n)$ 
end for

```

Of course, we must discretize in space to create full discretizations. Many possible variations of this approach exist.

The number of specified iterations in each step is an important factor. This approach can be interpreted as a fixed-point iteration; and if the fixed point iteration converges, then by iterating sufficiently, this discretization converges to the fully implicit discretization. On the other hand, this approach often is implemented with only one iteration. Doing so may have significant effects on accuracy. More worrisome is that if the fixed-point iteration actually fails to converge, using only one or a few iterations will fail to reveal this fact. The numerical results are questionable in this circumstance.

To illustrate some of the behavior that can arise in this approach, we consider simple linear heat equations with diffusion coefficients $\nu_1 = 1$ and $\nu_2 = 1/2$ on domains $\Omega_1 = [0, 1]$ and $\Omega_2 = [1, 2]$, respectively, coupled by imposing continuity of state and normal flux at $\Gamma_{12} = \{1\}$. The solution should converge to a steady state. We solve using Algorithm 4 with continuous second-order finite-element methods in space with $\Delta x = 0.01$ and implicit Euler with $\Delta t = 10^{-3}$ for each component. Each component is solved sufficiently accurately that the component errors are negligible.

We plot the total error of the approximation in Figure 20. We see the expected increase during the initial transient. But, the increase in error that occurs over longer time intervals is not expected, since the solution should be converging to a steady state. We can estimate the error that arises from incomplete iteration at each step (we use only one iteration) and the error in the numerical flux that is used in the solution of u_2 . Note that u_1 produces a numerical flux that is only first order. We see that the error from incomplete iteration is significant during the initial rapid transient but becomes unimportant over long time intervals. However, the error due to transferring a numerical flux increases as time passes, and becomes the dominant feature. In fact, the pollution from transferring the numerical flux causes a loss of order in spatial accuracy, as we show. Analyses of the errors, as well as examples in two dimensions with both matching and nonmatching grids along Γ , are given in [318, 319].

3.4.3. Solution of systems of elliptic equations with independent discretizations

In the next example, drawn from [320], we consider an even more nominally benign form of M³OD, which nonetheless produces interesting behavior. We consider a two-component, “one-way” coupled system.

$$\begin{cases} -\Delta u_1 = \sin(4\pi x) \sin(\pi y), & x \in \Omega \\ -\Delta u_2 = b \cdot \nabla u_1, & x \in \Omega, \\ u_1 = u_2 = 0, & x \in \partial\Omega, \end{cases} \quad b = \frac{2}{\pi} \begin{pmatrix} 25 \sin(4\pi x) \\ \sin(\pi x) \end{pmatrix}$$

3 ALGORITHMS FOR MULTIPHYSICS COUPLING

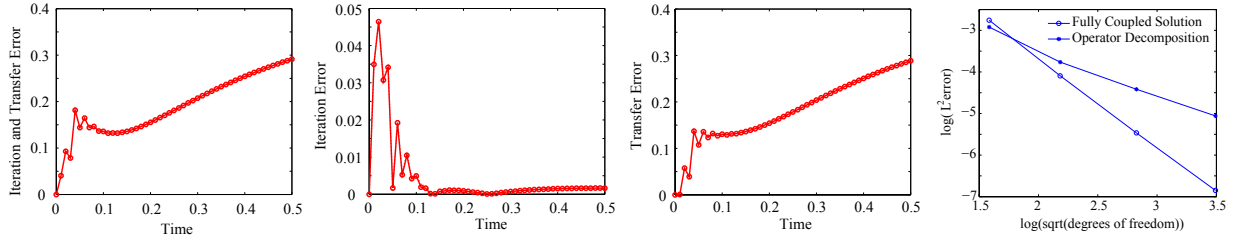


Figure 20: Results of the Gauss-Seidel algorithm for coupled parabolic problems. The component discretization errors are negligible. **Left:** the total error. **Second to left:** the contribution from incomplete iteration on each step. **Second to right:** the contribution from the passing numerical fluxes. **Right:** approximation showing loss of an order of accuracy in space as a result of the error in the numerical flux.

Here $\Omega = [0, 1] \times [0, 1]$. The “lower triangular” form of this system means that we can solve it as a coupled system or we can solve the first equation and then use the solution to generate the parameters for the second problem, as in Algorithm 5. In either case, we use independent discretizations of the two problems. Solutions obtained by using a

Algorithm 5

- Construct discretizations for the two equations
 - Compute the numerical solution u_1 of the first equation
 - Project u_1 onto the mesh for the second equation
 - Compute the numerical solution u_2 of the second equation
-

standard piecewise linear, continuous finite-element method on uniform, highly refined meshes and Algorithm 5 are shown in Figure 21.

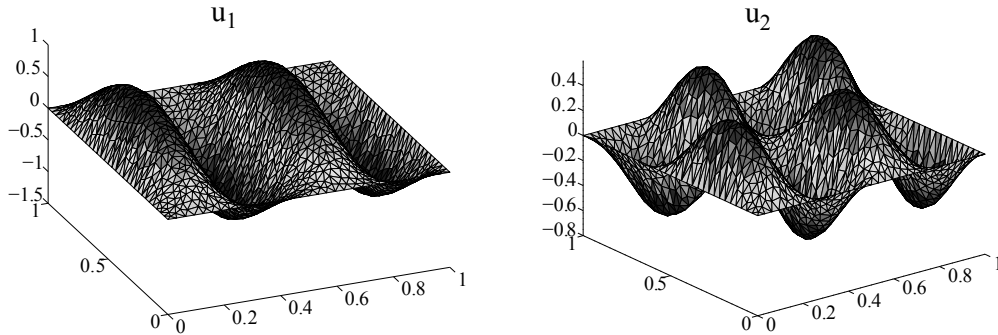


Figure 21: Solutions of the component problems of the elliptic system computed on uniform meshes.

Using uniform meshes, we evaluate the standard *a posteriori* error estimate for the second component problem and the quantity of interest $u_2(0.25, 0.25)$, ignoring any effect arising from error in the solution of the first component. We estimate the solution error to be ≈ 0.0042 . However, the true error is ≈ 0.0048 , giving a discrepancy of ≈ 0.0006 ($\approx 13\%$) in the estimate. This is significant for these kinds of accurate *a posteriori* error estimates since the solutions are smooth.

If we adapt the mesh for the solution of the second component based on the *a posteriori* error estimate of the error in that component while neglecting the effects of the decomposition (Figure 22), the discrepancy becomes alarmingly worse. For example, we can refine the mesh until the estimate of the error in the second component is ≈ 0.0001 . But, we find that the true error is ≈ 0.2244 .

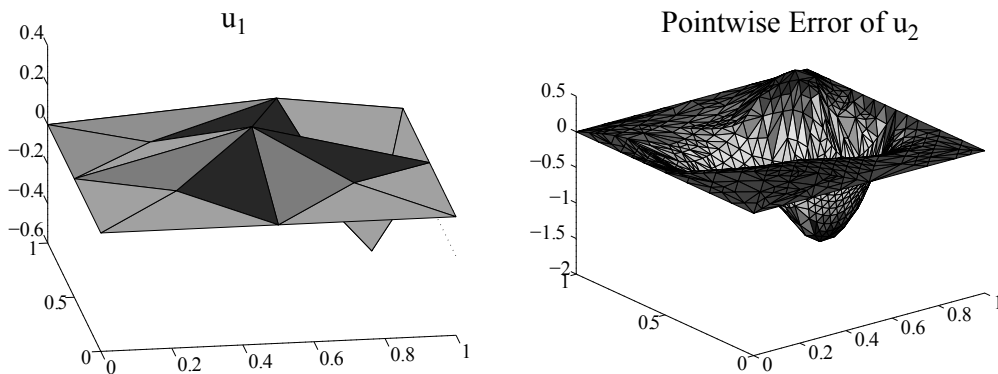


Figure 22: Results obtained after refining the mesh for the second component so that the *a posteriori* error estimate of the error only in the second component is less than 0.001. The mesh for the first component remains coarse; consequently, the error in the first component becomes relatively much larger.

3.5. Uncertainty Quantification

Uncertainty quantification (UQ) is a process for determining how random variation in model inputs, such as initial/boundary conditions or parameters in material properties or constitutive laws, affects the results of modeling and simulation. UQ is inherently important for using multiphysics models in engineering and scientific practice because of the complicated systems they attempt to model. In particular, multiphysics systems always present severe difficulties for experimental observation. Consequently, there is a heavy reliance on mathematical models to help supplement the meager availability of experimental data and observations. UQ can be used to characterize and reduce uncertainties in multiphysics applications as part of a process of continuous quantifiable improvement of simulation accuracy.

Random variation in model parameters can arise from a number of sources, including measurement error and natural variability. When sufficient data is available, this randomness can be characterized in the form of probability distributions, and the random parameter is treated stochastically. Such uncertainties are generally termed *aleatory*. On the other hand, *epistemic* uncertainties generally reflect a lack of knowledge, and probabilistic treatment is not always appropriate. In practice, the effects of epistemic uncertainty are sometimes explored by assuming a uniform distribution with suitable lower and upper bounds.

In either case, one of the principal concerns of UQ is to construct probability distributions of the model responses. Doing so requires multiple deterministic runs of the multiphysics model using samples drawn from the input distribution. Current UQ practice is dominated by “black-box,” nonintrusive strategies that require no modification of the application. Monte Carlo sampling and more efficient variations, such as Latin hypercube sampling, are most widely used. This approach can require thousands of deterministic realizations to obtain the output probability distribution, especially when a large number of uncertain parameters must be accounted for. This so-called curse of dimensionality can be mitigated by using spectral polynomial chaos expansions (PCE) [321]. While PCE can dramatically lower the cost of calculating the output probability distribution, it relies on a limited set of analytic input probability distributions for its approximation properties and requires a complete rewrite of the application. Stochastic collocation [322] is a nonintrusive alternative to PCE that is still much less expensive than Monte Carlo sampling but also relies on the same limited set of analytic input probability distributions for reliable uncertainty analysis.

While UQ is a relatively young field with many foundational research problems remaining to be tackled, UQ for multiphysics models is in its infancy. The complexity of a typical multiphysics model presents significant challenges. On the one hand, the complex behavior makes it difficult to obtain *informed, high-fidelity* statistical models of the system (when subject to random data and parameters), so that the black-box approach to UQ becomes problematic. On the other hand, propagation of random uncertainty through a multiphysics model encounters all the difficulties that have been outlined for the numerical solution. In addition, significant “stochastic feedback” often occurs between components of a multiphysics model.

4. Multiphysics Software

Software is the practical means through which collaboration on multiphysics algorithms and applications occurs; collaboration is unavoidable because the full scope of required functionality is broader than any single person or team can deeply understand. The following quote from Steve Jobs truly resonates for multiphysics [323]:

“The way you get programmer productivity is by eliminating lines of code you have to write.”

We could expand to say that the only way the multiphysics research community can address increasing challenges in modeling and algorithms is by making it possible to combine and leverage diverse code for the various phases of simulation. Software for multiphysics presents a host of difficulties beyond those faced in single physics contexts because of the compounded complexities of code interactions. Section 2.2 introduced issues from the perspective of scientific application teams, who face daunting challenges both in building application-specific infrastructure and in incorporating libraries, frameworks, and tools that are under development by many groups at universities and research laboratories. Section 4.1 expands on these issues with a closer look at the current status of multiphysics software, emphasizing capabilities needed for high-performance computing (HPC) environments. Section 4.2 discusses challenges in designing software for multiphysics algorithms as well as some possible approaches to help manage complexity and change. Section 4.3 considers practical difficulties in collaborative research software.

4.1. Status of Software for Multiphysics

We begin by surveying current multiphysics software practices. This information provides background for discussing common needs. We then present a few examples of HPC multiphysics software infrastructure.

4.1.1. Current practices

Current practices for multiphysics coupling in production codes, such those introduced in Section 2, must strike a balance among performance, software reuse, and numerical accuracy. Exactly which practices and approaches are favorable to each of these three goals is a subject of controversy and continued research. However, some common choices exist.

The importance of software reuse has strongly favored operator-splitting coupling methodologies, possibly at the expense of numerical accuracy. Typical workhorse codes in key physics areas such as computational fluid dynamics, particle transport, atomistic methods, and structural mechanics can easily have hundreds of developer-years behind them, including extensive test suites, comprehensive experience with the strengths and limitations of the underlying numerical methods, active user communities, and highly optimized kernels. Since most projects using these codes are driven by application goals (scientific or engineering discovery), any significant restructuring of the codes, let alone a complete rewrite, is extremely risky for application scientists, and often completely infeasible with current software technology. Moreover, many of these applications were not originally written with extensibility in mind and depend on outdated constructs and practices in software design. Thus, their integration into a more modern, modular, component-based software architecture is often more problematic than may initially appear.

Therefore, it is fair to say that the lowest threshold to early multiphysics results involves treating the coupling variables as input/output to these legacy codes. This approach allows some initial exploration into basic coupled phenomena with very limited modification of the original codes. We also note that higher than first-order accuracy may be achieved within this framework through various iterative techniques, as introduced in Section 3, albeit possibly at a higher performance cost than other choices.

While this approach sounds relatively straightforward, nontrivial implementation issues remain, particularly on HPC architectures, which beget another set of fundamental choices. These typically fall into two general classes.

The *service-oriented architecture* approach (e.g., the fusion IPS framework [324]), ensures the maximum independence of each code component by basing coupling on I/O rather than variable/grid-based data formats. The key idea is that interacting with each code at the I/O data level is more flexible and is independent of the internals of the codes to be coupled. Currently file I/O is commonly used for data transfer, but it will be necessary to move toward in-memory transfers for scalability. A standard for the meaning of the I/O format (e.g., which variables are located where on which mesh type) is still required for the coupling to be numerically consistent, but this approach avoids any direct “contact” with the internal data structures of each component. Note that, while not required, this approach

can be naturally carried out in a parallel environment when each component runs as its own executable on a disjoint set of MPI processes. This approach can be implemented on top of a higher-level I/O library.

In the *remote procedure call approach*, components share information by direct method calls based on a defined abstraction of the information that the components must share. This approach does not imply that the internal data representations of each code are directly modified to accommodate the coupling; rather, more typically, interface routines are used to transmit information between the components for loading into the internal representations of the communicating components. Clearly the more common the internal representations used by the components, the more efficient, in terms of memory and translation, the information communication process. This approach is more commonly followed with existing HPC codes and is more naturally implemented as a single MPI executable with either disjoint or overlapping processor configurations. When linking independently developed components with different internal data representations, this approach must employ a *lingua franca* perspective in which the information passed in the calls can be easily extracted from, and placed into, the internal data structures of the individual components. Such an approach has been adopted by the ITAPS SciDAC center for dealing with unstructured meshes and the field (solution) information defined over those meshes [325–327].

We note that there is no superior multiphysics software solution in general; the best approach for a particular circumstance depends on application-specific requirements. Usually, though, service-oriented architecture approaches will require much less development work but offer less flexibility and opportunity for performance optimization, in terms of both memory and execution time.

Increasingly, the commercial simulation community is endeavoring to provide “coupling toolkits” to ease the programming burden of building multiphysics simulations. Some of these toolkits, such as MpCCI [328], claim broad penetration and adoption in the commercial modeling world. Their basic philosophy is to allow users to specify their coupling problem at a high level of abstraction by providing the input/output variables and a description of the orchestration (time sequence of couplings). Such tools are focused on a loose coupling approach, providing the adapter code that can interpolate variables between different simulation packages. For this approach to work, the list of codes that can be coupled is hardwired into the toolkit—something that probably makes more sense in the commercial world where a larger fraction of the community uses a smaller number of codes, and thus it is feasible to support only the most popular ones. Limitations of this approach include little or no accommodation of parallel computing resources, particularly at the very high end, and lack of flexibility in potential candidate code components.

A number of projects have explored implicit coupling based on the Jacobian-free Newton-Krylov method, introduced in Section 3. This approach requires each of N_c code components to return a solution residual ($F_i(u)$, $i = 1, \dots, N_c$). The approach can be taken even one step further, with a broader goal of blending the best of both worlds by building in the option to employ an operator-splitting algorithm in cases where a tightly coupled approach is not needed and to employ an operator-splitting solution as a preconditioner when a fully implicit solution is deemed necessary. Section 4.1.3 discusses software to facilitate this approach.

For all these techniques, it is important at the implementation level to categorize the range of coupling types that one encounters across applications (see, e.g., [329, 330]). The frequency and volume of data movement depend greatly on these categories and go a long way in determining the difficulty of implementation or the other strategies described above.

4.1.2. Common needs

From this brief survey of current practices in production multiphysics codes, we identify common needs for HPC multiphysics software to achieve reasonable performance, robustness, portability, and flexibility, as well as to address the expectations of software developers and end users. These needs fall into several categories that are introduced briefly below and discussed in more depth in Section 4.2. Many of these needs also exist for good-quality HPC uniphysics software. However, while it may be possible to “get away with” some poor software engineering practices in a uniphysics context, addressing these issues is imperative for effective code sharing for multiphysics applications.

- *Interfaces with appropriate levels of encapsulation and granularity while permitting adequate performance.* In designing a multiphysics software package, it is critical to determine interfaces with appropriate levels of encapsulation. If too much is encapsulated, the interface is inflexible, and users cannot access needed data. If too little, objects can easily be used incorrectly, and functionality will not be great enough to encourage adoption. See Section 4.2.1.

- *Common infrastructure and ability to exploit commonality among different physics operators.* High-performance multiphysics software should also facilitate access to a broad set of common infrastructure that may be used by the physics codes being coupled to ease the coupling process. A clear example is mesh management [331–333], where the use of a common mesh and field database can allow multiple codes to access a mesh in a common way and to exchange data relating to that mesh. Common verification support, including order of accuracy verification for multiphysics simulations, would also be of great use. Other important types of infrastructure include coupling technology, such as parallel communication schedulers, interpolation, and mortar methods to exchange data on different meshes; nonlinear and timestepping solvers that support multiphysics algorithms; automatic differentiation technology for derivative computations; error estimation and uncertainty quantification software; and tools for data management, visualization, and I/O. A multiphysics software package might also help developers exploit commonality among different physics operators such as those in different types of convection-dominated systems. See Sections 4.2.2, through 4.2.4. A test-driven design that helps verify correct operation is important to manage the difficulties of incorporating such a broad range of tools that are developed by diverse multi-institutional teams and constantly changing. See Section 4.3.1.
- *Strong focus on meeting practical user expectations such as multilanguage support, ease of building, documentation, assurance of future support, and transparent open-source licensing.* In addition to functional requirements, multiphysics software must meet the practical needs of software developers and users if it is to be widely adopted. Some multiphysics software supports only C++, but Fortran and C interfaces are also necessary to accommodate the majority of physics codes, and Python is increasingly popular. Software must be sensitive to global scope issues (including symbol names and file usage). Multiphysics software must easily integrate with applications and must support all relevant hardware. Because many approaches to multiphysics software involve combining software from several different sources, clear interface specifications are needed. A related need is software documentation, including tutorial material and clear examples. See Section 4.3.
- *Extensibility to the largest current and anticipated future problems and to future extreme-scale systems.* Multiphysics approaches and algorithms that can be extended to very large problems on future extreme-scale architectures are required. We note that scalability by itself is not necessarily a useful measure of performance and that some very inefficient codes may exhibit relatively good scaling. For exascale architectures, memory use must be minimized. This circumstance may make it attractive to develop custom variants of implicit methods, and to run one type of physics on some nodes and other physics on other nodes, requiring careful attention to load-balancing issues. See Sections 4.2.5 and 5.1.

A longer-term and ambitious goal might be to allow domain experts who do not have extensive knowledge of multiphysics coupling and possibly have little advanced HPC expertise to write robust, coupled, parallel simulation codes that exhibit reasonably good performance.

4.1.3. Software successes

The development of multiphysics software is becoming a popular endeavor. Several commercial analysis packages contain the word “multiphysics” in their description, and Section 2 mentions some application-specific research software. Below we list a few frameworks, libraries, and other tools that incorporate *general-purpose, reusable algorithmic infrastructure in support of parallel PDE-based multiphysics applications* and aim to address the needs identified above. Of course, many additional projects address important aspects of multiphysics simulations. Such projects include approaches that facilitate working at higher levels of abstraction [334–338], as well as software such as Cactus [339], Coupler [340], DDEMA [341], FLASH [342], InterComm [343], Jade [344], Karma [345], MCT [346], MUSE [347], Overlink [348], Overture [349], PreCICE [350], Sierra [351], and XpressSpace [352].

BoxLib. BoxLib [353] is a software framework designed to facilitate the development of parallel, block-structured adaptive mesh refinement (AMR) applications. The block-structured approach to AMR used in BoxLib is based on representing the solution on a sequence of nested grids with successively finer resolution. The data at a given level of resolution is organized into a collection of logically rectangular grid patches each containing a large number of points. The aggregation of grid patches at a level, referred to as a MultiFab, is the fundamental parallel abstraction used in BoxLib. Coarse-grained parallelism is expressed by distribution of the grid patches in a MultiFab to nodes of the system. Operations at the node level are then specified in terms of iterators with an owner-computes rule. BoxLib supports a hybrid parallelization model in which OpenMP directives are used within a node to thread operations on

a grid patch. A number of applications have been implemented within the BoxLib framework. These include LMC for low Mach number combustion [354], CASTRO for radiation hydrodynamics [355, 356], MAESTRO for low Mach number astrophysical convection [357–361], PMAMR for subsurface flow [362] and NYX for computational cosmology [363].

Chombo. The Chombo library [364] provides a set of tools for implementing finite-volume methods for the solution of partial differential equations on block-structured, adaptively refined rectangular grids. Both elliptic and time-dependent modules are included. Chombo supports calculations in complex geometries with both embedded boundaries and mapped grids, as well as coupling with fast particle methods. Parallelization is currently achieved through a highly scalable domain decomposition strategy involving the distribution of rectangular grid blocks across processors [365]. Chombo has been used in a variety of multiphysics contexts, including biological flow, compressible gas dynamics, hydrology [366], ice sheet modeling [367], and plasma physics [368]. Of particular use in these multiphysics simulations is Chombo’s open source framework for temporal and spatial discretizations, which enables the scalable integration of multiple physical models, including external solvers, specialized physics codes, and custom algorithms.

Illinois Rocstar. Illinois Rocstar [369] is a general-purpose numerical solver for fully coupled, time-dependent fluid, structure, and combustion problems that grew out of the work in the Center for Simulation of Advanced Rockets (CSAR) [370]. Rocstar was designed from the beginning to be fully parallel and is routinely run on systems with over 1,000 processors. The Rocstar framework is built around Roccom, which manages the communication between the components. These include modules for timestepping, mesh modification, I/O, and performance profiling. While originally developed for solid propellant rocket motors, Rocstar has been applied to a variety of different problems, including simulations of the noise from helicopter blades and the effect of wind on tall buildings.

Escript/Finley. Escript [371, 372] is a Python-based programming tool for mathematical modeling based on nonlinear, time-dependent PDEs. Escript has been designed to give modelers an easy-to-use environment for developing and running complex and coupled models without accessing the underlying data structures directly. Escript employs the Finley finite-element method solver library. The code has been efficiently parallelized using MPI, OpenMP, and hybrid modes. Applications include modeling Earth mantle convection, earthquakes, porous media flow, reactive transport, plate subduction, and tsunamis.

LIME. The Lightweight Integrating Multiphysics Environment (LIME) is a software package for coupling multiphysics simulation codes [373, 374]. It is used when application codes for solving different parts of a multiphysics problem need to be coupled together in either a loose or a tight sense. LIME supports a number of coupling modes, ranging from simple, nonlinear block Gauss-Seidel and nonlinear block Jacobi to fully coupled Newton-based methods. In addition to the solution drivers, LIME provides an abstract interface to application physics codes based on the `Thyra::ModelEvaluator` in Trilinos. LIME has been successfully used on the BRISC project [373], to integrate six separate codes to assemble a virtual nuclear reactor model; the various couplings in the aggregate system include nonlinear block Gauss-Seidel, nonlinear Schur (nonlinear elimination), and Jacobian-free Newton-Krylov. LIME is currently used in the Consortium for Advanced Simulation of Light Water Reactors (CASL) Project [375].

MOAB. MOAB [56–58] is a library for query and modification of structured and unstructured mesh and field data associated with the mesh. MOAB provides a solution transfer capability for multiphysics approaches in which different physics are solved over the same volumetric domain, with each physics implemented on its own mesh and no assumptions made about mesh nesting or the two physics using the same element type or discretization. MOAB’s solution transfer tool [376] breaks the process into four stages: initialization, point location, interpolation, and conservation/normalization. During the initialization stage, the two meshes are instantiated in MOAB’s parallel representation, and spatial searching trees on the source mesh are constructed. Point location requires finding the source element and parametric location in that element of each target point (and the processor on which that element resides). Interpolation maps the source field to the target points, and conservation/normalization imposes global integral constraints on the mapped field. This capability is the basis for solution transfer being developed as part of the SHARP application discussed in Section 2.1.3 and also has been used to transfer radiation heat deposition results from the MCNP code [377] to a tetrahedral mesh for heat transfer/fluid flow computations [378]. Additional MOAB applications include Nek5000 [55], the Denovo neutron transport code [379], and an ice sheet modeling code [380, 381].

MOOSE. The Multiphysics Object-Oriented Simulation Environment (MOOSE) [48, 382] focuses on tightly coupled multiphysics applications using Jacobian-free Newton-Krylov methods and is centered on use of the residual abstraction, given by Equation (3). MOOSE uses a layered approach to provide core services in support of multiphysics applications. The heart of MOOSE is the Kernel abstraction, which represents a physics equation or term in an equation and provides core functionality for residual and Jacobian evaluation. One may consider a Kernel to be a mathematical operator, such as a Laplacian or a convection term in a PDE; Kernels may be swapped or coupled together to achieve different application goals. MOOSE employs the libMesh finite-element framework [383], which provides utilities for parallel finite-element computation, including mesh I/O, a finite-element library, and interfaces to nonlinear solvers in PETSc and Trilinos. Multiphysics applications using MOOSE include BISON [41, 42, 384] (models nuclear reactor fuel performance computation across diverse scales that involve both the continuum and sub-continuum [45, 46]); PRONGHORN (an analysis code that models coupled coolant flow and neutronics for very high temperature gas cooled reactors [385–388]); and geoscience analysis applications [389, 390].

OpenPALM. The main idea of the OpenPALM philosophy is to create applications by assembling independent components in sequences and letting them exchange data [391, 392]. The PALM framework performs the scheduling and synchronization of the components and the data exchanges. A PALM application can be described as a set of computational units arranged in a coupling algorithm. The different units are controlled by conditional and iterative constructs and belong to algorithmic sequences called computational branches. Inside a branch, the coupled independent programs are invoked as if they were subroutines of the branch program. The main features of PALM are the dynamic launching of the coupled components, the full independence of the components from the application algorithm, the parallel data exchanges with redistribution, and the separation of the physics from the algebraic manipulations performed by the PALM algebra toolbox. PALM multiphysics applications include oceanography, climate modeling, computational fluid dynamics, combustor cooling system optimization, and hydrology [393].

PETSc. The Portable Extensible Toolkit for Scientific computing (PETSc) [7, 8] incorporates generic parallel programming primitives, matrix and vector interfaces, and solvers for linear, nonlinear, and transient problems. PETSc emphasizes ease of experimentation so that every component has a plug-in architecture by which the user can provide first-class implementations. The strong encapsulation in this design facilitates runtime composition of hierarchical methods for coupled problems. Several relatively recent components have substantially improved composability for multiphysics and multilevel methods. `DMComposite`, a distribution manager for coupled problems, handles the algebraic aspects of “gluing” together function spaces, decomposing them for residual evaluation, and setting up matrices to act on those coupled function spaces. `DMComposite` contains support for the `Nest` matrix format, which stores the submatrix associated with each physics component independently. A common matrix assembly interface is available, such that individual physics modules can assemble parts of global matrices without needing global knowledge or specializing on matrix format (`Nest` or monolithic). The `FieldSplit` preconditioner solves linear block systems using either block relaxation or approximate block factorization (as in “physics-based” preconditioners for stiff waves [265] or block preconditioners for incompressible flow [394]); `FieldSplit` can be nested inside other preconditioners, including geometric or algebraic multigrid, with construction of the hierarchy and other algorithmic choices exposed as runtime options. Recent enhancements to the TS component support IMEX schemes such as introduced in Section 3.2.2. Multiphysics applications using these capabilities include lithosphere dynamics [395], subduction and mantle convection [396–398], ice sheet dynamics [399, 400], subsurface reactive flow [401], tokamak fusion [85–87, 402], and power networks [403].

Trilinos. Trilinos [11, 12] incorporates algorithms and enabling technologies within an object-oriented software framework for the solution of large-scale multiphysics problems. A unique design feature of Trilinos is its focus on stand-alone, light-weight packages for scalable linear algebra, linear and eigen solvers, nonlinear analysis tools (nonlinear solvers, optimization, stability, bifurcation, uncertainty quantification), I/O, discretizations, meshing, geometry, and load balancing. A key multiphysics capability is Thyra, an abstract operator/vector layer that can be used to build composite or blocked systems with contributions from different application codes. These block operators can be passed to Trilinos solvers, such as the NOX nonlinear solver for applications requiring Newton-based algorithms. Leveraging the blocked operator approach, the Teko package can be used to construct block-based and physics-based preconditioners. This capability has been demonstrated on CFD/MHD systems [404]. Additional multiphysics-related

support includes Moertel [49], which supplies capabilities for nonconforming mesh tying and contact formulations in two and three dimensions using mortar methods, and Phalanx [405], which provides a local field evaluation kernel that decomposes a complex problem into a number of simpler problems with managed dependencies [406]. Multiphysics applications using Trilinos include climate modeling, reactor fuel performance modeling [44], and Charon, which has been used to model magnetohydrodynamics [407], multiphase chemically reacting flows [408], and semiconductor drift diffusion modeling [409].

Uintah. A good example of a production-strength multiphysics framework is Uintah [410, 411], a set of software components and libraries that facilitate the solution of partial differential equations on structured adaptive mesh refinement grids using hundreds to thousands of processors [412]. Uintah employs an abstract task graph representation to describe computation and communication. It can integrate multiple simulation components, analyze the dependencies and communication patterns between them, and efficiently execute the resulting multiphysics simulation. The underlying methods inside Uintah are a combination of standard fluid-flow methods and material point (particle) methods. Uintah was developed within the Center for Simulation of Accidental Fires and Explosions [413] to solve reacting fluid-structure problems involving large deformations and fragmentation.

waLBerla. The Widely Applicable Lattice-Boltzmann from Erlangen (waLBerla) [414] is a massively parallel framework for simulating multiphase flows based on the lattice-Boltzmann method. The flow simulator can be coupled to a rigid body dynamics solver (PE, physics engine) [415] using an explicit fluid-structure interaction technique, with the goal of, for example, simulating particulate flows or suspensions as they appear in many technical, environmental, or biophysical applications. A highlight of the package is its careful performance-optimized design and good parallel scalability that has been demonstrated on more than 250,000 processor cores [416]. Additionally, the mesoscopic computational fluid dynamics method can be coupled to finite-element-based PDE solvers in the waLBerla package.

4.2. Challenges in Multiphysics Software Design

While current multiphysics applications (see Section 2.1) and software (see Section 4.1) have already demonstrated substantial advances, numerous software engineering challenges remain, providing opportunities for long-term research focusing on stable, accurate, robust, efficient, and scalable multiphysics algorithms, with easy extensibility for application-specific customization. In Sections 4.2.1 through 4.2.6 we discuss challenges in multiphysics software design and suggest some best practices. Related issues are discussed in [417, 418].

4.2.1. Enabling introduction of new models, algorithms, and data structures

Multiphysics software attempts to balance the competing goals of interface stability and software reuse with the ability to innovate algorithmically and develop new physical models. Interface stability is important for several reasons:

- Decoupling the software development process, which a single person or single team cannot manage alone;
- Enabling reuse of model components for different kinds of analysis, for example, transient, steady-state, sensitivity analysis, optimization, and uncertainty quantification;
- Enabling use of each “physics module” for stand-alone analysis and in various coupled configurations.

A related point is the need over time to re-examine interfaces in light of recent advances in mathematics and programming models. For example, providing (possibly optional) access to additional information (e.g., derivatives, not just fields, or other internal state) can inform library design for both couplers and solvers.

We discuss two important aspects of software design that promote interface stability and facilitate experimentation with existing approaches as well as the introduction of new capabilities over the lifetimes of multiphysics codes:

1. Library interfaces that are independent of physical processes and separate from choices of algorithms and data structures, and
2. Abstractions for mathematical objects such as vectors and matrices, which enable dealing with composite operators and changes in architecture.

Library interfaces. An important provision for extensibility is the separation of logically (or mathematically) distinct concerns into modules with *library* interfaces. Because a library cannot make assumptions about program startup or the use of state, it cannot seize control of “main” or assume `MPI_COMM_WORLD`; see Section 4.2.6 for additional software engineering issues. When preparing an application for use in a multiphysics context, a clear division should be made between the software components responsible for discretization of a physical model and the analysis being performed with the discrete model. The single-physics “analysis” component will typically be bypassed in a multiphysics context, so the interfaces between these components should be thought out carefully, especially with respect to ownership and hidden state.

A key concept is that any state data must be explicitly exchanged through an interface to maintain consistency; examples include common algebraic interfaces, such as vector and matrix classes available via PETSc and Trilinos. This approach may be used to specify inhomogeneous boundary conditions that would otherwise often be computed in a preprocessing stage. It is usually not critical which classes are used by a coupling interface, especially for vectors, because the memory backing a vector from one package can be wrapped as a vector from a different package without a copy (though GPU interaction and dynamic adaptivity can cause complications).

Interfaces that are independent of physical process. Whenever possible, interfaces should be independent of the physical process and discretization in use; otherwise, they will prevent reuse of analysis/solver components and be brittle when different physical processes or discretizations are employed. It is still important, however, to provide enough “hooks” to implement the desired algorithms for a specific problem and to avoid excessive overhead (in terms of memory or time).

A representative example is the choice of whether to expose a semi-discrete or fully discrete interface for time integration. In the former case, one assumes that a method-of-lines approach is used and writes $\partial_t u = f(u)$, as given in coupled form by Equations (2a–2b), where the vector u contains all state variables. If implicit methods are to be used, the semi-discrete form may also provide derivative information, such as approximations to $\partial f/\partial u$ and preconditioners. Responsibility for time stepping is relinquished by the physics module, enabling reuse of software for time integration schemes, including steady-state solvers and sensitivity or stability analysis methods, without modification or special support by the physics module. This approach is preferable from a software modularity perspective, but it has limitations—namely, little or no support for dynamically adaptive meshes and smaller, stable timesteps (for some transport discretizations). For multiphysics simulation where software reuse and algorithmic flexibility in coupling and analysis algorithms are often necessary, the semi-discrete form is usually desirable.

Interfaces that separate from choices of algorithms and data structures. Different data structures are better suited to different solution strategies, especially when implicit solvers are used as part of a multiphysics simulation. The best solution strategy is often not known *a priori* or depends on the problem and computer architecture. For example, coupled problems can be solved by using monolithic methods involving a single sparse matrix representing all components or by split methods in which each component manages its own linear algebra. A common hybrid approach is to adopt a global Newton-Krylov method (see Section 3.1), with the preconditioners involving splitting and perhaps approximate Schur complements. This case requires the action of operators on the global system, but most of the work is done preconditioning single-physics and auxiliary (e.g., approximate Schur complement) systems. When multigrid methods are used, the splitting can be performed at the top level with independent single-physics multigrid, or multigrid can be applied directly to the global problem (often with splitting used to define the smoother) [394, 419–422].

In the mathematical (algebraic) description of the solver, the distinction between these approaches involves a different order of operations for coarsening and extracting submatrices. Monolithic domain decomposition methods use different submatrices (e.g., corresponding to overlapping Dirichlet problems, non-overlapping Neumann problems, or combinations [201, 202, 423, 424]). When a suitable preconditioning algorithm is used, the data structures should exploit single-physics structure such as a constant block size or symmetry for memory and floating-point performance. While hierarchical or nested matrix formats generally require more messages to perform an operation with top-level semantics, they can also be more latency tolerant and involve operations on smaller communicators, which may be beneficial for performance.

High-level abstractions. As discussed above and in Section 4.1.3, mathematical abstractions from the fields of linear and nonlinear solution methods are effective software design elements for multiphysics. Also proving useful is template-based generic programming [425, 426] and related work on even higher-level abstractions, including varia-

tional forms available via FEniCS [427], Sundance [428], and Comsol [429]. Such abstractions enable higher-level library components (such as linear, nonlinear, and timestepping solvers) to be written as operations on mathematical objects (such as vectors and matrices), thereby minimizing the cost of developing software for each new architecture (and its variations) and enabling composite operators.

Dealing with changes in architecture. The end of frequency scaling for microprocessors is forcing significant changes in computer hardware to meet ever-increasing needs for performance. Some applications are exploring the use of GPUs; in the future, all high-performance computation applications will need to adapt to a range of new architectures. Mathematical objects such as vectors and matrices enable specific implementations to be employed for underlying use of GPUs, OpenMP, and so forth. Also useful in this context are guided code generators (see, e.g., [430]), where a mathematical representation is combined with a tool designed for particular operations.

Composite operators. Particular to multiphysics coupling is the ability to build block composite operators that aggregate physics applications for fully coupled solution techniques. For example, one could build a composite abstract vector from two physics applications that each supply their own residual vector. Generic tools for building composite block operations take as input abstract operators and produce new abstract operators that can be handed off to solvers that employ these.

Also useful is the approach of combining kernels of code to form terms and equations that contribute to the overall multiphysics system residual. These kernels also can supply an approximate Jacobian contribution to be assembled by the coordinating framework or library to be used for preconditioning operations. For example, in MOOSE, these kernels are usually composed of volume and surface integral terms from a variational form of a governing equation, written in residual form.

Another layer of abstraction is the interface to an application that a multiphysics driver requires to support various coupling techniques. Such interfaces can be difficult to define because of the need for balancing design requirements, including flexibility to support solution algorithm requirements, efficiency, and ease of integration. To this end, LIME [374] (see Section 4.1.3) provides a high-level domain model for multiphysics coupling, which supports the CASL project [375]. This domain model describes the analysis in terms of residuals,

$$F_i(\dot{u}_i, u_i, \{p_{i,l}\}, \{z_{i,k}\}, t) = 0,$$

transfer operators,

$$z_{i,k} = r_{i,k}(\{u_m\}, \{p_{m,n}\}),$$

and responses,

$$g_j(\{\dot{u}_i\}, \{u_i\}, \{z_{i,k}\}, \{p_{m,n}\}, t) = 0.$$

Here, F_i is the residual vector for physics i ; u_i is the vector of unknowns corresponding to physics i ; $\dot{u}_i = \partial u / \partial t$ is the vector of derivatives of unknowns corresponding to physics i with respect to time; $\{p_{i,l}\}$ is the set of all independent parameters for physics i ; $\{z_{i,k}\}$ is the set of coupling parameters for physics i ; and $r_{i,k}$ is the transfer function k for physics i , where physics i has $k = 1 \dots n_{i,k}$ transfer functions. A complete implementation of the domain model above should support all multiphysics coupling algorithms, ranging from loose to tight coupling. Based on the model, if less information is available from an application, the above interface is contracted, thereby limiting the available application coupling algorithms. Five forms of reduced domain models are identified [374]: the coupling elimination model $p \rightarrow g(p)$, state elimination model $p \rightarrow u(p)$, fully implicit timestep model $f(u, p) = 0$, transient explicitly defined ODE model $\dot{u} = f(u, p, t)$, and transient fully implicit DAE model $f(\dot{u}, u, p, t) = 0$.

4.2.2. Sharing methods and codes among application fields

One of the primary challenges with software development for multiphysics is sharing and integrating algorithms across coupled applications. As codes are integrated into a multiphysics framework, each application may bring in a number of required library dependencies. The reliance on such a large number of libraries can cause difficulties, since software engineering standards and portability can vary dramatically from code to code. Unifying or standardizing common software around libraries with exceptional software practices (domain-driven design, object-oriented design, test-driven design, coverage testing, etc.) and portability will be critical from a maintenance standpoint. Areas where sharing of tools and methods is important include meshing, solvers, data mapping, visualization, I/O, interpolation/mortar methods, and coupling libraries.

The current practice is to develop common components that are released as a framework or library. Two paths are typically followed. The first path is the monolithic Framework (Framework with a capital “F”; see Appendix A), which directs and controls all aspects of the problem, including memory management. This approach tends to allow rapid development, since all code is written to concrete implementations and much of the needed utility code is in place. However, the approach tends to lack flexibility because it forces applications to adopt the framework’s data structures and algorithms. The current trend is to move away from monolithic frameworks, since multiphysics coupling requires that applications “play well together”—they cannot insist on being in control of the overall coupled simulation. The second path is to develop frameworks that collaborate (framework with a lower case “f”; see Appendix A). These typically code to interfaces to allow for maximum flexibility. This approach allows users to provide their own concrete implementations of data structures at the cost of being somewhat more complex.

Data Layout. Developing data layout standards is a challenging problem. Optimal data structures can vary for different applications and for different hardware architectures (especially in light of GPU and exascale efforts). In addition, requirements for data structures for distributed computation can be vastly different from requirements for structures for local computation. While supplying a common data structure or interface is convenient from the standpoints of development and coupling, this approach can result in a compromise on performance—a tradeoff among flexibility, portability, and efficiency. Moreover, the integration with legacy codes where little knowledge or testing is in place may prohibit the adoption of new data structures.

Common infrastructure. From a software engineering standpoint, leveraging a common infrastructure for source code management, build, and test environments is challenging for multiphysics integrated systems. Unifying all application codes under a single build/test system can be difficult and time consuming. Further complexity is added if the application must support stand-alone builds outside the coupling framework. This could lead to inefficient resource allocation in maintaining multiple build systems. An example of a unifying build/test environment is the CASL project [375]. The approach taken with CASL was that if an application was under active development, then it was transitioned to natively use the build system and dependency tracking of the coupling framework, thereby easing integration issues.

Commonality in physics operators. Another issue is how to exploit commonalities among different physics operators. For example, convection-dominated systems can appear in a number of fields. Sharing both the mathematical methods and the software infrastructure for these types of problems can be difficult. An open research issue is how to decompose the operators. Additionally, as the nature of one physics operator changes, how does this affect the performance of the coupling and solution strategies? An example of leveraging similar operators is the Charon code [407] applied to both semiconductor device physics and magnetohydrodynamics. In this case, both systems were used for convection dominated flows. Not only were the fully implicit solution strategies and mesh database shared but also the stabilization algorithms.

4.2.3. Relating multiphysics spatial discretizations

Two general strategies are available for spatial integration within multiphysics applications, where the multiple physics models each apply across the entire domain [431, 432]. The first strategy is based on the use of a single mesh that spans the entire domain. Each physics model is hosted in terms of the basic operators supported on the mesh. The second strategy uses multiple meshes, in the extreme using a unique mesh for each physical model present in the code. This strategy is common when multiple codes are coupled; each code tends to use a mesh tailored to its requirements.

The use of a single mesh is straightforward in concept. However, this is almost always a compromise, since each physical model rarely has similar refinement requirements at a given location in the domain. If a single mesh is used, the mesh must be fine enough everywhere to support the most demanding physical model at that location. Moreover, the developers of the composite application must agree on a meshing strategy that meets the quality requirements and discretization assumptions of all concerned. On the other hand, if these issues can be overcome, this approach has several advantages. First, a single mesh incurs no error in transferring (or interpolating) data between models, since they all are hosted on a single mesh. Second, a single mesh is easier to implement, decompose for parallel execution, and load balance during a transient simulation.

The use of multiple meshes to host a multiphysics problem has nearly opposite considerations. Creating multiple meshes typically magnifies the cost, in person time, of running the application. In some cases, the problem domain may be sufficiently complex that generation of a single mesh is so expensive that the use of multiple meshes is not cost effective. In addition, as discussed in Section 3.2.1, care must be taken in communicating (transferring) the data between meshes and models in order to avoid sacrificing the accuracy of the final result [232, 433].

One option aimed at maintaining the advantages for both approaches is a multimesh assembly strategy [432] to achieve the spatial diversity advantages of a multiple mesh approach without sacrificing the accuracy of the single mesh approach. For example, Hermes code [434] is used to generate multiple hierarchically related meshes using *hp*-FEM, where *hp*-FEM provides for general adaptation and high-order convergence across the composite multiphysics problem. Each of these meshes hosts a physical model that is appropriately fine (in the *hp* sense) for desired accuracy of each model. The composite multiphysics model is also refined to meet a global accuracy requirement. The *hp*-FEM assembly process then integrates across the different models and meshes to provide a coupled solution without using data transfer. Thus, neither data transfer error or multiphysics coupling error is present in this approach.

In those cases where the individual physics analysis components being linked have independent meshes and internal data representations, a general *remote procedure call* approach must be taken. Key to the success of this approach is the abstractions used to define the mesh, solution parameters, and field information, defined over the entities in the mesh. Since the mesh is a discrete representation of the domain of the simulation, it is also desirable that the mesh entities maintain their relationship to the overall definition of the domain. Such relationships are particularly important when adaptive mesh control methods are used and the overall definition of the domain is a solid model as defined in a CAD system. The field information, which defines the parameters in the mathematical equations being solved, is best defined in terms of tensor quantities defined over the appropriate geometric entities. In cases where a high-level domain representation is used, the input parameters (loads, material properties, and boundary conditions) can be defined as distributions over the geometric model entities [435]. As part of the discretization process, these fields are mapped onto the mesh. The parameters that are solved for will be defined over the mesh and have known distributions as defined by the equation discretization processes. Maintaining knowledge of the distribution information is key for being able to properly map solution parameters between meshes in a multiphysics simulation.

The fundamental differences in representations used for *structured* and *unstructured* meshes indicate that the abstractions best suited for interacting with them are different. In the case of structured meshes, the two key components that must be addressed are the representation of the individual mesh blocks and the interactions between those blocks. In all cases, a compact representation of the mesh blocks in terms of a limited number of spacing and geometric mapping parameters is needed. The representation of the interactions between blocks varies strongly among the various methods (matched mapped blocks, overlapping Cartesian blocks, cut cells, etc.); thus, a separate abstraction tends to be used for each representation.

With unstructured meshes, topological entities from 0 to 3 dimensions, and the adjacencies of those entities [436, 437], define an ideal abstraction for defining the mesh as well and maintaining the relationship of the mesh entities to the geometric domain that is typically defined as a boundary representation. The various ITAPS iMesh implementations [325–327], and others, are based on such a topological abstraction of the mesh. Although it is possible to base the representations used on only the specific mesh entities and adjacencies used in the linkage of specific components, many implementations employ a complete representation in the sense that any of the desired 12 mesh adjacencies can be determined in $O(1)$ time (i.e., are not a function of the number of mesh entities) [436–438].

The parallel execution of multiphysics solution transfer requires the extension of the mesh representations that maintain appropriate linkages between the entities in the mesh as they are distributed over the memories of the parallel processors. With structured meshes, the mechanisms for doing this are a strong function of the methods used to understand the interactions between the mesh blocks. For example, the parallel implementation of overlapping mesh blocks will often have one or more individual mesh blocks on a process along with information on the mesh blocks on other processes with which those blocks interact, whereas with a tree-based decomposition a parallel tree representation might be used (e.g., a parallel octree). In the case of unstructured meshes, the common method of parallel distribution employs mesh patches of neighboring elements defined such that the numbers of mesh entities that share common boundaries with mesh entities on other processes is as small as possible. In this case the key component of parallelization is the addition of copies of the common bounding mesh entities on all neighboring processes with links between them [438]. A mechanism that can be employed to potentially support efficient parallel solution transfer is to ghost mesh entities, and their field data, from one process onto neighboring processes that happen to need that

information to support solution transfer operations. Since only one process is allowed to change the field data, specific care must be applied to ghosted information to be sure it is up to date on any process that will be using it.

4.2.4. Timestep control

Selection of an appropriate timestep for accuracy in a multiphysics calculation can be challenging. Software features with utility for multiphysics calculations span the needs of uniphysics codes, such as robust capabilities for checkpoint and restart and options for timestep control (e.g., manual, programmed, semi-automatic, and automatic timestep selection). When one adds multiple physics models, perhaps in multiple loosely linked modules or separate codes, timestep control across the aggregate application can become very involved. For example, in fission reactor fuel performance simulations, introduced in Section 2.1.2, the timescales of a macroscale structural mechanics model and a coupled mesoscale material model differ by over six orders of magnitude [46]. The mesoscale timescale cannot be stepped over, since the evolution of the material properties definitely affects the macroscale simulation. Taking picosecond-size timesteps is not practical for a fuel performance application. Thus one may need to use high-order time integration to support larger mesoscale timesteps, along with subcycling the mesoscale physics. In general, robust and accurate approaches to handle coupled problems with this level of time disparity are still being sought.

In a composite application, one might consider that a timestep is composed of three phases: initializing each package/module in preparation of taking the step, taking the timestep, and then finalizing the step. With multiple modules, the second phase might entail appreciable subiteration between modules, data and results transfer, internal model state updates, and so forth. The third phase includes checking result consistency across the modules, convergence behavior, error metrics, module intercommunication. The “failure” of a timestep might indicate the inability of subiteration between closely linked modules to converge, or some other error condition. In this case, logic would exist to reduce the timestep or otherwise recover from the error and retake the step. Related algorithmic issues are discussed in Section 3.2.2.

Unfortunately, automatic timestep selection is not robust for general multiphysics simulations and requires tuning for each problem class of interest. Analyst intuition and experience lead to a good understanding of the dynamical timescale of the problem of interest and are valuable in formulating effective approaches and “tuning” them for the application. This observation leads to one of the more effective timestep selection approaches. It is often effective to provide a general interface to a multiphysics code for the manual specification of a timestep history curve that the code should follow, in conjunction with providing residual information as postexecution feedback to the analyst. The analysis can then tune the timestep strategy to provide the desired accuracy and overall simulation behavior, given a few trial simulations.

4.2.5. Achieving performance

While performance is often a major issue in computational science codes, few applications have a quantitative estimate of the achievable performance; hence, they do not know whether the code is using the computing system efficiently. Gaining this understanding is important in judging the impact of different implementation choices and, in particular, the performance tradeoffs with different approaches to connecting multiphysics components together.

Achieving this level of understanding requires constructing a performance model good enough to provide an upper bound on performance. Such models are often a combination of instruction, memory reference, and floating-point performance and can identify the critical performance limiting resource (e.g., [218]). The “roofline” model, which simply looks at the performance limits due to a number of constraints (each contributing a line to the “roofline”), is a recent representation of this idea.

4.2.6. Software engineering issues for multiphysics integration

Previous sections have discussed some design issues for introducing new models and algorithms into multiphysics applications. In preparing to incorporate an application into a multiphysics environment, a number of issues must be addressed; the following simple suggestions can ease the integration process.

- *Program startup:* A general multiphysics library must not make assumptions about startup. It cannot declare the `main` routine, as this will most likely be declared by the driver or coupling framework. Applications expected to be used in multiphysics coupling should be written as a library; a stand-alone version of the application can simply wrap a driver (containing `main`) around the library. In addition, applications should not expect to be able

to directly parse input from the command line; multiple codes might interpret the same command-line argument for different purposes. Each application code should define an interface for accepting specific input parameters. The coordinating software (framework or library) can parse the command line and hand off specific data to each application.

- *MPI communicators*: Each application should accept an MPI communicator during startup through an interface. An application code should never directly code to `MPI_COMM_WORLD`. At startup a common multiphysics parallel decomposition pattern is to hand off blocks of processes (i.e., a subcommunicator) to each application. Creating a communicator on the fly by using `MPI_COMM_WORLD` inside an application will not work and will hang the coupled program.
- *Global scoping*: The naming of macros, variables, and functions in a code must be protected either by name-spacing or by prefixing. In our experience, almost every code that did not take care to apply such protections and polluted the global namespace required extensive refactoring because of collisions with other applications.
- *Third-party library dependencies*: Source code for third-party libraries should not be included in an application. If multiple coupled applications required the same third-party library and if the library were included with each, multiple instantiations would be present. While some compilers can link in this circumstance, others fail with an error if duplicates are present. Additional complexity arises if coupled application codes require different versions of the same library.
- *Output support*: All output for an application code should be allowed to be redirected to its own ostream/files. If two coupled applications are running concurrently and both dump output to the screen, the output from the applications will be mixed together, making the output unreadable.
- *Separation of algorithms from application physics*: When preparing an application for use in a multiphysics context, a clear division should be made between the software components responsible for discretization of a physical model and the analysis being performed with the discrete model; Section 4.2.1 discusses further details.
- *Level of granularity*: The extent to which data is exposed to the coupling framework will affect the choice in solution algorithms that can be applied to the coupled system; see Section 3.1.2. A code that cannot provide a residual or even expose a solution vector will have to be run as a black box, limiting couplings to nonlinear block Gauss-Seidel (see Algorithm 1) or nonlinear block Jacobi. If a residual can be supplied, then Jacobian-free Newton-Krylov methods can be employed (see Algorithm 3). If sensitivity information can be supplied, various preconditioning strategies become available. Viable coupling strategies can be severely restricted by the design of application components.
- *Memory management*: Low-level memory management and the handling of pointers are critical to robust multiphysics coupled codes. As data is transferred between applications, strong memory management tools such as reference counted smart pointers [439] will reduce memory issues and debugging efforts.
- *Error handling*: Safe error-handling policies are needed for a robust coupling capability. If components are aborting on certain threads/processes, other threads/processes of the coupled application may hang. Strong error-handling strategies and safe ways to shut down individual application components are needed. Exception-safe code is critical.

4.3. Difficulties in Collaborative Multiphysics Software

Building on the previous discussion of challenges in multiphysics software design, we now address practical difficulties in collaborative multiphysics software.

4.3.1. Maintenance of a stable and efficient development environment

Managing a multiphysics capability composed of individual application codes from multiple institutions places a number of additional hurdles on coupling infrastructure and application developers. A central issue is how to maintain a stable and efficient development environment in the face of distributed development teams, who may be working under intellectual property, export control, and cybersecurity constraints.

Since individual application codes are under active development at various institutions, procedures must be in place for synchronizing changes to the multiphysics infrastructure (sometimes called a framework). A common procedure is to wait for official releases of each application and then incorporate the changes into the multiphysics infrastructure. Depending on how frequently releases are made, updates based on application release cycles can allow

for large windows over which bugs can be introduced and compound. This circumstance makes updating applications difficult and time consuming, since developers must track down when and how the codes have changed.

A more recent software engineering approach is to use *continuous integration* (CI), where an automated system continuously checks for “commits” to individual application codes. When a new commit is found, the code is automatically downloaded and tested against the coupling infrastructure. This approach allows an automatic documentation as to which commit breaks a code. Automated emails can be sent to the developer owning the commit. Bugs can be caught and corrected the day they are introduced, leading to a much more efficient and stable development environment. Such a system requires a rigorous testing system as well as a disciplined development team willing to maintain 100% passing tests so that new failures are not masked by earlier failures.

An advantage of using an automated server for CI is in handling complexities associated with intellectual property and export control. For example, application developers may not be allowed to access the complete coupled code base, yet changes by any application developer may impact the coupled code base. An automated testing system is the first line of defense for notifying developers that their change broke a downstream dependency in code that they cannot access. Cybersecurity concerns can be addressed by locating the CI server outside firewalls.

A recent success story for the CI model is the CASL Nuclear Energy Hub [375]. The consortium is composed of 10 core institutions plus an additional 11 contributing partners, each developing their own application codes that are unified under a virtual reactor environment. Because of intellectual property constraints, application developers across the institutions are not allowed access to all application modules. Instead, a three-tiered approach is taken to protect the code base. The first line of defense is a precommit testing script that runs a fast set of unit tests on all codes to which a developer has access. Once this local testing is complete, the commits are pushed to the appropriate institution repositories. A CI server for the multiphysics infrastructure polls the institution application repositories every few minutes for changes and will launch a comprehensive test suite that includes all CASL partner codes whenever changes are detected. Results are submitted to a dashboard, and developers are notified by email if their commit introduced a failure. The execution time for tests run for precommit and CI testing is strictly limited so that users can obtain quick feedback that does not impede development efforts. The final tier of defense is a more comprehensive nightly test suite consisting of all precommit/CI tests plus any tests whose execution time is longer than appropriate for CI testing. This is run overnight and takes on the order of hours to complete. This testing can include features such as memory leak testing (see, e.g., valgrind). The procedures developed for CASL are an extension of the Trilinos build and test system.

4.3.2. Investment in software tools and techniques

A significant issue that is exacerbated by multiphysics coupling is the lack of investment in software engineering tools and strong software quality practices across application development teams. The use of software engineering tools (integrated development environments, version control, compilers, build/test and reporting tools) varies widely across institutions and research groups. Multiphysics additionally burdens the performance of the tools because of scaling from a single application to multiple, linked applications. For example, dependency tracking can cause unacceptable configure and build times if n^2 complexities exist in the configure system. Such issues might not be noticed in a single application where a smaller number of files and/or dependencies are present.

The lack of adoption of good software engineering tools impacts the development cycle of applications. One issue is that the tool might not be portable to all platforms of interest. Another issue is that a tool may not be funded or supported in the future, leaving users in a difficult position. Investments in open source software engineering tools should be a priority to support HPC applications. An example of a successful project along these lines is the CMake build system [440]. Tools like this are critical, not only for single applications, but especially for multiphysics coupled applications.

Poor software engineering practices can make integration of application components difficult. Some applications may have been developed to run only on a specific system or with a specific compiler with nonstandard conforming extensions. Porting to new platforms can be difficult. Memory management, error handling, and global scoping (see Section 4.2.6) can all be controlled with proper software coding guidelines and toolsets (see, e.g., [441]).

Multiphysics coupling invariably requires changes to each application code. The stronger the testing in place, the easier refactoring is. Fragile code bases with minimal testing often result in “code stagnation” as developers are wary of making large changes necessary for integration. The adoption of test-driven development and agile development strategies can address these issues.

4.3.3. *Adoption of third-party software*

Despite the many advantages of code reuse and success stories of the sort presented in Section 4.1.3, some application developers resist the adoption of third-party software. As introduced in Section 2.2.2, there are technical, programmatic, and personal reasons for this resistance. One is the risk, or perceived risk, of not being in control of the software that is being leveraged, or its long-term availability. Concerns include the long-term maintenance of the third-party software, the viability of the organization producing it, and the ability to influence the development of the software to enhance or maintain its relevance to the application being developed. Concerns may also exist about the “return on investment” of time in learning the software; it is not always clear that the time invested in learning the third-party tools will be recouped by increased development efficiency and by reuse in the longer term. Along these lines are concerns about the “opportunity cost,” or perceived cost, of using such tools. Deciding on a particular third-party capability may be hampered by concerns about making a poor selection, particularly if multiple third-party tools appear to meet the user’s requirements. The difficulty in comparing and assessing complex, high-consequence products can result in decision paralysis.

Another concern is the portability of third-party software. The software must span the hardware of interest to the user and must be both supported and tested on these platforms as part of the long-term maintenance strategy.

Another practical consideration is the quality of the third-party software. What processes will be used by the user and supplier to verify and validate the code base being leveraged? If the user makes changes to the third-party software, how do these changes get propagated back to the main repository? How are intellectual property and export control issues addressed?

Emotional considerations may also be a factor, such as the resistance to use software that is “not invented here.” Members of the user’s development team might feel security in developing their internal code and insecurity in using third-party capability. Indeed, a developer’s self-worth may become intertwined with the application over time.

4.3.4. *Licensing*

Additional impediments to using common multiphysics software are (1) the cost of commercial software, especially on machines with large numbers of nodes, and (2) unclear or unacceptable licensing terms for some open source software. Publicly available open source software is particularly attractive, and it is consistent with the long tradition within the HPC community. Not only does open source software remove a cost and inconvenience barrier, but it also allows users to determine exactly what the software does, promotes evolution according to user needs, encourages porting to many platforms, and ensures long-term availability of widely used software. It is important to use well-known licenses such as BSD, which confer broad rights to users. Licenses such as GPL, which restrict what users can do with code they have written, are unacceptable to many users, and unfamiliar licenses, regardless of terms, often require lengthy legal consultations at commercial companies and government laboratories.

5. Opportunities in Multiphysics Simulation Research

The preceding three sections have highlighted representative multiphysics applications, analyzed the structure of typical multiphysics systems algebraically, and explored implications for their implementation in software. Here we present a number of opportunities arising from these discussions that are ripe for research or development. Some of these opportunities call primarily for leveraging known concepts in new contexts. Some await breakthroughs. The prospect of predicting the behavior of complex systems combining multiple physical phenomena is one of the main motivations for extreme computing [13], so adapting the uniphysics components to architectural and scaling stresses is paramount. Since the number of such components may often be just two and rarely exceeds three to five in practice, issues of scaling flops, bytes, or transfers due to the multiphysics combination introduce only a modest factor beyond the orders of magnitude often called for in scaling up the components. However, difficulties still lurk in the combination. We recall the admonition of physicist A. S. Eddington [442]:

“We often think that when we have completed our study of one we know all about two, because ‘two’ is ‘one and one.’ We forget that we still have to make a study of ‘and.’ ”

After reviewing issues in forthcoming extreme computing environments in Section 5.1, we discuss in Section 5.2 concepts that are ripe for leveraging in multiphysics contexts, followed in Section 5.3 by broader multiphysics opportunities that require new breakthroughs in research and development.

5.1. Exascale Issues

For the most challenging multiphysics simulations, systems capable of sustaining an exaflop (10^{18} floating-point operations per second) or more will be necessary. But exascale systems are unlikely to look like current systems, even petascale systems. Very roughly, the power consumption of current peak petascale systems (meaning a system with a peak performance of roughly a petaflop) is already around 7 – 10 megawatts. At an average cost of \$1 million per megawatt-year, just providing power for current leadership-scale systems is a major part of the operating cost. For a sustained exascale system, computations must be about three orders of magnitude more efficient to keep the power bill alone under \$100 million per year. A more detailed analysis shows that data motion, whether it is between processors and local memory or between different processors, is a major contributor to the power cost. Consequently, successful programs on exascale machines may need to limit memory motion—current multiphysics approaches that move data representations between different physics components may not be feasible on exascale systems.

Power consumption is also limiting the clock speeds of individual processor cores; in fact, since 2006, processor clock speeds have been relatively stable and are not expected to increase much in the coming years. Consequently, higher performance will require much greater degrees of parallelism. An exascale system may have 10^9 cores, each of which runs at 1 – 2 GHz. Further, current DRAM (main) memory is a major source of both cost and power consumption. Without a change in technology, an exascale system is likely to have far less memory per core or per flop than is common today, where the rule of thumb is about 1 byte per flop. A likely ratio for an exascale system is more likely 0.01 – 0.1 bytes per flop. Adapting to this will require methods that are simultaneously more memory efficient (e.g., higher-order, compact schemes) and retain massive amounts of concurrency (easiest in a lower-order, noncompact schemes).

Furthermore, there is the issue of fault resilience. Some predictions about the likelihood of faults, based on the increasing number of processors cores, significantly overpredicted the failure rate. The reason is that faults are related more to the number of connections or parts than to the number of transistors. Thus, it has been possible to keep the failure rate of systems designed as massively parallel computers low enough that simple checkpointing strategies have been sufficient for applications. This may no longer be possible in an exascale system. While it may be possible to keep the number of parts roughly the same, other issues, such as approaches to reduce power consumption and increase memory size, may raise the probability of faults. Thus, applications may need to take a more active role in being resilient to faults (a general solution is possible but would consume extra power and may not be feasible).

5.2. Concepts Ripe for Leveraging

We describe some concepts that seem ripe for leveraging from existing applications and frameworks in the campaign to scale up uniphysics codes and to combine them.

5.2.1. Architecturally optimized, distributed, hierarchical data structures from solver packages

In coding for extreme architectures, there is much to learn from existing solver packages. Open source solvers such as hypre [188], PETSc [8], ScaLAPACK [443], and Trilinos [11] appear in inner loops in a vast variety of applications, have received attention from the global community for years, and are among the earliest pieces of software to explore and be ported to new architectures. They have benefited from autotuning in their use of the memory hierarchy and multicore structure and from sophisticated manipulation of tradeoffs in data transfers and in the use of local and global representations to satisfy the demands of working at different points on the performance/programmability spectrum. Examples of such issues include cache and register blocking, loop unrolling and jamming, coloring, message aggregation, and communication-computation overlap. Some use the data structure implementations in such libraries directly for customized ends; others call the solvers built upon them at a higher level. In high-end physics applications if each component controls a separate set of data structures, there may be little chance for strong performance at extreme scale, because of the overwhelming electrical power and execution time costs of transfers in the inner loops between different data structures. On the other hand, as discussed in Section 4.2, there is a benefit from separation of concerns: if the solver is intermingled with the physics, it is difficult for innovations in either domain to be introduced. A fairly commonly used interface between the physics and the solver is a call-back subroutine such as PETSc's `FormFunction`. Through this interface, PETSc provides state data, u , on a local data structure (submesh or sublist) that includes data owned by the local process and data gathered from interacting processes needed for the current step. The user locally evaluates the local piece of the residual or tendency function vector, $F(u)$, and returns it to the solver. The user code within `FormFunction` may be legacy or may be rewritten or generated automatically from an abstract problem description language for optimal cache and register use.

5.2.2. High-accuracy interpolation

Multiphysics codes often require interpolation between continua represented on nonconforming meshes at interfaces or throughout bulk regions, or between different formulations, such as particles and fields. High-accuracy multivariate interpolation is a well-studied problem, though high-performance implementations that make use of existing data structures without creating extensive intermediate working sets may be too much to ask for in library form. Nevertheless, multiphysics applications should not be limited by interpolation errors. In cases where the discrete data structures are solution adaptive or track a dynamically evolving interface, it may be necessary to preserve the underlying geometry of interface and query it directly to adapt the discretization.

5.2.3. Multiphysics as a form of adaptation

Multiple physical models often arise as a form of solution adaptivity, for instance, in atomistic-continuum models of crack propagation as introduced in Section 2.1.4, in which a continuum description that would possess a mathematical singularity at the crack tip is replaced with a more first-principles atomistic description. Such models may fit the classic design paradigm of local nested h -type mesh refinement, as an abstraction. Instead of refining the mesh in a subregion, one refines the model. Decades of experience with adaptive mesh refinement frameworks such as Boxlib [353] may be useful in this context. They possess built-in features such as timestep control, subcycling, and refluxing that may be leverageable in far more general contexts; see Section 3.2.2.

5.2.4. Applications of classical numerical analysis

Multiphysics coupling of components that are individually high-order in time are formally limited to first-order in time for standard splittings into two sequential half-steps, though formal second-order temporal accuracy may be achieved in a two-component problem by Strang splitting [2], which is equivalent to reading out the states at the half-step of one of the component steps, or at a quarter-step overall. First- or second-order splitting can reduce the order of accuracy of the multiphysics integration below that of the components; and unless the componentwise operators satisfy various commuting relationships, which would be rare apart from certain space dimensional splittings in linear constant coefficient problems, this would be the end of the road. One solution is to use Richardson extrapolation in time. Another approach is to use spectral deferred correction (see Section 3.2.2). A third option is to go fully implicit with all of the components.

Often, slow convergence in a multiphysics iteration is due to the degree of coupling. In a loosely coupled solution procedure, numerical effort to reduce the residual in one physics component may lead to an increase in the residual

of other components. Finding ways to improve convergence behavior, short of going fully implicit, can reduce this inefficiency.

If going fully implicit, the coupled system may possess linear stiffness and/or “nonlinear stiffness” that is much greater than any of the individual components. Ordinary, nonlinear globalization strategies, such as backtracking, are designed to prevent stepwise increase in the residual; however, this strategy fails when increases in norm of the nonlinear residual function need to rise in order to pass from one local minimum to another. Different globalization criteria, such as pseudo-transient continuation [444], may be needed for cases where simple globalizations fail.

5.2.5. *Exploitation of additional concurrency*

Hardware developed in the campaign for extreme computing has made flops relatively cheap and memory and memory bandwidth more scarce than ever, while vastly expanding concurrency. This may have enormous implications for multiphysics simulations. Time-parallel methods, e.g., parareal [445] are probably contraindicated as a source of additional concurrency, since they store many additional copies of the state vector at different time levels (essentially one new copy of global state per processor/thread engaged in the temporal direction). Additional concurrently available flops may encourage physics and solver modules that store many fewer intermediate common subexpressions and compute more on the fly.

Arithmetic intensity (roughly the ratio of flops to loads/stores required to execute an algorithm) is boosted for formulations that increase the interaction between elements of a given state at each timestep. This is typically the case for implicit methods, which may also be employed to reduce the number of timesteps relative to explicit methods in many systems, where stability, and not accuracy, sets the timestep limits. Arithmetic intensity may also be boosted for formulations that increase state (which has detrimental storage implications) but increase useful flops faster—superlinearly in the size of the state. This may be true of many of the advanced formulations of exascale simulations, in which the additional “state” is actually auxiliary to the physical state, such as sensitivities in uncertain quantification problems or adjoints in optimization problems. Stochastic error analysis is perhaps newly affordable.

5.2.6. *Exploitation of “tricks” for implicit algorithms*

Computations that are implicit often use a Jacobian-free Newton-Krylov algorithm (see Section 3.1.2), which avoids formulation and storage of a Jacobian matrix (which can otherwise contribute to the dominant-sized and least reusable object in the solver) and relies on evaluation of many residual functions, $F(u)$ for different states u presented by the Krylov method, as given by Equation (12). Normally, a second evaluation at a perturbed argument $F(u + i\sigma v)$ is necessary to estimate by finite differences $F'(u)v$, in order to generate the next Krylov vector. It is, however, possible for many functions F to evaluate $F(u)$ and $F'(u)v$ simultaneously from real and imaginary parts of an overloaded complex-valued residual evaluation. Implicit capability can also be employed to capture stationary periodic behavior, for which we expect $u(t + T) = u(t)$, in any transient model with a black-box propagator, $u(t + \tau) = P(u(t), \tau)$. One can apply Newton’s method directly to the fixed point $F(u) \equiv u - P(u, T) = 0$.

5.2.7. *Provision of special hardware*

The extra transistors on leading-edge chips may be dedicated to functions that are not used often but significantly reduce the latency of other operations they replace in hardware. One example of such a “luxury” may be random number hardware, which would be useful for multiphysics models that include stochastic physics components or stochastic forms of uncertainty quantification.

5.2.8. *Software maintenance best practices*

Multiphysics simulation codes can be tested and maintained with reference to the tests and maintenance practices of the uniphysics component codes; however, as discussed in Section 4, there may be subtleties. Manufactured solutions used to verify the components, which assume that $\frac{\partial F_1}{\partial u_2} = 0$ and $\frac{\partial F_2}{\partial u_1} = 0$, should be generalized to take into account the variation of the cross-components. A regression test suite for the multiphysics problem should hierarchically combine the uniphysics regression tests, again to take into account the new cross-component dependencies in the flow from inputs to outputs.

5.3. Required Breakthroughs

In contrast to the opportunities above, which should be straightforward developments of existing techniques in new contexts, we describe here some concepts that seem to require new breakthroughs in mathematical understanding, algorithmic development, or software design and implementation.

5.3.1. Formulation of objective functions useful for optimizing simulations

Excluding the human cost of programming and operating supercomputers, general figures of merit in assessing the impact and efficiency of simulations may often take the form of “scientific impact per unit energy” or “scientific impact per unit time.” Scientific impact is obviously hard to measure in absolute terms but may be represented by a proxy such as accuracy. If the goal of the extreme-scale simulation is beyond the domain of original insight and is related to the quality of prediction, then “accuracy per joule” is a reasonable desideratum. Like the Drake equation of astronomy, this estimate can be factored into a number of ratios coming from more basic assumptions, in this case from numerical analysis and computer architecture, for instance: accuracy per stencil evaluation, stencil evaluations per flop, and flops per joule, which can be subfactored, in turn, given problem- and machine-specific details. Each of these factors depends on myriad assumptions about smoothness, convergence rate, resources of the memory system, and so forth. As difficult and subjective as it may be to define a good objective function, we cannot hope to optimize it without quantifying it.

5.3.2. Characterization of uncertainty

Each component in a multiphysics application needs to have an uncertainty characterization of its own in quantities of ultimate interest. These may not be the primitive variables of the simulation but functionals of components of the state vector, such as the integral over a surface of a normal derivative of a field representing its flux. Uncertainties that depend on quantities coming from other components, in the form of boundary conditions, interior forcing terms, constituent laws, and so forth, must also be characterized. This characterization may be particularly interesting for quantities of different types such as particles and fields. The best ways of quantifying uncertainty for simulations of multiphysics processes whose models have nondeterministic aspects (e.g., bubble formation in foams, surface tension effects, cracks) remain to be seen. In general (see, e.g., Section 3.5), one cannot ignore cross-component effects in quantifying uncertainty, since uncertainties propagate across couplings, and there may be cyclical effects. Consequently, multiphysics applications will need to have the ability to sample from the output of individual physics components, and uncertainty quantification will need to become more intrusive.

5.3.3. Equidistribution and control of error

Mature uniphysics component codes come with error control mechanisms, such as variable-order, variable-size timestepping, h -, p -, and r -type spatial refinement, particle generation and removal, and model substitutions, which respond to error estimators, typically of a *a posteriori* type (see, e.g., Section 3.4). However, multiphysics applications must evaluate the cost and benefit of refinement strategies in a much broader context. It is possible to be deceived into spending nearly all of the execution resources in a component whose accuracy is already much better than the limiting accuracy in the system, which could be due to a different component, or at some interface between components. Context-appropriate premiums must be placed on gross checks like integral conservation, and consistency of representation of fluxes at boundaries. Of course, error requirements may sometimes be exceeded when parameters are chosen with respect to some other criterion, like stability, which may be more stringent for a coupled system than for any of its components. These thresholds of limiting parameters should be identified because they could trigger dynamic adaptations of the algorithms.

5.3.4. Dynamic characterization of cross-physics interaction strength

Though cross-coupling enters into performance efficiency considerations in subtle ways like arithmetic intensity, which may be better for a coupled system than for its components, making the coupling effectively “free,” efficient multiphysics algorithms will preferably allocate the extra work of enforcing tight coupling only where the interaction is strong and will default to loose coupling where the interaction is weak, or sequential where it is unidirectional. Algorithmic selection, which may be dynamic over the duration of the computation, therefore requires dimensionless or relative measures of cross-variable sensitivities. Just as error estimators are required to inform adaptive discretizations, some measure of interaction strengths between the components is required to inform appropriate coupling tightness.

5.3.5. Reduction of data transfers

Multiphysics simulation, as an extreme-scale simulation, will generally execute on hardware on which the energy costs of data transfers, between processor-memory units or within the memory hierarchy of a single unit, is at a premium. Many new programming paradigms are under consideration for such hardware, which represents the end of a successful era of correctness-portable, relatively performance-portable, MPI-based SIMD code design. There is discussion of paradigms such as “moving data to code, rather than moving code to data,” even though this cannot be perfectly realized for applications—like nearly all physical world models—with all-to-all or even one-to-all data dependencies. Whatever successful paradigms are developed for uniphysics problems may be vulnerable to high energy cost of transfers if the uniphysics codes are combined in a black-box manner, since this normally requires the repackaging and redistribution of data. We expect that significant efforts in codesign will be required in order to keep data transfer costs manageable in the multiphysics applications for which extreme computing environments are built.

5.3.6. Relaxation of synchrony

Tomorrow’s extreme-scale computing systems will have less performance reliability than do their contemporary counterparts for a number of reasons well documented in the IESP exascale roadmap [446]. Energy-efficient, low-voltage differentials between logical “0” and “1”, for instance, may mean that some arithmetic operations may be corrupted by noise and need to be redone. Because of the decreasing performance reliability, it will not be possible to load balance with sufficiently high precision for today’s SIMD bulk-synchronous algorithms to run efficiently. Less-synchronous algorithms with work stealing will need to be developed within components and across components. Today we synchronize artifactually with too heavy a reliance on nested loop structures, and far more frequently than is required for physically causally correct computations. In many places, it should not be difficult to beat back to causal limits. However, asynchronous algorithms typically compromise robustness relative to strongly synchronous counterparts. Both programming paradigms and mathematics will need to be revisited. For instance, global `MPI_COMM_WORLD` will be replaced with local communicators with the smallest required scope to allow greater independence of threads, allowing different physics components to complement each other in cycle and resource scavenging without interference. Algorithmically, data that is causally required but unavailable may often be predicted and the computations dependent on such data rolled back if, when the data becomes available, it is determined that the output would be too sensitive a function of the discrepancy between the actual and the predicted data. In practice, low-frequency components of the predicted data may be accurate, and high-frequency discrepancies may be demonstrably irrelevant because they are rapidly decaying. There are numerous such opportunities for relaxation of artifactual synchronization or too strict an observance of bona fide synchronization, both within and between components.

5.3.7. User-specified checkpointing

At a level of unreliability beyond individual operations, executing long-running multiphysics codes on extreme-scale computing systems with many points of potential hardware and software failure will need user-specified checkpointing. Operating system checkpointing must be conservative and save complete program state, without appreciation for the fact that much of the bulk of the computational state can be reproduced from a much smaller physical state. For example, a Jacobian matrix does not need to be saved for restart since it can be computed from a (typically much smaller) current state vector. Data related to multiphysics coupling may be reproducible from the union of the uniphysics states. There is opportunity to specify minimum-size checkpoints for reproducible restarts as a union of checkpoints from individual physics packages so that the overall simulation is robust. This may also be good software discipline for debugging complex multiphysics codes on contemporary machines.

5.3.8. Combining continuum-based and discrete phenomena

Many multiphysics problems span scales from atomistic (or other discrete) phenomena to the continuum. The coupling of discrete and continuous formulations transcends all scales of computation and remains a theoretical challenge; see Section 3.3. To perform judiciously truncated computations in either the discrete or continuum-based world requires quantification of their associated errors and the errors of coupling so that errors can be “equidistributed” and no one phase, including the coupling, is driven too hard. There are several opportunities to be investigated here, including how best to transfer state, fluxes, and forces between the formulations, how to estimate and control errors

incurred in the transfers, and how to develop a unified formulation, perhaps involving penalties or discrete conservation principles.

5.3.9. *Multiphysics software design*

The time is ripe for the multiphysics community to extend the successful software approaches introduced in Section 4, as well as to leverage the features of emerging computer architectures, so that new algorithms and approaches can be explored. Innovative next-generation multiphysics software has the power to transform computational science, with a goal of enabling fully predictive simulations.

5.3.10. *Multiphysics performance models*

The performance of many computational science applications has been dominated for many years by the cost of moving data rather than raw floating-point rates. For a large multiphysics applications, data typically needs to be moved between components; it may also need to be reorganized if the data structure details are different. These factors raise numerous questions:

- What is the cost of moving data between components? If the cost is high, are there approaches that can mitigate the overhead, such as multistep communications or other communication minimizing or avoiding techniques?
- What is the overhead of translating between representations? Is it more efficient to translate or to select a common representation (which may not be optimal for any individual component but is optimal for the overall simulation)?
- In a parallel program, what is the impact on process placement? For example, should the different components be laid out in compact sections on the parallel system or intermingled to reduce communication costs? Is dynamic load balancing practical?
- What is the overhead of using library routines (particularly a sequence of routines) over using custom code?

Relatively simple performance models can be developed to address these questions. For example, models that determine the number of compulsory cache misses can be used to determine an upper bound on performance and are still sensitive to the parameters of the system's hardware. Similarly, these performance models can be used to help design and optimize the use of hybrid programming models, such as MPI with OpenMP.

6. Insertion Paths for Algorithms and Software into Multiphysics Applications

As many of the challenges presented in the previous section are tackled, new algorithms and software designs will be available for use in existing multiphysics applications. The next set of challenges will lie in how to incorporate these new technologies into simulations in ways that will minimize the implementation time and maximize their effectiveness. The issues that arise involve both technical and nontechnical areas, and we examine considerations from both of these viewpoints. We also present ideas for successfully incorporating new technologies into multiphysics application areas.

6.1. Mechanisms

Before looking at factors that affect insertion of new technologies, we outline three common mechanisms for such insertions.

6.1.1. Leveraging of interfaces

One successful insertion path is that of engineering current software to leverage interfaces. Software infrastructure that encapsulates functionality in terms of well-defined interfaces is able to easily adopt new technologies by swapping capabilities behind those interfaces; see Section 4.2. This design allows a new technology to “look” like the old from the point of view of the application. In this framework, direct comparisons between the old and new technologies can easily be made. In addition, well-defined interfaces often allow the adoption of external libraries, in turn providing the ability to investigate the performance of many methods before choosing a final one for use. Encapsulation of functionality behind good interfaces provides an effective way to transition between technologies that adhere to those interfaces. As a result, more options are available to a given application.

We note that well-defined interfaces can sometimes prohibit extreme performance tuning. For mature technologies, where one technology shows a clear benefit, the code framework may be more closely tied to the technology implementation in order to benefit more from efficiency gains. However, the tradeoff often means that interfacing with any further new technologies is much more difficult.

6.1.2. Small, side effort

In many cases, a transition to a new technology starts with a small, side effort. Often this activity is a splinter from a main, funded effort and is conducted by a student or junior staff member who can devote significant time to the project. Within this effort new technologies can be tried in a specific branch or simplified code. If this proof-of-principle effort proves successful, the technology can then be moved into the main code. The principal risk with this mechanism is that since junior staff (typically, temporary or shorter term) are often the ones conducting most of the work, the people who are most knowledgeable about the new technology may not be present after the initial testing and use.

6.1.3. Full rewrites

An effective way to port new technologies into multiphysics applications is to fully redesign the software considering the demands of the existing multiphysics setting; unfortunately, this is also the most expensive in terms of time required to begin conducting simulations. Additionally, there is no guarantee that for the present project a full rewrite will produce significantly better results than will reusing the existing code. Also, with a full rewrite, it will likely take longer to begin initial simulations.

Even so, fully rewriting software is the best approach if (1) the existing software is out of date by today’s software engineering standards, (2) the mechanism for coupling this software to the multiphysics application is fragile, and (3) someone is willing to pay for a full rewrite. By totally redesigning the technology with a multiphysics mindset, coupling and scalability issues related to data traffic can be addressed effectively. Also, the software can be designed extensibly to take advantage of the most modern algorithms and libraries, as well as to prepare for future evolutions; this approach will help delay any future rewrites and encourage incorporation of new technologies across the community by demonstrating a successful implementation. The value of this approach is in the long-term viability of the multiphysics application achieved by designing it to adapt to existing and coming software practices while still leveraging the knowledge present in the original code. Having the original designer of the technology available is helpful in understanding the underlying mentality of the original design and ensuring that the most effective aspects

still appear in the rewrite. We further note that one cannot assume that the new code is as correct as the original; therefore, a full verification is often required when a rewrite takes place.

6.2. *Factors Affecting Successful Insertion Paths*

Successful integration of new technologies into existing multiphysics (and often multidisciplinary) applications depends on adequate funding, a solid collaboration among participants that is well supported by the home institutions, new technologies meeting well-motivated needs, and software engineering that allows for flexibility. All these conditions rarely are present simultaneously. Below we list some key considerations that have given rise to successful integrations of new technology into simulation areas. Our intention is that when all the above conditions are not present, participants can look at these areas as ways to make good progress.

6.2.1. *Clear motivation*

The most successful integrations of new technologies into existing simulation codes occur when the new technology or enhancement is clearly needed. However, understanding the goals and hence needs of applications research codes can sometimes be difficult. In addition, assessing the current and potential performance of the code and any new technologies is not straightforward.

A typical research code is often used to analyze or simulate multiple and very different physical phenomena over its lifetime. As a result, the goals of such codes are fluid, and identifying specific requirements can be challenging.

When specific goals and needs are identified, however, a new factor comes into play: the determination of whether a new technology can help the simulation code meet a goal. In order to address this question, performance analyses are required. Multiphysics applications are often large and sometimes difficult to manage, making performance analysis on multiphysics software complex. This complexity, in turn, can make it challenging to quantify the computational goals of multiphysics applications (e.g., memory demands, data traffic, computational speed, order of accuracy). One approach to conducting this analysis is to produce “skeleton codes” that allow for simpler identification and classification of contributing costs. These skeleton codes have the same semantic structure as the complete multiphysics codes but without the full complexity, both computationally and physically.

In order to use this approach, simplified models of the relevant component physics are developed and then included in the skeleton simulation. These skeleton simulations are then profiled, with the goal of identifying the issues that will be present in the full multiphysics simulation and produce clear targets for further progress. These activities invariably help focus the needs and expose the motivations for use of potential new technologies.

6.2.2. *Collaboration*

One of the most significant impediments to multiphysics applications also produces the most potential for great progress: collaboration between interdisciplinary groups. While the initial implementation of the technology may have been designed only by an application developer, moving to the multiphysics setting—and hence more complex simulations—will likely require the input of software developers to optimize efficiency and mathematicians to study stability and accuracy.

The value of communication goes both ways, though. Math and computer science experts must understand the problems that are involved in the simulation and the assumptions that went into the underlying models. Without this input there is likely no way that full analysis or optimization can be done, and time will be lost while necessary physical assumptions or conditions may be violated during the development. All members of the interdisciplinary development team must understand both their role and the role of the other participants. The most successful collaborations will produce multiphysics simulations that allow for new physical insight obtained with computational efficiency and mathematical accuracy.

Multidisciplinary collaborations also tend to be most successful when each member takes the time to understand the reward structure of other members’ institutions and fields. Authorship order, choice of journals, choice of conferences to present the work, and software releases can all be important factors, but of different weight in different disciplines. It behooves team members to understand these and, to the extent possible, help the members meet their institutional requirements.

6.2.3. *Discretization and meshes*

The development of new multiphysics capabilities for which solution parameters must be related between physics components requires understanding and interacting with the meshes, as well as distribution of the field parameters over the mesh, as indicated by the equation discretization process. The complexity of developing the needed coupling strongly depends on the degree to which the types of meshes, physics discretizations, and data structures used are, or are not, matched. The insertion options associated with various approaches are outlined in Section 4.2.3.

6.2.4. *Software packages*

The approaches for addressing multiphysics software challenges given in Section 4 apply to both application-specific projects and to software packages. The burden on package developers is greater, however, because the success of a package depends on its integration into other applications.

Good software engineering alone does not address the reasons given in Section 4.3.3 for resistance to adopting external packages. The developers of software packages need to minimize the real and perceived costs of using their work. Documentation has always been a necessity; and as projects become more complex, that documentation needs to be not just thorough but well structured. Equally important are examples of the intended use of the package. Examples, particularly multiphysics examples, can demonstrate relevance, give application developers an idea of how to match their codes to the new technology, and drive home the advantages that new technology has over what it is intended to replace.

The chances of an external package being successfully integrated into an application increase if its developers are able to provide support during its use. The developers of multiphysics software often have experience with more varied high-performance computing resources than do their users; hence, their help with migrating to new platforms, troubleshooting, and profiling can be extremely valuable. Support should also include a willingness to adapt software packages to the needs of applications. Just as applications should adopt software packages with new technologies to stay competitive, software packages should incorporate the feature requests of applications to stay relevant.

6.2.5. *Funding*

Funding is a key element of successful integration of new technologies. Often, funding in scientific applications areas is provided for use of the simulation to discover new science or conduct an analysis of a given phenomenon. Funding agencies tend to be reluctant to fund new code development without new science. On the algorithms side, funding is usually available only for development of novel methods or software libraries containing algorithm implementations done in a general way so as to benefit multiple application areas. Funding for integration of a specific method with a specific application area is often difficult to obtain. The recent NSF program on Software Infrastructure for Sustained Innovation and the DOE SciDAC program pose opportunities for such specific insertions of new technologies.

Clear motivation for new technologies often leads to greater chances of success for their integration. It also leads to greater success in funding that integration, because justifications are more readily made.

One reason science applications and their funding agencies give less priority to adopting new mathematical and computer science technologies is that new technologies can be less robust and hence less reliable. While a new technology may present an opportunity for greater speed or accuracy, the adopting code will need to go through a new verification study. In addition, input files and models may need to be adjusted in order to produce accurate simulations with the new technologies. Even when funding is available for the initial insertion, it is often not available farther down the road for new input files and verifications. Regularly run, detailed verification suites (not solely unit tests) with documented results significantly help in this endeavor. With these suites, new technologies can be verified more quickly, since the full machinery for doing verification with the code is already developed and functioning.

6.2.6. *Personnel*

Another important consideration for insertion of new technologies is personnel. Often, integration of new technologies requires some software engineering expertise. However, a small to moderate-sized application effort may not have a dedicated software engineer. Hiring such expertise can be difficult for a short-term position. As a result, scientists often perform large software engineering tasks themselves. However, reward systems at many institutions do not take such work into account. In addition, the work is often not in the scientists' main line of interest. As a

result, software engineering performed in such an environment may not take advantage of some best practices, causing possible delays downstream. Furthermore, a scientist may be less able to tune an algorithm or, when an algorithm fails, determine whether there is a bug in the implementation or an algorithm failure for a problem.

On the mathematics and computer science side, a similar trend is often seen. Again, because institutional reward systems value publishing in disciplinary fields, staff are most interested in implementing and developing new methods, not in placing known methods into a scientific code. This interest is in direct conflict with the scientists' need for proven and reliable methods to move the science simulation forward. In addition, while some funding exists to package software for use across applications, often package writers are not funded for longer-term support.

A further hindrance to effective cross-disciplinary work from the software developer side is the issue of recognition. Often, software development is seen as a necessary requirement to move a simulation's capability forward. However, since the main new science is done in modeling and interpreting code results, the software work is not given credit in publications nor in most criteria used for promotions at most institutions. As a result, recruiting for these tasks is difficult, and only rarely have packages and software been supported for long terms.

For many of these reasons, implementations often become the task of junior researchers, since they often are the only ones with large enough blocks of time to do complex implementations. These junior staff tend to be less familiar with effective software practices. In addition, junior staff have a high turnover rate and thus cannot be relied on to provide long-term knowledge and support of the software implementations.

6.3. Observations

Based on participant experiences and discussions during the workshop, we present the following observations on ways that are most likely to ensure success in integrating new technologies into existing multiphysics simulation areas.

1. Ongoing verification and software support activities require adequate scoping, budgeting, and funding from funding agencies and principal investigators.
2. Software development activities aimed at increased simulation speed and accuracy (and not just new capabilities) are important aspects of software development and merit consideration in proposal and publication reviews.
3. New technologies often bring in new models and new numerics. The most successful projects engage experts in validation and verification as well as in uncertainty quantification.
4. Successful long-term projects engage a software architect who ensures use of best practices and enforces standards.
5. Good meshing technologies enable complex simulations and allow the necessary discretization schemes to facilitate flexibility in algorithms.
6. Compartmentalization of code functionalities leads to extensibility for future technologies. Mesh, visualization, interpolation, solvers, and uncertainty quantification are areas where high-quality external code can be incorporated to facilitate software coupling.

7. Conclusions

The preceding sections demonstrate a vast variety of approaches to high-performance multiphysics coupling of PDE-based models contributed from the applications, mathematics, and computer science communities in developing the science, the methods, and the computational infrastructure necessary for bringing together complex simulations of multiple physics processes in parallel computing environments. As introduced in Section 1 and further detailed for various applications in Section 2, numerical coupling strategies range from loosely coupled Gauss-Seidel and operator splitting approaches to tightly coupled Newton-based techniques. Researchers have made substantial progress in understanding coupling issues for multiphysics components in space and time, including aspects of problem formulation, discretization, meshing, multidomain interfaces, interpolation, partitioned timestepping, operator-specific preconditioning, as well as difficulties that can arise in multiscale, multiphysics, multidiscretization operator decomposition (see Section 3). Building on the experience of teams that have developed application-specific multiphysics software and more general infrastructure in support of parallel PDE-based multiphysics problems, Section 4 identifies software features that facilitate robustness, extensibility, and easy of use – essential to support research on multiphysics algorithmic issues, architecture-specific enhancements, and changes over time.

While the computational science community has made significant progress in advancing scientific knowledge and engineering practice through these methods, we have outlined in Sections 5 and 6 numerous challenges that await the focused attention and critical thought necessary for further advancements. Many of these issues have received attention in venues of note. In particular, the 2008 report by D. Brown et al. [1], integrates input from a broad spectrum of scientific and engineering applications to highlight key challenges from an applied mathematics perspective in the predictive modeling and simulation of complex systems. Included are the following two research objectives:

- Advance the fidelity, predictability, and sophistication of modeling and simulation methodologies by developing the mathematical tools needed for the analysis and simulation of complex systems characterized by combinations of multiple length and timescales, multiple processes or components.
- Develop analytical and computational approaches needed to understand and model the behavior of complex, multiphysics and multiscale phenomena.

These two points underscore the significant and still remaining key need for conducting the mathematical analysis necessary to ensure that coupling schemes are accurate, stable, robust, consistent, and are implemented correctly – as a prerequisite for developing predictive multiphysics simulations. Moreover, emerging extreme-scale computational environments require fundamentally rethinking approaches to multiphysics modeling and algorithms, with attention to issues of data motion, data structure conversions, and overall software design.

Clearly, much work remains before we can fully realize the potential of multiphysics simulations on emerging high-performance architectures for high-impact policy and decision support. We expect that over the next decade, such investigations will unify practitioners of multiphysics simulations and advance many areas from the common current practice of “two scientists meet at a conference and decide to combine their codes” to fully integrated, numerically accurate, efficient simulations. It is our hope that by documenting the current state from the complementary perspectives of multiple applications areas, applied mathematics, and computer science, we will enable significant further progress through cross-fertilization of ideas that will lead to a next generation of integrated and accurate multiphysics simulations that enable scientific discovery.

ACKNOWLEDGMENTS

Acknowledgments

We gratefully acknowledge the Institute of Computing in Science (ICiS) [447] for support of this workshop. We thank all ICiS staff, especially Jutta Strate-Meyer and Cheryl Zidel. We also thank Satish Balay for establishing a Mercurial repository to facilitate collaborative writing of this document, as well as Gail Pieper for her masterful copy editing. We are grateful to Phil Colella and Hans Johansen for providing information about Chombo, and to Tim Tautges for providing information about MOAB.

Appendix A. Glossary

A continuing issue on the frontier dividing science and applied mathematics is the cultural barrier of language. The growth of multiphysics exacerbates this problem since it engages scientists and mathematicians from increasingly diverse areas. A better common language is clearly called for if science–math collaborations are to thrive. This glossary defines terms as used in this report.

- **strong (versus weak) coupling of physical models:** We use strong (versus weak) coupling to refer to strong (versus weak) interactions between different physics models that are intrinsic between the different physics in a natural process. A strong coupling could be due to large overlap of geometric domains or due to a strong influence of one model on another. For example, in climate modeling, the coupling between the ocean and sea-ice models is strong, while the coupling between ocean and land models is weak. Mathematically, the off-diagonal block of the Jacobian matrix of a strongly coupled multiphysics model may be nearly full or may be sparse but contain relatively larger entries. In contrast, a weakly coupled multiphysics model may contain relatively few or relatively small off-diagonal entries. See Section 2 (page 8) for further discussion.
- **tight (versus loose) coupling of numerical models:** We use tight (versus loose) coupling to refer to the algorithmic aspect of multiphysics coupling schemes (or numerical models) in terms of whether the state variables across different models are well synchronized. A tightly coupled scheme (sometimes referred to as a strongly coupled scheme by some authors) would keep all the state variables as synchronized as possible across different models at all times, whereas a loosely coupled scheme may allow the state variables to be shifted by one timestep or be staggered by a fraction of time steps (see Figure 1). For example, a monolithic fluid-structure approach (e.g., [448]) and a full iterative scheme (e.g., [449]) are tightly coupled, whereas a sequential staggered scheme (e.g., [39]) is loosely coupled. Note that there is not a direct one-to-one correspondence between strong (versus weak) physical coupling and tight (versus loose) numerical coupling, since a tight or a loose numerical coupling scheme may be used for strongly or weakly coupled physical models. See Section 2 (page 8) for further discussion.
- **operator splitting:** Multiphysics operator splitting is a loose coupling scheme that evolves one timestep for each physics component in sequence, as shown by Algorithm 2.
- **subcycling:** Given two time-dependent variables w_1 and w_2 , assume their time evolution is coupled between time levels t^{n-1} and t^n . Let $u_i \approx w_i$ for $i = 1, 2$ be decoupled approximations satisfying

$$\begin{aligned} \partial_t u_1 + f_1(u_1, u_2(t_{n-1})) &= 0 \\ \text{and } \partial_t u_2 + f_2(u_1(t_{n-1}), u_2) &= 0. \end{aligned}$$

(Note that Algorithm 2 may be written in this form by redefining f_2 .) It is said that u_2 is subcycled M times relative to u_1 if $u_2(t_n)$ is approximated by any timestepping method requiring M timesteps to evolve from t_{n-1} to t_n .

- **explicit methods:** In an explicit method, future states are given by a formula involving known computed quantities that can be evaluated directly. The simplest example is forward Euler: $y_{n+1} = y_n + \Delta t F(y_n)$. Explicit methods are inexpensive per iteration because they require a fixed number of function evaluations to complete the step; however, they typically have severe stability restrictions and, therefore, are suited only for nonstiff problems. See Section 3.2.2.
- **implicit methods:** A fully implicit method advances the solution in time by using current information and an inversion process with no explicit steps. The simplest example is backward Euler: $y_{n+1} = y_n + \Delta t F(y_{n+1})$. Fully implicit methods are relatively expensive because of the need to solve (non)linear systems at each step; however, they have very favorable stability properties [246, 247]. See Section 3.2.2.
- **semi-implicit methods:** Inclusion of both implicit and explicit elements gives rise to semi-implicit methods [248, 249]. See Section 3.2.2.

- **Frameworks and frameworks:** In the context of software, a framework provides a generic functionality (such as solving a system of linear equations) that can be customized by the user to provide an application-specific solution (such as a specific matrix and preconditioner). Frameworks typically handle the flow of control, which distinguishes them from the use of library routines. Some frameworks take complete control of the application; we called these Frameworks with a capital F. Other frameworks assert control only over the task they perform and can be combined in an application with other code or even other frameworks; we called these frameworks with a lower case f. See also http://en.wikipedia.org/wiki/Software_framework. See Section 4.

Appendix B. Multiphysics Exemplars and Benchmarks

A collection of multiphysics exemplars and benchmarks that encompass features present in many applications would provide a resource for motivating and validating new research on multiphysics algorithms and software. Goals of an overall collection include diversity of mathematical formulation (e.g., fields and particles), simplicity of code size to describe, range of scale and parametric challenge, and range of applicability to different user communities. Initial pointers to multiphysics test problems collected during the 2011 ICiS Multiphysics Workshop are available via the workshop website [450].

This appendix takes a first step at presenting a few multiphysics exemplars in a unified format, based on that used by Kamm et al. [451], though with a few modifications to convey additional information that we consider relevant in multiphysics contexts. A longer-term goal is to engage the community to contribute a broader range of representative multiphysics test cases and codes that could be published on a webpage that outlives an archival fixed publication and could be updated from time to time. We initially consider the following examples:

- **Advection-reaction** (Appendix B.1): This example is a prototypical reactive transport problem, and it is a good test for the convergence of IMEX time integration schemes and order reduction phenomena due to stiffness.
- **Heat transfer coupled across an interface** (Appendix B.2): A simple linear heat equation with constant diffusivity is posed on two neighboring, non-overlapping subdomains that meet along an interface.
- **Conjugate heat transfer** (Appendix B.3): This example illustrates the flow of a fluid around a cylinder of a different temperature, where the regions are coupled through the heat equations by continuity of state and normal flux. The flow of the fluid guarantees that the effects of error are felt downstream from the object. Coupling can be handled in various ways. Note that in general there will be different methods for the fluid equations and the heat equations, providing another difficulty.

Appendix B.1. Advection-reaction

- Contributor: Emil Constantinescu

- Conceptual Description

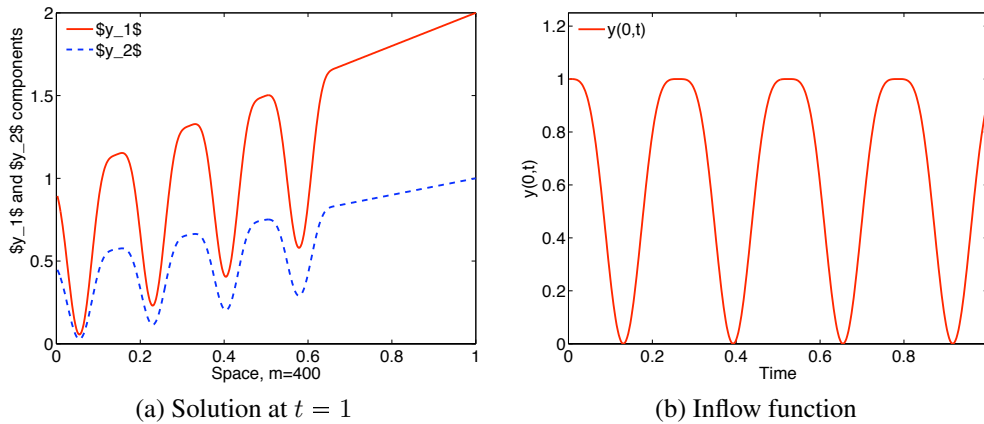
- General: This is a typical reactive-transport problem exposing two physical processes that share the same spatial domain: advection and a linear kinetic process. The two processes have different temporal scales—the advection process is relatively slow, whereas the kinetic term is numerically stiff. Therefore, this problem is suitable for partitioned time integration methods and, in particular, IMEX schemes as discussed in Section 3.2.2. This problem can be made very stiff, in which case the order reduction phenomenon and stiffness leakage can be analyzed. The governing equations model the dynamics of two chemical species by

$$\begin{aligned} \frac{\partial y_1}{\partial t} + \alpha_1 \frac{\partial y_1}{\partial x} &= -k_1 y_1 + k_2 y_2 + s_1 \\ \frac{\partial y_2}{\partial t} + \alpha_2 \frac{\partial y_2}{\partial x} &= k_1 y_1 - k_2 y_2 + s_2 . \end{aligned}$$

- Processes modeled: advective transport and chemical kinetics on the same spatial domain. We can write the problem as $y' = f(y) + g(y)$, where f represents advection and g (possibly stiff) reaction. The solution at $t = 1$ is illustrated in Figure B.23.a.
- Specification of initial conditions:

$$y_1(x, 0) = 1 + s_2 x, \quad y_2(x, 0) = \frac{k_1}{k_2} y_1(x, 0) + \frac{1}{k_2} s_2$$

- Specification of boundary conditions: $y_1(0, t) = 1 - \sin(12t)^4$ (see Fig. B.23.b). The right boundary condition is an outflow boundary. If $\alpha_2 = 0$, then there is no need to specify boundary conditions for the second component.
- Benchmark type: convergence of time integration strategies and order reduction.
- Description of quantities in which accuracy is important: For convergence analysis, all quantities are important in this problem. However, one may consider different quantities of interest, such as the total mass of one of the species.


 Figure B.23: Illustration of (a) the advection-reaction solution and (b) the boundary inflow term ($y_1(0, t)$).

- Principal Features Tested

- Discretization: This test is aimed at temporal discretization. Thus, we provide a simple, high-order (order 4) spatial discretization that can be used with this example. The spatial grid is defined as $x_i = i\Delta x$, $i = 1 \dots m$, $\Delta x = 1/m$; and the advective term can be discretized as

$$\begin{aligned} f(x_1) &= \alpha \frac{1}{\Delta x} \left(\frac{1}{3}y(0, t) + \frac{1}{2}y(x_1) - y(x_2) + \frac{1}{6}y(x_3) \right) \\ f(x_2) &= \alpha \frac{1}{\Delta x} \left(-\frac{1}{12}y(0, t) + \frac{2}{3}y(x_1) - \frac{2}{3}y(x_3) + \frac{1}{12}y(x_4) \right) \\ f(x_j) &= \alpha \frac{1}{\Delta x} \left(-\frac{1}{12}y(x_{j-2}) + \frac{2}{3}y(x_{j-1}) - \frac{2}{3}y(x_{j+1}) + \frac{1}{12}y(x_{j+2}) \right), \quad j = 3 \dots m-2 \\ f(x_{m-1}) &= \alpha \frac{1}{\Delta x} \left(-\frac{1}{6}y(x_{m-3}) + y(x_{m-2}) - \frac{1}{2}y(x_{m-1}) - \frac{1}{3}y(x_m) \right) \\ f(x_m) &= \alpha \frac{1}{\Delta x} (-y(x_m) + y(x_{m-1})). \end{aligned}$$

- Algorithms: This problem can be integrated in time with a fully implicit method as well as a partitioned method in which f is integrated with an explicit integrator and g with an implicit one.
- Solver: Multiple solvers can be used to test different aspects related to algorithmic efficiency and effectiveness. However, because this problem is a small-scale experiment, convergence and relative implementation and computational cost between fully implicit and IMEX methods can be explored.
- Code: The code can be easily implemented in Matlab. It is implemented in PETSc (which allows a simple switch between fully implicit and IMEX using additive Runge-Kutta schemes).

- Mathematical Description

- Problem domain: This is a one-dimensional experiment with $m = 400$ grid points for each component. The problem domain is delimited by

$$0 < x < 1, \quad 0 < t \leq t_{\max}.$$

- Ranges of data and parameters: Typical ranges for the parameters that yield a stiff problem are

$$\begin{aligned} \alpha_1 &= 1, \quad k_1 = 10^6, \quad s_1 = 0, \\ \alpha_2 &= 0, \quad k_2 = 2k_1, \quad s_2 = 1. \end{aligned}$$

- Accuracy Assessment: A reference solution can be computed by using a small timestep with classical stiff integrators.
- Testing Procedures: The stability and accuracy testing procedure (i.e., the convergence and its rate) can be achieved by measuring the error when using different time steps. For multistage methods (Runge-Kutta) the convergence rate beyond a certain error level may drop to the method's stage order—the order reduction phenomenon. Fully implicit methods suffer from the same peril.
- Additional User Information: When using nonlinear discretization schemes for f (shock-capturing schemes (e.g., flux-limited WENO), it is convenient to employ an IMEX procedure that does not need a Jacobian approximation of this component. Moreover, by using a splitting strategy one can take advantage of the fact that the stiff (implicitly integrated) component is linear and does not require nonlinear solves.
- References: [253, 254]; contact Emil Constantinescu (emconsta@mcs.anl.gov).

Appendix B.2. Heat transfer coupled across an interface

- Contributor: Brendan Sheehan

- Conceptual Description

- General: A simple linear heat equation with constant diffusivity is posed on two neighboring, non-overlapping subdomains that meet along an interface. Equality of state and normal flux are enforced along the interface. The subscripts “c” and “e” stand for core and edge and are used below to differentiate the subdomains.

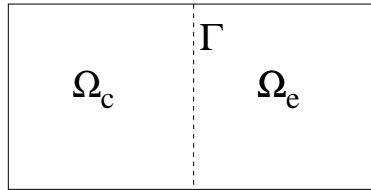


Figure B.24: Illustration of neighboring subdomains

$$\begin{aligned}
 \frac{\partial}{\partial t} \mathcal{U}_c - \nabla \cdot (\nabla \mathcal{U}_c) &= \mathcal{S}_c \quad \text{on } \Omega_c \\
 \frac{\partial}{\partial t} \mathcal{U}_e - \nabla \cdot (\nabla \mathcal{U}_e) &= \mathcal{S}_e \quad \text{on } \Omega_e \\
 \mathcal{U}_c &= \mathcal{U}_e \quad \text{on } \Gamma \\
 \hat{n} \cdot \nabla \mathcal{U}_c &= \hat{n} \cdot \nabla \mathcal{U}_e \quad \text{on } \Gamma
 \end{aligned}$$

- Processes modeled: The only physics modeled is heat transfer, but it is the error caused by the exchange of information across the interface that is of interest.
 - Specification of initial conditions: Each subdomain requires a standard initial condition $U_0 = U(\vec{x}, t) |_{t=0}$.
 - Specification of boundary conditions: Each subdomain requires standard Dirichlet or Neumann boundary conditions on $\Omega_c \setminus \Gamma$ and $\Omega_e \setminus \Gamma$, respectively.
 - Benchmark type: stability and accuracy of coupling algorithms.
 - Description of quantities in which accuracy is important: The quantity of interest is the difference between the computed solution and the implicitly coupled solution (see accuracy assessment below), measured in a suitable norm. Special attention is paid to the error near the interface, where inaccuracies often appear and grow in time. Also, the error at the end time of longer simulations is more likely to reveal any instability. One could consider (a) the temperature at the midpoint of one domain at a specific time T and time averaged from 0 to T, and (b) the average spatial temperature over one domain at a specific time T and time averaged from 0 to T.
- Principal Features Tested
 - Discretization: The discretization within the subdomains is something standard, like finite volume in space and backward Euler in time. The discretization of the interface conditions (equality of state and normal derivative on the interface) can be varied and tested.
 - Algorithms: The solution algorithm is a blockwise iterative procedure at each timestep. The two subdomains swap information and are then solved independently. This procedure can be done once at each timestep or repeatedly at each timestep in hopes of converging to a better solution for that timestep.
 - Solver: The solver consists of the most accurate direct solve available, with no concern for computational cost. The accuracy and stability issues introduced by the coupling can then be examined without error from inexact solves in the subdomains entering the picture.

- Code: Matlab
- Mathematical Description
 - Problem domain: In 1D, the domain is $x \in [0, 1]$ with the interface between the subdomains at $x = .5$. In 2D the subdomains are two neighboring squares.
 - Ranges of data and parameters: The diffusivity is left at a constant equal to one. The forcing term is set to match the manufactured solution. The size of the grid cells, the size of the timestep, and the number of iterations performed at each timestep are the parameters to be varied.
- Accuracy Assessment: The problem is constructed by starting with a known continuous solution and calculating the corresponding source term. Two forms of discrete solution are then produced. One, called the implicitly coupled solution, is computed by directly inverting the coupled discrete system. The second, called the computed solution, is computed by performing the block iteration described above at each timestep. The error between these solutions can then be measured directly.
- Testing Procedures: The accuracy and stability of the algorithm for various timesteps and numbers of iterations per timestep are tested by direct measurement of error.
- Additional User Information: It is also possible to do multirate examples, where the timestep in one subdomain is an integer multiple of the timestep in the other subdomain.
- References: Contact Brendan Sheehan (brendansheehan6@gmail.com).

Appendix B.3. Conjugate heat transfer

- Contributors: Don Estep and Brendan Sheehan
- Conceptual Description
 - General: This is the flow of a fluid around a cylinder of a different temperature, where the regions are coupled through the heat equations by continuity of state and normal flux. The flow of the fluid guarantees that the effects of error are felt downstream from the object.

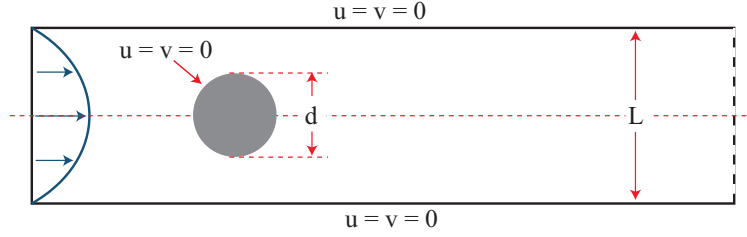


Figure B.25: Illustration of conjugate heat transfer problem

$$\begin{cases} -\mu\Delta\mathbf{u} + \rho_0(\mathbf{u} \cdot \nabla)\mathbf{u} + \nabla p + \rho_0\beta T_F\mathbf{g} = \rho_0(1 + \beta T_0)\mathbf{g}, & x \in \Omega_F, \\ -\nabla \cdot \mathbf{u} = 0, & x \in \Omega_F, \\ -k_F\Delta T_F + \rho_0 c_p(\mathbf{u} \cdot \nabla T_F) = Q_F, & x \in \Omega_F, \\ \begin{cases} T_S = T_F, \\ k_F(\hat{n} \cdot \nabla T_F) = k_S(\hat{n} \cdot \nabla T_S), \end{cases} & x \in \Gamma_I, \\ -k_S\Delta T_S = Q_S, & x \in \Omega_S, \end{cases}$$

The symbols ρ_0 and T_0 are reference values for the density and temperature respectively, μ is the molecular viscosity, β is the coefficient of thermal expansion, c_p is the specific heat, k_F and k_S are the thermal conductivities of the fluid and solid respectively, Q_F and Q_S are source terms, and \hat{n} is the unit normal vector directed into the fluid.

- Processes modeled: Heat transfer and fluid flow.
- Specification of initial conditions: Not required for steady-state problem; required for time-dependent version.
- Specification of boundary conditions: The fluid subdomain requires a known profile at the inflow boundary and uses a zero velocity condition at the outer walls and at the interface with the solid. The temperature/heat flux boundary conditions must be known on $\partial\Omega_F \setminus \Gamma_I$ and are provided on Γ_I by information exchange between the subdomains. This process starts with the initial conditions in the time-dependent case and requires some initial guess in the steady-state case.
- Benchmark type: Coupling two different PDEs discretized with two different discrete methods, evaluating both stability and accuracy.
- Description of quantities in which accuracy is important: Examples of quantities of interest include the average temperature in the solid, or the temperature at a specific point in the solid, or fluid downstream of the solid. One could consider (a) the temperature at the midpoint of the solid at a specific time T and time averaged from 0 to T , (b) the average spatial temperature over the solid at a specific time T and time averaged from 0 to T , and (c) the temperature in the fluid downstream a small ways from the solid at a specific time T and time averaged from 0 to T .

- Principal Features Tested

- Discretization: Different discretizations may be needed for the fluid and solid subdomains, and this is likely to put some stress on the accuracy of the solution.
- Algorithm/solver/code: The subdomains are coupled through an iteration in which temperature and heat flux information is passed from one subdomain to the other. Furthermore, the fluid subdomain has heat transfer coupled to fluid flow. Therefore, either the component solution in the fluid either must come from a solver that can handle this all at once, or an additional iteration between a fluid code and a heat transfer code must take place within the fluid subdomain.

- Mathematical Description

- Problem domain: 2D pipe in which the fluid flows.
- Ranges of data and parameters: The material properties of the fluid and solid are the parameters to be varied, including the thermal conductivity of the solid and fluid, and the viscosity of the fluid.

- Accuracy Assessment/Testing Procedures: A posteriori error analysis is used to assess the accuracy in a quantity of interest. Alternatively, a reference solution on a very fine grid could be used to assess the accuracy of the solution on a coarse grid.

- References: Contact Brendan Sheehan (brendansheehan6@gmail.com).

REFERENCES

References

- [1] D. Brown (Chair), Applied Mathematics at the U.S. Department of Energy: Past, Present, and a View to the Future, Office of Science, U.S. Department of Energy (2008).
URL http://science.energy.gov/~media/ascr/pdf/program-documents/docs/Brown_report_may_08.pdf
- [2] G. Strang, On the construction and comparison of difference schemes, *SIAM J. Numer. Anal.* 5 (1968) 506–517.
- [3] L. F. Richardson, The approximate arithmetical solution by finite differences of physical problems involving differential equations, with an application to the stresses in a masonry dam, *Philosophical Transactions of the Royal Society of London, Series A* 210 (1911) 307–357. doi:10.1098/rsta.1911.0009.
- [4] L. F. Richardson, J. A. Gaunt, The deferred approach to the limit, *Philosophical Transactions of the Royal Society of London, Series A* 226 (1927) 636–646. doi:10.1098/rsta.1927.0008.
- [5] D. A. Knoll, D. E. Keyes, Jacobian-free Newton-Krylov methods: a survey of approaches and applications, *J. Comput. Phys.* 193 (2) (2004) 357–397. doi:10.1016/j.jcp.2003.08.010.
- [6] X.-C. Cai, D. E. Keyes, Nonlinearly preconditioned inexact Newton algorithms, *SIAM J. Sci. Comput.* 24 (2002) 183–200. doi:10.1137/S106482750037620X.
- [7] S. Balay, W. D. Gropp, L. C. McInnes, B. F. Smith, Efficient management of parallelism in object oriented numerical software libraries, in: E. Arge, A. M. Bruaset, H. P. Langtangen (Eds.), *Modern Software Tools in Scientific Computing*, Birkhäuser Press, 1997, pp. 163–202.
- [8] S. Balay, J. Brown, K. Buschelman, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, B. F. Smith, H. Zhang, *PETSc users manual*, Tech. Rep. ANL-95/11 - Revision 3.2, Argonne National Laboratory (2011).
URL <http://www.mcs.anl.gov/petsc>
- [9] A. C. Hindmarsh, P. N. Brown, K. E. Grant, S. L. Lee, R. Serban, D. E. Shumaker, C. S. Woodward, SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers, *ACM Trans. Math. Softw.* 31 (2005) 363–396. doi:10.1145/1089014.1089020.
URL <http://doi.acm.org/10.1145/1089014.1089020>
- [10] C. Woodward et al., SUNDIALS Web page, <https://computation.llnl.gov/casc/sundials/>.
- [11] 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, K. S. Stanley, An overview of the Trilinos project, *ACM Trans. Math. Softw.* 31 (2005) 397–423. doi:10.1145/1089014.1089021.
URL <http://doi.acm.org/10.1145/1089014.1089021>
- [12] M. Heroux et al., Trilinos Web page, <http://trilinos.sandia.gov/>.
- [13] D. Brown, P. Messina (Chairs), Scientific Grand Challenges: Crosscutting Technologies for Computing at the Exascale, Office of Science, U.S. Department of Energy (2010).
URL http://extremecomputing.labworks.org/SumReps/Crosscutting-SumRep-PNNL_20168.pdf
- [14] H. Simon, T. Zacharia, R. Stevens (Co-Chairs), Modeling and Simulation at the Exascale for Energy and the Environment, Office of Science, U.S. Department of Energy (2007).
URL <http://science.energy.gov/~media/ascr/pdf/program-documents/docs/Townhall.pdf>
- [15] W. Washington (Chair), Scientific Grand Challenges: Challenges in Climate Change Science and the Role of Computing at the Extreme Scale, Office of Science, U.S. Department of Energy (2008).
URL <http://extremecomputing.labworks.org/workshops.stm>
- [16] W. Tang, D. Keyes (Chairs), Scientific Grand Challenges: Fusion Energy and the Role of Computing at the Extreme Scale, Office of Science, U.S. Department of Energy (2009).
URL <http://extremecomputing.labworks.org/workshops.stm>
- [17] R. Rosner (Chair), The Opportunities and Challenges of Exascale Computing, Office of Science, U.S. Department of Energy (2010).
URL http://science.energy.gov/~media/ascr/ascac/pdf/reports/Exascale_subcommittee_report.pdf
- [18] C. Farhat, M. Lesoinne, P. L. Tallec, Load and motion transfer algorithms for fluid/structure interaction problems with non-matching discrete interfaces: Momentum and energy conservation, optimal discretization and application to aeroelasticity, *Computer Methods in Applied Mechanics and Engineering* 157 (1-2) (1998) 95–114. doi:10.1016/S0045-7825(97)00216-8.
URL <http://www.sciencedirect.com/science/article/pii/S0045782597002168>
- [19] A. de Boer, A. van Zuijlen, H. Bijl, Review of coupling methods for non-matching meshes, *Comput. Methods Appl. Mech. Eng.* 196(8) (2007) 1515–1525.
- [20] R. Jaiman, X. Jiao, P. Geubelle, E. Loth, Conservative load transfer along curved fluid-solid interface with non-matching meshes, *J. Comp. Phys.* 218(1) (2006) 372–397.
- [21] H.-J. Bungartz, B. Gatzhammer, M. Mehl, T. Neckel, Partitioned simulation of fluid-structure interaction on Cartesian grids, in: H.-J. Bungartz, M. Mehl, M. Schäfer (Eds.), *Fluid-Structure Interaction—Modelling, Simulation, Optimisation, Part II*, LNCSE, Springer, Berlin, Heidelberg, 2010, pp. 255–284.
- [22] D. Scholz, S. Kollmannsberger, A. Düster, E. Rank, Thin solids for fluid-structure interaction, in: H.-J. Bungartz, M. Schäfer (Eds.), *Fluid-Structure Interaction*, number 53 in *Lecture Notes in Computational Science and Engineering*, Springer, Berlin, Heidelberg, 2006, pp. 294–335.
- [23] F. Baaijens, A fictitious domain/mortar element method for fluid-structure interaction, *International Journal for Numerical Methods in Fluids* 35(7) (2001) 743–761.
- [24] T. Klöppel, A. Popp, U. Kttler, W. Wall, Fluid-structure interaction for non-conforming interfaces based on a dual mortar formulation, *Computer Methods in Applied Mechanics and Engineering* Accepted.
- [25] M. Ross, Coupling and simulation of acoustic fluid-structure interaction systems using localized Lagrange multipliers, PhD thesis, University of Colorado (2006).
- [26] B. Irons, R. Tuck, A version of the Aitken accelerator for computer iteration, *International Journal of Numerical Methods in Engineering* 1 (1969) 275–277.

REFERENCES

- [27] M. Schäfer, D. Stempel, G. Becker, P. Pironkov, Efficient numerical simulation and optimization of fluid-structure interaction, in: H.-J. Bungartz, M. Mehl, M. Schäfer (Eds.), *Fluid-Structure Interaction II*, number 73 in *Lecture Notes in Computational Science and Engineering*, Springer, Berlin, Heidelberg, 2010, pp. 131–158.
- [28] W. Wall, D. Mok, E. Ramm, Accelerated iterative substructuring schemes for instationary fluid-structure interaction, *Computational Fluid and Solid Mechanics (2001)* 1325–1328.
- [29] J. Degroote, K. Bathe, J. Vierendeels, Performance of a new partitioned procedure versus a monolithic procedure in fluid-structure interaction, *Comput. Struct.* 87 (2009) 793–801.
- [30] E. H. van Brummelen, Added mass effects of compressible and incompressible flows in fluid-structure interaction, *Journal of Applied Mechanics* 76 (2) (2009) 021206. doi:10.1115/1.3059565.
- [31] E. van Brummelen, Partitioned iterative solution methods for fluid-structure interaction, *Int. J. Numer. Meth. Fluids* 65 (2010) 3–27.
- [32] P. Geuzaine, G. Brown, C. Harris, C. Farhat, Aeroelastic dynamic analysis of a full F-16 configuration for various flight conditions, *AIAA Journal* 41 (2003) 363–371.
- [33] C. Farhat, P. Geuzaine, G. Brown, Application of a three-field nonlinear fluid-structure formulation to the prediction of the aeroelastic parameters of an F-16 fighter, *Computers and Fluids* 32 (2003) 3–29.
- [34] C. Farhat, A. Rallu, K. Wang, T. Belytschko, Robust and provably second-order explicit-explicit and implicit-explicit staggered time-integrators for highly nonlinear fluid-structure interaction problems, *International Journal for Numerical Methods in Engineering* 84 (2010) 73–107.
- [35] N. Maman, C. Farhat, Matching fluid and structure meshes for aeroelastic computations: A parallel approach, *Computers & Structures* 54 (4) (1995) 779–785.
- [36] C. Farhat, K. Pierson, C. Degand, Multidisciplinary simulation of the maneuvering of an aircraft, *Engineering with Computers* 17 (2001) 16–27.
- [37] C. Farhat, C. Degand, B. Koobus, M. Lesoinne, Torsional springs for two-dimensional dynamic unstructured fluid meshes, *Computer Methods in Applied Mechanics and Engineering* 163 (1-4) (1998) 231–245. doi:10.1016/S0045-7825(98)00016-4. URL <http://www.sciencedirect.com/science/article/pii/S0045782598000164>
- [38] K. Wang, A. Rallu, J.-F. Gerbeau, C. Farhat, Algorithms for interface treatment and load computation in embedded boundary methods for fluid and fluid-structure interaction problems, *International Journal for Numerical Methods in Fluids* In press.
- [39] C. Farhat, G. van der Zee, P. Geuzaine, Provably second-order time-accurate loosely-coupled solution algorithms for transient nonlinear computational aeroelasticity, *Computer Methods in Applied Mechanics and Engineering* 195 (2006) 1973–2001.
- [40] G. Hansen, C. Newman, D. Gaston, C. Permann, An implicit solution framework for reactor fuel performance simulation, in: *20th International Conference on Structural Mechanics in Reactor Technology (SMIRT 20)*, paper 2045, Espoo (Helsinki), Finland, 2009.
- [41] C. Newman, D. Gaston, G. Hansen, Computational foundations for reactor fuel performance modeling, in: *American Nuclear Society 2009 International Conference on Advances in Mathematics, Computational Methods, and Reactor Physics*, Saratoga Springs, NY, 2009.
- [42] C. Newman, G. Hansen, D. Gaston, Three dimensional coupled simulation of thermomechanics, heat, and oxygen diffusion in UO₂ nuclear fuel rods, *Journal of Nuclear Materials* 392 (2009) 6–15. doi:10.1016/j.jnucmat.2009.03.035.
- [43] R. L. Williamson, D. A. Knoll, Enhancing the ABAQUS thermomechanics code to simulate steady and transient fuel rod behavior, in: *Proceedings of Top Fuel, 2009*, paper 2072.
- [44] G. Hansen, A Jacobian-free Newton Krylov method for mortar-discretized thermomechanical contact problems, *Journal of Computational Physics* 230 (2011) 6546–6562. doi:10.1016/j.jcp.2011.04.038.
- [45] M. R. Tonks, G. Hansen, D. Gaston, C. Permann, P. Millett, D. Wolf, Fully-coupled engineering and mesoscale simulations of thermal conductivity in UO₂ fuel using an implicit multiscale approach, *Journal of Physics: Conference Series* 180 (1) (2009) 012078. doi:10.1088/1742-6596/180/1/012078.
- [46] M. Tonks, D. Gaston, C. Permann, P. Millett, G. Hansen, D. Wolf, A coupling methodology for mesoscale-informed nuclear fuel performance codes, *Nuclear Engineering and Design* 240 (10) (2010) 2877–2883. doi:10.1016/j.nucengdes.2010.06.005.
- [47] D. Gaston, C. Newman, G. Hansen, D. Lebrun-Grandié, MOOSE: A parallel computational framework for coupled systems of nonlinear equations, *Nuclear Engineering and Design* 239 (2009) 1768–1778. doi:10.1016/j.nucengdes.2009.05.021.
- [48] D. Gaston, G. Hansen, S. Kadioglu, D. Knoll, C. Newman, H. Park, C. Permann, W. Taitano, Parallel multiphysics algorithms and software for computational nuclear engineering, *Journal of Physics: Conference Series* 180 (1) (2009) 012012. doi:10.1088/1742-6596/180/1/012012.
- [49] M. W. Gee, G. A. Hansen, Moertel mortar methods package. URL <http://trilinos.sandia.gov/packages/moertel>
- [50] A. Siegel, T. Tautges, A. Caceres, D. Kaushik, P. Fischer, G. Palmiotti, M. Smith, J. Ragusa, Software design of SHARP, in: *Joint International Topical Meeting of Mathematics & Computation and Supercomputing in Nuclear Applications, 2007*.
- [51] P. Fischer, J. Lottes, A. Siegel, G. Palmiotti, Large eddy simulation of wire wrapped fuel pins, in: *Joint International Topical Meeting on Mathematics & Computation and Supercomputing in Nuclear Applications, 2007*.
- [52] P. Fischer, J. Lottes, D. Pointer, A. Siegel, Petascale algorithms for reactor hydrodynamics, *Journal of Physics: Conference Series* 125 (1) (2008) 012076. URL <http://stacks.iop.org/1742-6596/125/i=1/a=012076>
- [53] D. Kaushik, M. Smith, A. Wollaber, B. Smith, A. Siegel, W. S. Yang, Enabling high-fidelity neutron transport simulations on petascale architectures, in: *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis, SC '09, ACM, New York, NY, USA, 2009*, pp. 67:1–67:12. doi:<http://doi.acm.org/10.1145/1654059.1654128>.
- [54] P. Fischer, G. W. Kruse, F. Loth, Spectral element methods for transitional flows in complex geometries, *J. Sci. Comput.* 17 (2002) 81–98.
- [55] P. Fischer, Nek5000. URL <https://nek5000.mcs.anl.gov>
- [56] T. J. Tautges, R. Meyers, K. Merkle, C. Stimpson, C. Ernst, MOAB: a Mesh-Oriented database, SAND2004-1592, Sandia National Laboratories, report (Apr. 2004).

REFERENCES

- [57] T. J. Tautges, MOAB wiki.
URL <http://trac.mcs.anl.gov/projects/ITAPS/wiki/MOAB>
- [58] T. J. Tautges, J. Kraftcheck, N. Bertram, V. Sachdeva, J. Magerlein, Mesh interface resolution and ghost exchange in a parallel mesh representation, in: 26th IEEE International Parallel & Distributed Processing Symposium, Shanghai, China, 2011.
- [59] R. Rosner (PI), Center for Exascale Simulation of Advanced Reactors, Website.
URL <http://cesar.mcs.anl.gov/>
- [60] G. Lu, E. Kaxiras, Overview of multiscale simulations of materials, in: M. Rieth, W. Schommers (Eds.), Handbook of Theoretical and Computational Nanotechnology, Vol. X, 2005, pp. 1–33.
- [61] F. Abraham, J. Broughton, N. Bernstein, E. Kaxiras, Spanning the length scales in dynamic simulation, *Comput. in Phys.* 12 (1998) 538–546.
- [62] J. P. Hirth, J. Lothe, Theory of dislocations, Krieger Publishing Co., Malabar, Florida, 1992.
- [63] R. Peierls, The size of a dislocation, *Proc. Phys. Soc. London* 52 (1940) 34–37.
- [64] F. R. N. Nabarro, Dislocations in a simple cubic lattice, *Proc. Phys. Soc. London* 59 (1947) 256–272.
- [65] B. Devincere, L. Kubin, Mesoscopic simulations of dislocations and plasticity, *Mater. Sci. Eng.* 8 (1994) A234–236.
- [66] V. Bulatov, F. F. Abraham, L. Kubin, S. Yip, Connecting atomistic and mesoscopic simulations of crystal plasticity, *Nature* 391 (1998) 669–681.
- [67] R. Car, M. Parrinello, Unified approach for molecular-dynamics and density-functional theory, *Phys. Rev. Lett.* 55 (1985) 2471–2474.
- [68] E. B. Tadmor, M. Ortiz, R. Phillips, Quasicontinuum analysis of defects in solids, *Philos. Mag. A* 73 (1996) 1529–1563.
- [69] G. S. Smith, E. B. Tadmor, N. Bernstein, E. Kaxiras, Multiscale simulation of silicon nanoindentation, *Acta Materialia* 49 (2001) 4089.
- [70] E. B. Tadmor, U. V. Waghmare, G. S. Smith, E. Kaxiras, Polarization switching in pbtio3: an ab initio finite element simulation, *Acta Materialia* 50 (2002) 2989.
- [71] A. F. Voter, A method for accelerating the molecular dynamics simulation of infrequent events, *J. Chem. Phys.* 106 (1997) 4665–4677.
- [72] A. F. Voter, Hyperdynamics: Accelerated molecular dynamics of infrequent events, *Phys. Rev. Lett.* 78 (1997) 3908–3911.
- [73] J. J. Kasianowicz, E. Brandin, D. Branton, D. Deamer, Characterization of individual polynucleotide molecules using a membrane channel, *Proc. Nat. Acad. Sci. (USA)* 93 (1996) 13770–13773.
- [74] A. Meller, L. Nivon, E. Brandin, J. Golovchenko, D. Branton, Rapid nanopore discrimination between single polynucleotide molecules, *Proc. Nat. Acad. Sci.* 97 (2000) 1079–1084.
- [75] J. Li, M. Gershow, D. Stein, E. Brandin, J. Golovchenko, DNA molecules and configurations in a solid-state nanopore microscope, *Nat. Mater.* 2 (2003) 611–615.
- [76] A. J. Storm, C. Storm, J. Chen, H. Zandbergen, J.-F. Joanny, C. Dekker, Fast DNA translocation through a solid-state nanopore, *Nanolett* 5 (2005) 1193–1197.
- [77] C. Dekker, Solid state nanopores, *Nature Nanotech.* 2 (2007) 209–215.
- [78] M. Fyta, S. Melchionna, S. Succi, Translocation of biomolecules through solid-state nanopores: Theory meets experiments, *J. Polymer Sci. B: Polymer Phys.* 49 (2011) 985–1011.
- [79] S. Succi, The Lattice Boltzmann Equation for Fluid Dynamics and Beyond, Clarendon Press, Oxford, 2001.
- [80] R. Benzi, S. Succi, M. Vergassola, The lattice boltzmann equation: theory and applications, *Physics Reports* 222 (1992) 145–197.
- [81] M. G. Fyta, S. Melchionna, E. Kaxiras, S. Succi, Multiscale coupling of molecular dynamics and hydrodynamics: Application to dna translocation through a nanopore, *Multiscale Model. & Sim.* 5 (2006) 1156–1173.
- [82] A. Peters, S. Melchionna, E. Kaxiras, J. Latt, J. Sircar, M. Bernaschi, M. Bisson, S. Succi, Multiscale simulation of cardiovascular flows on the IBM Bluegene/P: full heart-circulation system at near red-blood cell resolution, in: 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis, New Orleans, 2010. doi:10.1109/SC.2010.33.
- [83] International Thermonuclear Experimental Reactor (ITER) homepage, <http://www.iter.org> (2011).
- [84] D. E. Post, D. B. Batchelor, R. B. Bramley, J. R. Cary, R. H. Cohen, P. Colella, S. C. Jardin, Report of Fusion Simulation Project Steering Committee (August 2004).
URL <http://w3.pppl.gov/usjapanim/FSPReport.pdf>
- [85] J. R. Cary, A. Hakim, M. Miah, S. Kruger, A. Pletzer, S. Shasharina, S. Vadlamani, A. Pankin, R. Cohen, T. Epperly, T. Rognlien, R. Groebner, S. Balay, L. McInnes, H. Zhang, FACETS - a Framework for Parallel Coupling of Fusion Components, in: The 18th Euromicro International Conference on Parallel, Distributed and Network-Based Computing, Pisa, Italy, 2009.
- [86] A. H. Hakim, T. D. Rognlien, R. Groebner, J. Carlsson, J. Cary, S. Kruger, M. Miah, A. Pankin, A. Pletzer, S. Shasharina, S. Vadlamani, R. Cohen, T. Epperly, Submitted for publication to *Physics of Plasmas* (2011).
- [87] J. Cary (PI), Framework Application for Core-Edge Transport Simulations (FACETS), proto-FSP project, see <http://www.facetsproject.org>.
- [88] J. Rubin, Transport of reacting solutes in porous media: Relation between mathematical nature of problem formulation and chemical nature of reactions, *Water Resources Research* 19(5) (1983) 1231–1252. doi:10.1029/WR019i005p01231.
- [89] G. A. Cederberg, R. L. Street, J. O. Leckie, A groundwater mass transport and equilibrium chemistry model for multicomponent systems, *Water Resources Research* 21(8) (1985) 1095–1104. doi:10.1029/WR021i008p01095.
- [90] P. C. Lichtner, Continuum model for simultaneous chemical reactions and mass transport in hydrothermal systems, *Geochimica et Cosmochimica Acta* 49(3) (1985) 779–800. doi:10.1016/0016-7037(85)90172-3.
- [91] G. T. Yeh, V. S. Tripathi, A critical evaluation of recent developments in hydrogeochemical transport models of reactive multichemical components, *Water Resources Research* 25(1) (1989) 93–108. doi:10.1029/WR025i001p00093.
- [92] C. I. Steefel, K. T. MacQuarrie, Approaches to modeling of reactive transport in porous media, *Reviews in Mineralogy* 34 (1996) 83–129.
- [93] A. J. Valocchi, M. Malmstead, Accuracy of operator splitting for advection-dispersion-reaction problems, *Water Resources Research* 28(5) (1992) 1471–1476. doi:10.1029/92WR00423.
- [94] D. A. Barry, C. T. Miller, P. Culligan-Hensley, Temporal discretisation errors in non-iterative split-operator approaches to solving chemical reaction groundwater transport models, *Journal of Contaminant Hydrology* 22(1–2) (1996) 1–17. doi:10.1016/0169-7722(95)00062-3.

REFERENCES

- [95] J. Kanney, C. Miller, D. Barry, Convergence of iterative split-operator approaches for approximating nonlinear reactive transport problems, *Advances in Water Resources* 26(3) (2003) 247–261. doi:10.1016/S0309-1708(02)00162-8.
- [96] J. Kanney, C. Miller, D. Barry, Comparison of fully coupled approaches for approximating nonlinear transport and reaction problems, *Advances in Water Resources* 26(4) (2003) 353–372. doi:10.1016/S0309-1708(02)00188-4.
- [97] J. Carayrou, R. Mose, P. Behra, Operator-splitting procedures for reactive transport and comparison of mass balance errors, *Journal of Contaminant Hydrology* 68(3–4) (2004) 239–268. doi:10.1016/S0169-7722(03)00141-4.
- [98] B. A. Robinson, H. S. Viswanathan, A. J. Valocchi, Efficient numerical techniques for modeling multicomponent ground-water transport based upon simultaneous solution of strongly coupled subsets of chemical components, *Advances in Water Resources* 23(4) (2000) 307–324. doi:10.1016/S0309-1708(99)00034-2.
- [99] G. Hammond, P. Lichtner, Field-scale modeling for the natural attenuation of uranium at the hanford 300 area using high performance computing, *Water Resources Research* 46 (2010) W09527. doi:10.1029/2009WR008819.
- [100] G. Hammond, A. Valocchi, P. Lichtner, Application of Jacobian-free Newton-Krylov with physics-based preconditioning to biogeochemical transport, *Advances in Water Resources* 28(4) (2005) 359–376. doi:10.1016/j.advwatres.2004.12.001.
- [101] M. T. Balhoff, K. E. Thompson, M. Hjortso, Coupling pore-scale networks to continuum-scale models of porous media, *Computers & Geosciences* 33 (2007) 393–410.
- [102] A. M. Tartakovsky, D. M. Tartakovsky, T. D. Scheibe, P. Meakin, Hybrid simulations of reaction-diffusion systems in porous media, *SIAM Journal on Scientific Computing* 30 (6) (2008) 2799–2816.
- [103] M. T. Balhoff, S. G. Thomas, M. F. Wheeler, Mortar coupling and upscaling of pore-scale models, *Computers & Geosciences* 12 (2008) 15–27.
- [104] I. Battiato, D. M. Tartakovsky, A. M. Tartakovsky, T. D. Scheibe, Hybrid models of reactive transport in porous and fractured media, *Advances in Water Resources* 34 (9) (2011) 1140–1150. doi:10.1016/j.advwatres.2011.01.012.
- [105] A. M. Tartakovsky, T. D. Scheibe, Dimension reduction method for advection-diffusion-reaction systems, *Advances in Water Resources* In press. doi:10.1016/j.advwatres.2011.07.011.
- [106] W. E. B. Engquist, Z. Y. Huang, Heterogeneous multiscale method: A general methodology for multiscale modeling, *Physical Review B* 67 (9) (2003) 092101.
- [107] R. A. Freeze, R. L. Harlan, Blueprint for a physically-based digitally simulated, hydrologic response model, *Journal of Hydrology* 9 (1969) 237–258.
- [108] J. E. VanderKwaak, K. Loague, Hydrologic-response simulations for the r-5 catchment with a comprehensive physics-based model, *Water Resources Research* 37 (4) (2001) 999–1013.
- [109] S. Panday, P. S. Huyakorn, A fully coupled physically-based spatially-distributed model for evaluating surface/subsurface flow, *Advances in Water Resources* 27 (4) (2004) 361–382.
- [110] J. P. Jones, E. A. Sudicky, A. E. Brookfield, Y. J. Park, An assessment of the tracer-based approach to quantifying groundwater contributions to streamflow, *Water Resour. Res.* 42 (2006) 10.1029/2005wr004130.
- [111] S. J. Kollet, R. M. Maxwell, Integrated surface-groundwater flow modeling: A free-surface overland flow boundary condition in a parallel groundwater flow model, *Advances in Water Resources* 29 (7) (2006) 945–958. doi:10.1016/j.advwatres.2005.08.006.
- [112] Y. Qu, C. J. Duffy, A semidiscrete finite volume formulation for multiprocess watershed simulation, *Water Resources Research* 43(W08419) (2007) 18.
- [113] S. J. Kollet, R. M. Maxwell, Capturing the influence of groundwater dynamics on land surface processes using an integrated, distributed watershed model, *Water Resources Research* 44 (2) (2008) 18.
- [114] R. M. Maxwell, N. L. Miller, Development of a coupled land surface and groundwater model, *Journal of Hydrometeorology* 6 (3) (2005) 233–247.
- [115] R. M. Maxwell, F. K. Chow, S. J. Kollet, The groundwater-land-surface-atmosphere connection: Soil moisture effects on the atmospheric boundary layer in fully-coupled simulations, *Advances in Water Resources* 30 (12) (2007) 2447–2466. doi:10.1016/j.advwatres.2007.05.018.
URL <http://www.sciencedirect.com/science/article/pii/S0309170807000954>
- [116] R. M. Maxwell, J. K. Lundquist, J. D. Mirocha, S. G. Smith, C. S. Woodward, A. F. B. Tompson, Development of a coupled groundwater-atmosphere model, *Monthly Weather Review* 139 (1) (2011) 96116. doi:10.1175/2010MWR3392.1.
- [117] B. A. Ebel, B. B. Mirus, C. S. Heppner, J. E. VanderKwaak, K. Loague, First-order exchange coefficient coupling for simulating surface water-groundwater interactions: Parameter sensitivity and consistency with a physics-based approach, *Hydrological Processes* 23 (13) (2009) 1949–1959. doi:10.1002/Hyp.7279.
- [118] M. Camporese, C. Paniconi, M. Putti, P. Saladin, A comparison of data assimilation techniques for a coupled model of surface and subsurface flow, *Vadose Zone Journal* 8 (4) (2009) 837–845.
- [119] M. Kumar, C. J. Duffy, K. M. Salvage, A second order accurate, finite volume based, integrated hydrologic modeling (fihm) framework for simulation of surface and subsurface flow, *Vadose Zone Journal* 8 (4) (2009) 873–890.
- [120] Y. J. Park, E. A. Sudicky, S. Panday, G. Matanga, Implicit sub-time stepping for solving the nonlinear equations of flow in an integrated surface-subsurface system, *Vadose Zone Journal* 8 (4) (2009) 825–836.
- [121] S. J. Kollet, R. M. Maxwell, C. S. Woodward, S. Smith, J. Vanderborcht, H. Vereecken, C. Simmer, Proof of concept of regional scale hydrologic simulations at hydrologic resolution utilizing massively parallel computer resources, *Water Resources Research* 46 (2010) W04201. doi:10.1029/2009wr008730.
- [122] S. F. Ashby, R. D. Falgout, A parallel multigrid preconditioned conjugate gradient algorithm for groundwater flow simulations, *Nucl. Sci. Eng.* 124 (1) (1996) 145–159.
- [123] J. E. Jones, C. S. Woodward, Newton-Krylov-multigrid solvers for large-scale, highly heterogeneous, variably saturated flow problems, *Adv. Water Res.* 24 (7) (2001) 763–774.
- [124] Y. Dai, X. Zeng, R. E. Dickinson, I. Baker, G. B. Bonan, M. G. Bosilovich, A. S. Denning, P. A. Dirmeyer, P. R. Houser, G. Niu, K. W. Oleson, C. A. Schlosser, Z.-L. Yang, The common land model, *Bull. Am. Meteorol. Soc.* 84 (8) (2003) 1013–1023.

REFERENCES

- [125] K. Mitchell, The community noah land-surface model user's guide, release 2.7.1, http://www.emc.ncep.noaa.gov/mmb/gcp/noah_lsm/Noah_LSM_USERGUIDE_2.7.1.htm (2005).
- [126] S. Solomon, D. Qin, M. Manning, Z. Chen, M. Marquis, K. Avery, M. Tignor, e. H. L. Miller, Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, Cambridge University Press, Cambridge, United Kingdom, 2007.
- [127] J. M. Dennis, M. Vertenstein, P. Worley, A. A. Mirin, A. P. Craig, R. Jacob, S. Mickelson, Computational performance of ultra-high resolution capability in the community earth system model, *Int. J. High Perf. Comput. Appl.* Accepted.
- [128] P. R. Gent, G. Danabasoglu, L. J. Donner, M. M. Holland, E. C. Hunke, S. R. Jayne, D. M. Lawrence, R. B. Neale, P. J. Rasch, M. Vertenstein, P. H. Worley, Z.-L. Yang, M. Zhang, The Community Climate System Model version 4, *Journal of Climate* Under review.
- [129] T. Ringler, D. W. Jacobsen, M. Gunzburger, L. Ju, M. Duda, W. Skamarock, Exploring a multi-resolution modeling approach within the shallow-water equations, *Monthly Weather Review* In press, (pdf).
- [130] K. J. Evans, D. W. I. Rouson, A. G. Salinger, M. A. Taylor, W. Weijer, J. B. White, A scalable and adaptable solution framework within components of the community climate system model, in: *ICCS (2)*, 2009, pp. 332–341.
- [131] L. W. Horowitz, S. Walters, D. L. Mauzerall, L. K. Emmons, P. J. Rasch, C. Granier, X. Tie, J.-F. Lamarque, M. G. Schultz, G. S. Tyndall, J. J. Orlando, G. P. Brasseur, A global simulation of tropospheric ozone and related tracers: Description and evaluation of mozart, version 2, *J. Geophys. Res.* 108(D24) (2003) 4784. doi:10.1029/2002JD002853.
- [132] G. C. Pomraning, *The Equations of Radiation Hydrodynamics*, Pergamon Press, Oxford, New York, 1973.
- [133] D. Mihalas, B. W. Mihalas, *Foundations of Radiation Hydrodynamics*, Oxford University Press, New York, Oxford, 1984.
- [134] J. I. Castor, *Radiation Hydrodynamics*, Cambridge University Press, Cambridge, UK, 2004.
- [135] R. P. Drake, F. W. Doss, R. G. McClarren, M. L. Adams, N. Amato, D. Bingham, C. C. Chou, C. DiStefano, K. Fidkowski, B. Fryxell, T. I. Gombosi, M. J. Grosskopf, J. P. Holloway, B. van der Holst, C. M. Huntington, S. Karnia, C. M. Krauland, C. C. Kuranza, E. Larsen, B. van Leer, B. Mallick, D. Marion, W. Martin, J. E. Morel, E. S. Myra, V. Nair, K. G. Powell, L. Raushberger, P. Roe, E. Rutter, I. V. Sokolov, Q. Stout, B. R. Torralva, G. Toth, K. Thornton, A. J. Visco, Radiative effects in radiative shocks in shock tubes, *High Energy Density Physics* 7 (2011) 130–140.
- [136] E. S. Myra, B. Fryxell, M. J. Grosskopf, B. van der Holst, G. Toth, I. V. Sokolov, A study of shock morphology in x-ray-driven radiative shock-tube simulations, *High Energy Density Physics* submitted.
- [137] F. Graziani, Radiation diffusion: an overview of physical and numerical concepts, in: A. Mezzacappa, G. M. Fuller (Eds.), *Open Issues in Core Collapse Supernova Theory*, World Scientific Publishing Co., Singapore, 2005, pp. 29–66.
- [138] R. L. Bowers, J. R. Wilson, A numerical model for stellar core collapse calculations, *Astrophysical Journal Supplement Series* 50 (1982) 115–159.
- [139] S. W. Bruenn, Stellar core collapse—numerical model and infall epoch, *Astrophysical Journal Supplement Series* 58 (1985) 771–841.
- [140] E. S. Myra, S. A. Bludman, Y. Hoffman, N. Lichtenstadt, I. and Sack, K. A. van Riper, The effect of neutrino transport on the collapse of iron stellar cores, *Astrophysical Journal* 318 (1987) 744–759.
- [141] N. J. Turner, J. M. Stone, A module for radiation hydrodynamic calculations with zeus-2d using flux-limited diffusion, *Astrophysical Journal Supplement Series* 135 (2001) 95–107.
- [142] F. D. Swesty, E. S. Myra, A Numerical Algorithm for Modeling Multigroup Neutrino-Radiation Hydrodynamics in Two Spatial Dimensions, *Astrophysical Journal Supplement Series* 181 (2009) 1–52.
- [143] B. van der Holst, G. Toth, I. V. Sokolov, K. G. Powell, J. P. Holloway, E. S. Myra, Q. Stout, M. L. Adams, J. E. Morel, R. P. Drake, Crash: A block-adaptive-mesh code for radiative shock hydrodynamics—implementation and verification, *Astrophysical Journal Supplement Series* 194 (2011) 23.
- [144] E. E. Lewis, J. W. F. Miller, *Computational Methods of Neutron Transport*, American Nuclear Society, Inc., La Grange Park, IL, USA, 1993.
- [145] D. R. Reynolds, F. D. Swesty, C. S. Woodward, A Newton-Krylov solver for implicit solution of hydrodynamics in core collapse supernovae, *Journal of Physics: Conference Series* 125 (1) (2008) 012085.
- [146] A. B. Watts, *Isostasy and flexure of the lithosphere*, Cambridge University Press, 2001.
- [147] G. Schubert, D. L. Turcotte, P. Olson, *Mantle convection in the earth and planets*, Cambridge University Press, 2001.
- [148] P. E. van Keken, S. D. King, H. Schmeling, U. R. Christensen, D. Neumeister, M. Doin, A comparison of methods for the modeling of thermochemical convection, *Journal of Geophysical Research*.
URL <http://www.agu.org/pubs/crossref/1997/97JB01353.shtml>
- [149] S. D. King, C. Lee, P. E. v. Keken, W. Leng, S. Zhong, E. Tan, N. Tosi, M. C. Kameyama, A community benchmark for 2-D cartesian compressible convection in the earth's mantle, *Geophysical Journal International* 180 (1) (2010) 73–87. doi:10.1111/j.1365-246X.2009.04413.x.
- [150] J. Tarduno, H. Bunge, N. Sleep, U. Hansen, The bent Hawaiian-Emperor hotspot track: Inheriting the mantle wind, *Science* 324 (5923) (2009) 50–53. doi:10.1126/science.1161256.
URL <http://www.sciencemag.org/cgi/content/abstract/324/5923/50>
- [151] D. R. Davies, C. R. Wilson, S. C. Kramer, Fluidity: A fully unstructured anisotropic adaptive mesh computational modeling framework for geodynamics, *Geochemistry Geophysics Geosystems* 12 (2011) 20 PP. doi:201110.1029/2011GC003551.
URL <http://www.agu.org/pubs/crossref/2011/2011GC003551.shtml>
- [152] U. R. Christensen, D. A. Yuen, Layered convection induced by phase transitions, *Journal of Geophysical Research* 90 (B12) (1985) PP. 10,291–10,300. doi:198510.1029/JB090iB12p10291.
URL <http://www.agu.org/pubs/crossref/1985/JB090iB12p10291.shtml>
- [153] G. A. Glatzmaier, P. H. Roberts, A three-dimensional convective dynamo solution with rotating and finitely conducting inner core and mantle, *Physics of the Earth and Planetary Interiors* 91 (1-3) (1995) 63–75. doi:10.1016/0031-9201(95)03049-3.
URL <http://elsevier-apps.sciverse.com/GoogleMaps/index.jsp?doi=10.1016/0031-9201%2895%2903049-3>
- [154] W. Kuang, J. Bloxham, Numerical modeling of magnetohydrodynamic convection in a rapidly rotating spherical shell: Weak and strong field dynamo action, *Journal of Computational Physics* 153 (1) (1999) 51–81. doi:10.1006/jcph.1999.6274.

REFERENCES

- [155] B. J. Kaus, H. Mhlhaus, D. A. May, A stabilization algorithm for geodynamic numerical simulations with a free surface, *Physics of the Earth and Planetary Interiors* 181 (1-2) (2010) 12–20. doi:10.1016/j.pepi.2010.04.007.
URL <http://www.sciencedirect.com/science/article/pii/S0031920110000877>
- [156] B. Blankenbach, F. Busse, U. Christensen, L. Cserepes, D. Gunkel, U. Hansen, H. Harder, G. Jarvis, M. Koch, G. Marquart, D. Moore, P. Olson, H. Schmeling, T. Schnaubelt, A benchmark comparison for mantle convection codes, *Geophysical Journal International* 98 (1) (1989) 23–38. doi:10.1111/j.1365-246X.1989.tb05511.x.
URL <http://onlinelibrary.wiley.com/doi/10.1111/j.1365-246X.1989.tb05511.x/abstract>
- [157] H. Bunge, M. A. Richards, C. Lithgow-Bertelloni, J. R. Baumgardner, S. P. Grand, B. A. Romanowicz, Time scales and heterogeneous structure in geodynamic earth models, *Science* 280 (5360) (1998) 91–95. doi:10.1126/science.280.5360.91.
URL <http://www.sciencemag.org/content/280/5360/91.abstract>
- [158] M. Spiegelman, Flow in deformable porous media. part 1. simple analysis, *Journal of Fluid Mechanics* 247 (17) (1993) 38.
- [159] M. Spiegelman, Flow in deformable porous media. part 2. numerical analysis the relationship between shock waves and solitary waves, *Journal of Fluid Mechanics* 247 (39) (1993) 63.
- [160] R. Katz, M. Knepley, B. Smith, M. Spiegelman, E. Coon, Numerical simulation of geodynamic processes with the portable extensible toolkit for scientific computation, *Physics of The Earth and Planetary Interiors* 163 (1-4) (2007) 52–68. doi:10.1016/j.pepi.2007.04.016.
URL <http://www.sciencedirect.com/science/article/B6V6S-4NPG0F1-4/2/5dc82d8bcbe047c302c9d889ba3f45c1>
- [161] R. F. Katz, M. Spiegelman, C. H. Langmuir, A new parameterization of hydrous mantle melting, *Geochemistry Geophysics Geosystems* 4 (2003) 19. doi:200310.1029/2002GC000433[Citation].
URL <http://www.agu.org/pubs/crossref/2003/2002GC000433.shtml>
- [162] J. Jin, *The Finite Element Method in Electromagnetics*, John Wiley & Sons, New York, 2002.
- [163] R. B. Morgan, On restarting the arnoldi method for largenon-symmetric eigenvalue problems, *Mathematics of Computation* 65 (215) (1996) 1213–1230.
- [164] A. Gupta, G. Karypis, V. Kumar, A highly scalable parallel algorithm for sparse matrix factorization, *IEEE Transactions on Parallel and Distributed Systems* 8 (5) (1997) 502D520.
- [165] Y. Sun, N. Folwell, Z. Li, G. Golub, High precision accelerator cavity design using the parallel eigensolver omega3p, *Proc. of the 18th Annual Review of Progress in Applied Computational Electromagnetics*, 2002.
- [166] G. Cheng, H. Wang, D. Smithe, Rf-thermal combined simulations of a superconducting hom coaxial coupler, in: *Proceedings of 2011 Particle Accelerator Conference*, New York, NY, USA, 2011, TUP107.
- [167] C. E. Reece, E. F. Daly, G. K. Davis, W. R. Hicks, T. Rothgeb, H. L. Phillips, J. Preble, H. Wang, G. Wu, Diagnosis, analysis, and resolution of thermal stability issues with hom couplers on prototype cebaf srf cavities, in: *Proceedings of SRF 2007*, Beijing, October 2007, WEP62, 2007, p. 656.
- [168] V. Akcelik, L.-Q. Lee, Z. Li, C.-K. Ng, K. Ko, G. Cheng, R. Rimmer, H. Wang, Thermal analysis of srf cavity couplers using parallel multiphysics tool TEM3P, in: *Particle Accelerator Conference*, Vancouver, Canada, 2009.
- [169] V. Akcelik, A. C. ana A. Kabel, L.-Q. Lee, Z. Li, C.-K. Ng, L. Xiao, K. Ko, Parallel computation of integrated electromagnetic, thermal and structural effects for accelerator cavities, *Tech. Rep. SLAC-PUB-13280*, Stanford Linear Accelerator Center (SLAC) (2008).
- [170] M. Liepe, W. D. Moeller, S. N. Simrock, Dynamic Lorentz force compensation with a fast piezoelectric tuner, in: *Proceedings of the 2001 Particle Accelerator Conference*, Chicago, 2001.
- [171] L.-Q. Lee, L. Ge, M. Kowalski, Z. Li, C.-K. Ng, G. Schussman, M. Wolf, K. Ko, Solving large sparse linear systems in end-to-end accelerator structure simulations, in: *Proceedings of the 18th International Parallel and Distributed Processing Symposium*, Santa Fe, New Mexico, 2004.
- [172] B.-S. Liao, Z. Bai, L.-Q. Lee, K. Ko, Solving large scale nonlinear eigenvalue problem in next-generation accelerator design, *Tech. Rep. SLAC-PUB-12137*, Stanford Linear Accelerator Center (SLAC) (2006).
- [173] Z. Li, et al., Towards simulations of electromagnetics and beam physics at the petascale, in: *Proc. PAC2007*, Albuquerque, New Mexico, 2007.
- [174] L.-Q. Lee, Z. Li, C. Ng, K. Ko, Omega3p: A parallel finite-element eigenmode analysis code for accelerator cavities, *Tech. Rep. SLAC-PUB-13529*, Stanford Linear Accelerator Center (SLAC) (2009).
- [175] S.-H. Kim, Simulation of quench dynamics in SRF cavities under pulsed operation, in: *Proceedings of the 2003 Particle Accelerator Conference*, 2003.
- [176] J. Lesrel, S. Bousson, T. Junquera, A. Caruette, M. Fouaidy, Study of thermal effects in srf cavities, in: *Proceedings of the 1997 Workshop on RF Superconductivity*, Abano Terme (Padova), Italy, 1997.
- [177] V. Akcelik, G. Biros, O. Ghattas, D. Keyes, K. Ko, L. Lee, E. G. Ng, Adjoint methods for electromagnetic shape optimization of the low-loss cavity for the International Linear Collider, *Journal of Physics: Conference Series* 16 (2005) 435–445.
- [178] V. Akcelik, K. Ko, L.-Q. Lee, Z. Li, C.-K. Ng, L. Xiao, Shape determination for deformed electromagnetic cavities, *Journal of Computational Physics* 227 (2008) 1722D1738.
- [179] K. H. Lee, Design optimization of periodic flows using a time-spectral discrete adjoint method, Ph.D. thesis, Stanford University (2010).
- [180] S. Choi, M. Potsdam, K. H. Lee, G. Iaccarino, J. J. Alonso, Helicopter rotor design using a time-spectral and adjoint-based method, in: *12th AIAA/ISSMO Multidisciplinary and Optimization Conference Sep. 10–12, 2008*, British Columbia, Canada, 2008.
- [181] L. Xiao, K. Ko, K. H. Lee, C. Ng, M. Liepe, N. Valles, Effects of elliptically deformed cell shape in the cornell ERL cavity, in: *15th International Conference on RF Superconductivity*, Chicago IL, July 25–29, 2011, 2011.
- [182] A. Brandt, Multi-level adaptive technique (MLAT) for fast numerical solutions to boundary value problems, in: *Proc. 3rd Int. Conf. on Numerical Methods in Fluid Mechanics*, Vol. 18 of *Lecture Notes in Physics*, Springer-Verlag, 1973, pp. 82–89.
- [183] A. Brandt, Multiscale scientific computation: Review 2001, *Lecture Notes in Computational Science and Engineering* 20 (2002) 3–96.
- [184] U. Trottenberg, C. Oosterlee, A. Schüller, *Multigrid*, Academic Press, 2001.
- [185] A. H. Baker, R. D. Falgout, T. V. Kolev, U. M. Yang, *High Performance Scientific Computing: Algorithms and Applications – A Tribute to*

REFERENCES

- Prof. Ahmed Sameh, Springer, 2011, Ch. Scaling hypre's Multigrid Solvers to 100,000 Cores, also available as LLNL Tech. Rep. LLNL-JRNL-479591.
- [186] V. E. Henson, U. M. Yang, Boomeramg: a parallel algebraic multigrid solver and preconditioner, *Appl. Numer. Math.* 41 (2000) 155–177.
- [187] M. W. Gee, C. M. Siefert, J. J. Hu, R. S. Tuminaro, M. G. Sala, ML 5.0 Smoothed Aggregation User's Guide, Tech. Rep. SAND2006-2649, Sandia National Laboratories (2006).
- [188] R. Falgout et al., hypre users manual, Tech. Rep. Revision 2.7, Lawrence Livermore National Laboratory (2011).
URL <http://www.llnl.gov/CASC/hypre>
- [189] M. R. Hestenes, E. Stiefel, Methods of conjugate gradients for solving linear systems, *J. Res. Nat. Bur. Standards* 49 (6) (1952) 409–436.
- [190] J. K. Reid, On the method of conjugate gradients for the solution of large sparse systems of linear equations, in: J. K. Reid (Ed.), *Large sparse sets of linear equations*, Academic Press, 1971, pp. 231–254.
- [191] O. Axelsson, A generalized SSOR method, *BIT* 12 (1972) 443–467. doi:10.1007/BF01932955.
- [192] P. Concus, G. H. Golub, D. P. O'Leary, A generalized conjugate gradient method for the numerical solution of elliptic partial differential equations, in: J. Bunch, D. Rose (Eds.), *Sparse Matrix Computations*, Academic Press, 1976, pp. 309–332.
- [193] G. H. Golub, D. O'Leary, Some history of the conjugate-gradient and Lanczos algorithms - 1948-1976, *SIAM Rev.* 31 (1) (1989) 50–102. doi:10.1137/1031003.
- [194] R. W. Freund, G. H. Golub, N. M. Nachtigal, Iterative solution of linear systems, *Acta Numerica* (1992) 87–100.
- [195] Y. Saad, H. A. van der Vorst, Iterative solution of linear systems in the 20th century, *J. Comput. Appl. Math.* 123 (1-2, SI) (2000) 1–33. doi:10.1016/S0377-0427(00)00412-X.
- [196] Y. Saad, M. H. Schultz, GMRES: A generalized minimum residual algorithm for solving nonsymmetric linear systems, *SIAM J. Sci. Statist. Comput.* 7 (1986) 856–869.
- [197] H. A. van der Vorst, BI-CGSTAB: A fast and smoothly converging variant of BI-CG for the solution of nonsymmetric linear systems, *SIAM J. Sci. Statist. Comput.* 13 (1992) 631–644.
- [198] V. Faber, T. Manteuffel, Necessary and sufficient conditions for the existence of a conjugate-gradient method, *SIAM J. Numer. Anal.* 21 (2) (1984) 352–362. doi:10.1137/0721026.
- [199] J. Meijerink, H. A. van der Vorst, An iterative solution method for linear systems of which the coefficient matrix is a symmetric M-matrix, *Math. Comp.* 31 (1977) 148–162.
- [200] R. Kettler, Analysis and comparison of relaxation schemes in robust multigrid and preconditioned conjugate-gradient methods, *Lecture Notes In Math.* 960 (1982) 502–534.
- [201] B. Smith, P. Bjørstad, W. Gropp, *Domain Decomposition – Parallel Multilevel Methods for Elliptic Partial Differential Equations*, Cambridge University Press, 1996.
- [202] A. Toselli, O. B. Widlund, *Domain decomposition methods—algorithms and theory*, Springer series in computational mathematics, Springer, 2005.
- [203] M. Dryja, O. Widlund, An additive variant of the Schwarz alternating method for the case of many subregions, Tech. Rep. 339, Courant Institute (1987).
- [204] X.-C. Cai, Some domain decomposition algorithms for nonselfadjoint elliptic and parabolic partial differential equations, Tech. Rep. 461, Courant Institute (1989).
- [205] T. Chan, J. Zou, Additive Schwarz domain decomposition methods for elliptic problems on unstructured meshes, *Numer. Algorithms* 8 (1994) 329–346.
- [206] X.-C. Cai, M. Sarkis, A restricted additive Schwarz preconditioner for general sparse linear systems, *SIAM J. Sci. Comput.* 21 (1999) 792–797.
- [207] J. Bramble, J. Pasciak, J. Xu, Parallel multilevel preconditioners, *Math. Comp.* 55 (1990) 1–21.
- [208] C. Farhat, F. X. Roux, A method of finite element tearing and interconnecting and its parallel solution algorithm, *Int. J. Num. Meth. Eng.* 32 (1991) 1205–1227.
- [209] J. Mandel, Balancing domain decomposition, *Comm. Num. Meth. Eng.* 9 (1993) 233–241.
- [210] C. R. Dohrmann, A preconditioner for substructuring based on constrained energy minimization, *SIAM J. Sci. Comput.* 25 (1) (2003) 246–258.
- [211] S. Banach, Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales, *Fund. Math.* 3 (1922) 133–181.
- [212] A. Granas, J. Dugundji, *Fixed point theory*, Springer monographs in mathematics, Springer, 2003.
- [213] R. S. Dembo, S. C. Eisenstat, T. Steihaug, Inexact Newton methods, *SIAM J. Numer. Anal.* 19 (1982) 400–408.
- [214] S. C. Eisenstat, H. F. Walker, Globally convergent inexact Newton methods, *SIAM J. Optimization* 4 (1994) 393–422.
- [215] S. C. Eisenstat, H. F. Walker, Choosing the forcing terms in an inexact Newton method, *SIAM J. Sci. Comput.* 17 (1996) 16–32.
- [216] P. N. Brown, A local convergence theory for combined inexact-Newton/finite-difference projection methods, *SIAM J. Numer. Anal.* 24 (1987) 407–434.
- [217] C. T. Kelley, *Iterative methods for linear and nonlinear equations*, Frontiers in applied mathematics, Society for Industrial and Applied Mathematics, 1995.
- [218] W. K. Anderson, W. D. Gropp, D. K. Kaushik, D. E. Keyes, B. F. Smith, Achieving high sustained performance in an unstructured mesh CFD application, in: *Proceedings of the 1999 ACM/IEEE conference on Supercomputing (CDROM)*, Supercomputing '99, ACM, New York, NY, USA, 1999. doi:<http://doi.acm.org/10.1145/331532.331600>.
- [219] W. Gropp, D. Keyes, L. C. McInnes, M. D. Tidriri, Globalized Newton-Krylov-Schwarz algorithms and software for parallel implicit CFD, *Int. J. High Perf. Comput. Applics.* 14 (2000) 102–136.
- [220] M. Pernice, M. D. Tocci, A multigrid-preconditioned Newton-Krylov method for the incompressible Navier-Stokes equations, *SIAM J. Sci. Comput.* 23 (2001) 398–418. doi:10.1137/S1064827500372250.
- [221] L. Chacón, D. A. Knoll, A 2d high- β Hall MHD implicit nonlinear solver, *Journal of Computational Physics* 188 (2) (2003) 573 – 592. doi:10.1016/S0021-9991(03)00193-1.
URL <http://www.sciencedirect.com/science/article/pii/S0021999103001931>

REFERENCES

- [222] V. A. Mousseau, Implicitly balanced solution of the two-phase flow equations coupled to nonlinear heat conduction, *J. Comput. Phys.* 200 (2004) 104–132. doi:10.1016/j.jcp.2004.03.009.
URL <http://portal.acm.org/citation.cfm?id=1039334.1039340>
- [223] D. A. Knoll, V. A. Mousseau, L. Chacón, J. Reisner, Jacobian-Free Newton-Krylov methods for the accurate time integration of stiff wave systems, *J. Sci. Comput.* 25 (2005) 213–230. doi:10.1007/s10915-004-4640-8.
URL <http://portal.acm.org/citation.cfm?id=1104587.1104599>
- [224] D. R. Reynolds, R. Samtaney, C. S. Woodward, A fully implicit numerical method for single-fluid resistive magnetohydrodynamics, *J. Comput. Phys.* 219 (2006) 144–162. doi:10.1016/j.jcp.2006.03.022.
URL <http://portal.acm.org/citation.cfm?id=1238654.1238664>
- [225] A. Yeckel, L. Lun, J. J. Derby, An approximate block Newton method for coupled iterations of nonlinear solvers: Theory and conjugate heat transfer applications, *J. Comput. Phys.* 228 (2009) 8566–8588. doi:10.1016/j.jcp.2009.08.003.
URL <http://portal.acm.org/citation.cfm?id=1630172.1630409>
- [226] J. P. Whiteley, K. Gillow, S. J. Tavener, A. C. Walter, Error bounds on block Gauss-Seidel solutions of coupled multiphysics problems, *International Journal for Numerical Methods in Engineering* 88 (12) (2011) 1219–1237. doi:10.1002/nme.3217.
URL <http://dx.doi.org/10.1002/nme.3217>
- [227] L. B. Wigton, N. J. Yu, D. P. Young, GMRES acceleration of computational fluid dynamics codes, in: *AIAA 7th Computational Fluid Dynamics Conference*, 1985.
- [228] L. Chacón, D. A. Knoll, J. M. Finn, An implicit, nonlinear reduced resistive MHD solver, *J. Comput. Phys.* 178 (2002) 15–36.
- [229] H. C. Elman, V. E. Howle, J. N. Shadid, R. S. Tuminaro, A parallel block multi-level preconditioner for the 3D incompressible Navier-Stokes equations, *J. Comput. Phys.* 187 (2003) 504–523. doi:10.1016/S0021-9991(03)00121-9.
URL <http://portal.acm.org/citation.cfm?id=846060.846067>
- [230] K. Mosthaf, K. Baber, B. Flemisch, R. Helmig, A. Leijnse, I. Rybak, B. Wohlmuth, A coupling concept for two-phase compositional porous-medium and single-phase compositional free flow, *Water Resour. Res.* 47. doi:10.1029/2011WR010685.
- [231] P. Bochev, M. Shashkov, Constrained interpolation (remap) of divergence-free fields, *Comput. Methods Appl. Mech. Eng.* 194 (2-5) (2005) 511–530. doi:10.1016/j.cma.2004.05.018.
URL <http://linkinghub.elsevier.com/retrieve/pii/S0045782504003482>
- [232] X. Jiao, M. T. Heath, Common-refinement-based data transfer between nonmatching meshes in multiphysics simulations, *Internat. J. Numer. Methods Eng.* 61 (14) (2004) 2402–2427.
- [233] R. Jaiman, P. Geubelle, E. Loth, X. Jiao, Transient fluid-structure interaction with non-matching spatial and temporal discretization, *Comput. & Fluids* 50 (2011) 120–135. doi:10.1016/j.compfluid.2011.07.001.
- [234] J. Grandy, Conservative remapping and region overlays by intersecting arbitrary polyhedra, *J. Comput. Phys.* 148 (2) (1999) 433–466. doi:10.1006/jcph.1998.6125.
URL <http://linkinghub.elsevier.com/retrieve/pii/S0021999198961253>
- [235] R. Jaiman, P. Geubelle, E. Loth, X. Jiao, Combined interface boundary condition method for unsteady fluid-structure interaction, *Computer Methods in Applied Mechanics and Engineering* 200 (2011) 27–39. doi:10.1016/j.cma.2010.06.039.
- [236] X. Jiao, M. T. Heath, Overlaying surface meshes, part I: algorithms, *Int. J. Comput. Geometry Appl.* 14 (6) (2004) 379–402.
- [237] X. Jiao, M. T. Heath, Overlaying surface meshes, part II: topology preservation and feature matching, *Int. J. Comput. Geometry Appl.* 14 (6) (2004) 403–419.
- [238] P. Solin, J. Cerveny, L. Dubcova, D. Andrs, Monolithic discretization of linear thermoelasticity problems via adaptive multimesh FEM, *J. Comput. Appl. Math.* 234 (7) (2010) 2350 – 2357, *icetitle*;Fourth International Conference on Advanced Computational Methods in Engineering (ACOMEN 2008);*icetitle*. doi:10.1016/j.cam.2009.08.092.
URL <http://www.sciencedirect.com/science/article/pii/S0377042709005731>
- [239] P. Chidyagwai, B. Rivière, Numerical modelling of coupled surface and subsurface flow systems, *Adv. in Water Resour.* 33 (1) (2010) 92 – 105. doi:10.1016/j.advwatres.2009.10.012.
URL <http://www.sciencedirect.com/science/article/pii/S0309170809001742>
- [240] V. Girault, B. Rivière, DG approximation of coupled Navier-Stokes and Darcy equations by Beaver-Joseph-Saffman interface condition, *SIAM J. Numer. Anal.* 47 (2009) 2052–2089. doi:10.1137/070686081.
URL <http://portal.acm.org/citation.cfm?id=1654814.1654834>
- [241] P. Chidyagwai, B. Rivière, A two-grid method for coupled free flow with porous media flow, *Adv. Water Resour.* In Press, Corrected Proof (2011)–. doi:10.1016/j.advwatres.2011.04.010.
URL <http://www.sciencedirect.com/science/article/pii/S0309170811000741>
- [242] C. Bernardi, Y. Maday, A. T. Patera, A new nonconforming approach to domain decomposition: the mortar element method, in: H. Brezis, J. Lions (Eds.), *Nonlinear PDEs and Their Applications*, Longman, 1994.
- [243] C. Bernardi, Y. Maday, A. Patera, Domain decomposition by the mortar element method, in: H. K. et al. (Ed.), *Asymptotic and numerical methods for partial differential equations with critical parameters*, Reidel, Dordrecht, 1993, pp. 269–286.
- [244] B. Wohlmuth, Variationally consistent discretization schemes and numerical algorithms for contact problems, *Acta Numerica* (2011) 569–734.
- [245] B. Flemisch, M. Kaltenbacher, S. Triebenbacher, B. I. Wohlmuth, The equivalence of standard and mixed finite element methods in applications to elasto-acoustic interaction, *SIAM J. Sci. Comput.* 32 (2010) 1980–2006. doi:10.1137/090758507.
- [246] L. F. Shampine, C. W. Gear, A user’s view of solving stiff ordinary differential equations, *SIAM Rev.* 21 (1) (1979) 1–17.
- [247] D. J. Higham, L. N. Trefethen, Stiffness of ODEs, *BIT* 33 (2) (1993) 285–303.
- [248] X. Zhong, Additive semi-implicit Runge-Kutta methods for computing high speed nonequilibrium reactive flows, *J. Comput. Phys.* 128 (1996) 19–31.
- [249] F. X. Giraldo, Semi-implicit time-integrators for a scalable spectral element atmospheric model, *Quarterly Journal of the Royal Meteorological Society* 131 (610) (2005) 2431–2454.

REFERENCES

- [250] U. M. Ascher, S. J. Ruuth, R. J. Spiteri, Implicit-explicit Runge-Kutta methods for time-dependent partial differential equations, *Appl. Numer. Math.* 25 (1997) 151–167. doi:10.1016/S0168-9274(97)00056-1.
URL <http://portal.acm.org/citation.cfm?id=266992.267001>
- [251] J. G. Verwer, B. P. Sommeijer, An implicit-explicit Runge-Kutta-Chebyshev scheme for diffusion-reaction equations, *SIAM J. Sci. Comput.* 25 (5) (2004) 1824–1835.
- [252] L. F. Shampine, B. P. Sommeijer, J. G. Verwer, IRKC: an IMEX solver for stiff diffusion-reaction PDEs, *J. Comp. Appl. Math.* 196 (2) (2006) 485–497.
- [253] W. Hundsdorfer, S. J. Ruuth, IMEX extensions of linear multistep methods with general monotonicity and boundedness properties, *J. Comput. Phys.* 225 (2007) 2016–2042. doi:10.1016/j.jcp.2007.03.003.
URL <http://portal.acm.org/citation.cfm?id=1280287.1280444>
- [254] E. M. Constantinescu, A. Sandu, Optimal explicit strong-stability-preserving general linear methods, *SIAM J. Sci. Comput.* 32 (2010) 3130–3150. doi:10.1137/090766206.
- [255] J. M. Connors, A. Miloua, Partitioned time discretization for parallel solution of coupled ODE systems, *BIT* 51 (2) (2011) 253–273. doi:10.1007/s10543-010-0295-z.
- [256] F. X. Giraldo, J. B. Perot, P. F. Fischer, A spectral element semi-lagrangian (SESL) method for the spherical shallow water equations, *J. Comput. Phys.* 190 (2) (2003) 623–650.
- [257] D. Estep, V. Ginting, D. Ropp, J. N. Shadid, S. Tavener, An a posteriori-a priori analysis of multiscale operator splitting, *SIAM J. Numer. Anal.* 46 (2008) 1116–1146. doi:10.1137/07068237X.
URL <http://portal.acm.org/citation.cfm?id=1362718.1362720>
- [258] F. X. Giraldo, M. Restelli, High-order semi-implicit time-integrators for a triangular discontinuous Galerkin oceanic shallow water model, *Internat. J. Numer. Methods Fluids* (in press) (2008) 24.
- [259] P. N. Blossley, D. R. Durran, Implicit-explicit multistep methods for nonhydrostatic atmospheric models, In Preparation.
- [260] F. X. Giraldo, M. Restelli, M. Lauter, Semi-implicit formulations of the Navier-Stokes equations: Application to nonhydrostatic atmospheric modeling, *SIAM J. Sci. Comput.* 32 (6) (2010) 3394–3425.
- [261] S. J. Ruuth, Implicit-explicit methods for reaction-diffusion, *J. Math. Biol.* 34 (1995) 148–176.
- [262] E. M. Constantinescu, A. Sandu, Multirate timestepping methods for hyperbolic conservation laws, *J. Sci. Comput.* 33 (3) (2007) 239–278. doi:10.1007/s10915-007-9151-y.
- [263] A. Sandu, E. Constantinescu, Multirate explicit Adams methods for time integration of conservation laws, *J. Sci. Comput.* 38 (2) (2009) 229–249. doi:10.1007/s10915-008-9235-3.
- [264] E. M. Constantinescu, A. Sandu, G. Carmichael, Modeling atmospheric chemistry and transport with dynamic adaptive resolution, *Comput. Geosci.* 12 (2) (2008) 133–151. doi:10.1007/s10596-007-9065-7.
- [265] D. A. Knoll, L. Chacón, L. G. Margolin, V. A. Mousseau, On balanced approximations for time integration of multiple time scale systems, *J. Comput. Phys.* 185 (2003) 583–611. doi:10.1016/S0021-9991(03)00008-1.
URL <http://portal.acm.org/citation.cfm?id=775807.775820>
- [266] C. A. Kennedy, M. H. Carpenter, Additive Runge-Kutta schemes for convection-diffusion-reaction equations, *Appl. Numer. Math.* 44 (2003) 139–181. doi:10.1016/S0168-9274(02)00138-1.
URL <http://portal.acm.org/citation.cfm?id=639155.639164>
- [267] S. Y. Kadioglu, D. A. Knoll, R. B. Lowrie, R. M. Rauenzahn, A second order self-consistent IMEX method for radiation hydrodynamics, *J. Comput. Phys.* 229 (2010) 8313–8332. doi:10.1016/j.jcp.2010.07.019.
- [268] C. W. Gear, D. R. Wells, Multirate linear multistep methods, *BIT* 24 (1984) 484–502.
- [269] A. Bartel, M. Günther, A multirate W-method for electrical networks in state-space formulation, *J. Comput. Appl. Math.* 147 (2) (2002) 411–425. doi:10.1016/S0377-0427(02)00476-4.
- [270] A. Logg, Multi-adaptive Galerkin methods for ODEs I, *SIAM J. Sci. Comput.* 24 (6) (2003) 1879–1902.
- [271] A. Logg, Multi-adaptive Galerkin methods for ODEs II: Implementation and applications, *SIAM J. Sci. Comput.* 25 (4) (2003) 1119–1141.
- [272] A. Logg, Multi-adaptive Galerkin methods for ODEs III: Existence and stability, Tech. Rep. 2004-04, Chalmers Finite Element Center Preprint Series (2004).
- [273] A. Logg, Multi-adaptive Galerkin methods for ODEs III: A priori estimates, *SIAM J. Numer. Anal.* 43 (6) (2006) 2624–2646.
- [274] V. Savcenko, W. Hundsdorfer, J. Verwer, A multirate time stepping strategy for stiff ODEs, *BIT* 47 (1) (2007) 137–155. doi:10.1007/s10543-006-0095-7.
- [275] M. Crouzeix, Une methode multipas implicite-explicite pour l’approximation des equations d’evolution parabolique, *Numer. Math.* 35 (1980) 257–276.
- [276] F. Filbet, S. Jin, A class of asymptotic preserving schemes for kinetic equations and related problems with stiff sources, *J. Comput. Phys.*
- [277] R. LeVeque, *Finite volume methods for hyperbolic problems*, Cambridge University Press, 2002.
- [278] U. M. Ascher, S. J. Ruuth, B. T. R. Wetton, Implicit-explicit methods for time-dependent partial differential equations, *SIAM J. Numer. Anal.* 32 (3) (1995) 797–823.
- [279] W. B. Gragg, H. J. Stetter, Generalized multistep predictor-corrector methods, *J. ACM* 11 (2) (1964) 188–209. doi:10.1145/321217.321223.
- [280] J. Frank, W. Hundsdorfer, J. G. Verwer, On the stability of implicit-explicit linear multistep methods, *Appl. Numer. Math.* 25 (2-3) (1997) 193–205. doi:10.1016/S0168-9274(97)00059-7.
- [281] D. L. Ropp, J. N. Shadid, C. C. Ober, Studies of the accuracy of time integration methods for reaction-diffusion equations, *J. Comput. Phys.* 194 (2004) 544–574. doi:10.1016/j.jcp.2003.08.033.
URL <http://portal.acm.org/citation.cfm?id=1008428.1008435>
- [282] D. L. Ropp, J. N. Shadid, Stability of operator splitting methods for systems with indefinite operators: Advection-diffusion-reaction systems, *J. Comput. Phys.* 228 (2009) 3508–3516. doi:10.1016/j.jcp.2009.02.001.
URL <http://portal.acm.org/citation.cfm?id=1518324.1518454>

REFERENCES

- [283] E. Constantinescu, A. Sandu, Extrapolated implicit-explicit time stepping, *SIAM Journal on Scientific Computing* 31 (6) (2010) 4452–4477. doi:10.1137/080732833.
- [284] J. Connors, J. Howell, W. Layton, Partitioned timestepping for a parabolic two domain problem, *SIAM Jour. Num. Analysis* 47 (5) (2009) 3526–3549.
- [285] M. S. Day, J. B. Bell, Numerical simulation of laminar reacting flows with complex chemistry, *Combustion Theory and Modelling* 4 (4) (2000) 535–556. arXiv:<http://www.tandfonline.com/doi/pdf/10.1088/1364-7830/4/4/309>, doi:10.1088/1364-7830/4/4/309. URL <http://www.tandfonline.com/doi/abs/10.1088/1364-7830/4/4/309>
- [286] A. S. Almgren, J. B. Bell, P. Colella, L. H. Howell, M. L. Welcome, A conservative adaptive projection method for the variable density incompressible Navier-Stokes equations, *J. Comput. Phys.* 142 (1998) 1–46.
- [287] R. E. Ewing, R. D. Lazarov, P. S. Vassilevski, Local refinement techniques for elliptic problems on cell-centered grids. I: Error analysis, *Math. Comp.* 56 (1991) 437–461.
- [288] D. I. Choi, J. D. Brown, B. Imbiriba, J. Centrella, P. MacNeice, Interface conditions for wave propagation through mesh refinement boundaries, *J. Comput. Phys.* 193 (2) (2004) 398–425. doi:10.1016/j.jcp.2003.07.036.
- [289] M. J. Berger, P. Colella, Local adaptive mesh refinement for shock hydrodynamics, *J. Comput. Phys.* 82 (1989) 64–84.
- [290] R. B. Pember, L. H. Howell, J. B. Bell, P. Colella, W. Y. Crutchfield, W. A. Fiveland, J. P. Jessee, An adaptive projection method for unsteady, low-mach number combustion, *Combustion Science and Technology* 140 (1-6) (1998) 123–168. doi:10.1080/00102209808915770.
- [291] R. Frank, W. Ueberhuber, Iterated defect correction for the efficient solution of stiff systems of ordinary differential equations, *BIT* 17 (2) (1977) 146–159.
- [292] A. Dutt, L. Greengard, V. Rokhlin, Spectral deferred correction methods for ordinary differential equations, *BIT* 40 (2000) 241–266. doi:10.1023/A:1022338906936.
- [293] M. L. Minion, Semi-implicit spectral deferred correction methods for ordinary differential equations, *Comm. Math. Sci.* 1 (3) (2003) 471–500.
- [294] M. L. Minion, Semi-implicit projection methods for incompressible flow based on spectral deferred corrections, *Appl. Numer. Math.* 48 (3-4) (2004) 369–387, Workshop on Innovative Time Integrators for PDEs, Amsterdam, Netherlands, Nov. 25-27, 2002. doi:10.1016/j.apnum.2003.11.005.
- [295] J. Fish (Ed.), *Multiscale Methods: Bridging the Scales in Science and Engineering*, Oxford University Press, 2010.
- [296] X. Luo, T. Stylianopoulos, V. H. Barocas, M. S. Shephard, Multiscale computation for soft tissues with complex geometries, *Engineering with Computers* 25 (1) (2009) 87–96.
- [297] A. Bensoussan, J. Lions, G. Papanicolaou, *Asymptotic Analysis for Periodic Structures*, North Holland, 1978.
- [298] E. Sanchez-Palencia, *Non-homogeneous Media and Vibration Theory*, Vol. 127 of *Lecture Notes in Physics*, Springer-Verlag, 1980.
- [299] A. Li, R. Li, J. Fish, Generalized mathematical homogenization: From theory to practice, *Comput. Methods Appl. Mech. Eng.* 197 (41–42) (2008) 3225–3248.
- [300] J. Knap, M. Ortiz, An analysis of the quasicontinuum method, *J. Mech. Phys. Solids* 49 (2001) 1899–1923.
- [301] R. E. Miller, E. B. Tadmor, The Quasicontinuum Method: Overview, applications and current directions, *Journal of Computer-Aided Materials Design* 9 (3) (2002) 203–239.
- [302] T. Belytschko, S. P. Xiao, Coupling methods for continuum model with molecular model, *Int. J. Mult. Comp. Eng.* 1 (2003) 115–126.
- [303] J. Fish, M. A. Nuggehally, M. S. Shephard, C. R. Picu, S. Badia, M. L. Parks, M. Gunzburger, Concurrent AtC coupling based on a blend of the continuum stress and the atomistic force, *Computer Methods in Applied Mechanics and Engineering* 196 (45-48) (2007) 4548 – 4560. doi:10.1016/j.cma.2007.05.020. URL <http://www.sciencedirect.com/science/article/pii/S0045782507002277>
- [304] G. Wagner, W. Liu, Coupling of atomic and continuum simulations using a bridging scale decomposition, *J. Comput. Phys.* 190 (2003) 249–274.
- [305] S. Xiao, T. Belytschko, A bridging domain method for coupling continua with molecular dynamics, *Comput. Methods Appl. Mech. Eng.* 193 (2004) 1645–1669.
- [306] P. T. Bauman, J. T. Oden, S. Prudhomme, Adaptive multiscale modeling of polymeric materials: Arlequin coupling and Goals algorithm, *Comput. Methods Appl. Mech. Eng.* 198 (5–8) (2008) 799–818.
- [307] M. Nuggehally, C. Picu, M. Shephard, J. Fish, Adaptive model selection procedure for concurrent multiscale problems, *Int. J. Mult. Comp. Eng.* 5 (5) (2007) 369–386.
- [308] J. T. Oden, P. T. Bauman, S. Prudhomme, Multiscale modeling of physical phenomena: adaptive control of models, *SIAM J. Sci. Comput.* 28 (6) (2006) 2359–2389.
- [309] W. E, B. Engquist, X. Li, W. Ren, E. Vanden-Eijnden, Heterogeneous multiscale methods: A review, *Communications in Computational Physics* 2 (3) (2007) 367–450. URL http://www.global-sci.com/openaccess/v2_367.pdf
- [310] I. Kevrekidis, C. Gear, J. Hyman, P. Kevrekidis, O. Runborg, C. Theodoropoulos, Equation-free coarse-grained multiscale computation: enabling microscopic simulators to perform system-level tasks, *Commun. Math. Sci.* 1 (4) (2003) 715–762.
- [311] M. Ainsworth, J. Oden, A posteriori error estimation in finite element analysis, *Computer Methods in Applied Mechanics and Engineering* 142 (1-2) (1997) 1–88. doi:10.1016/S0045-7825(96)01107-3.
- [312] Z. Cai, S. Zhang, Recovery-based error estimator for interface problems: conforming linear elements, *SIAM J. Numer. Anal.* 47 (3) (2009) 2132–2156. doi:10.1137/080717407.
- [313] Z. Cai, S. Zhang, Recovery-based error estimators for interface problems: mixed and nonconforming finite elements, *SIAM J. Numer. Anal.* 48 (1) (2010) 30–52. doi:10.1137/080722631.
- [314] Z. Cai, S. Zhang, Flux recovery and a posteriori error estimators: conforming elements for scalar elliptic equations, *SIAM J. Numer. Anal.* 48 (2) (2010) 578–602. doi:10.1137/080742993.

REFERENCES

- [315] D. J. Estep, M. G. Larson, R. D. Williams, Estimating the error of numerical solutions of systems of reaction-diffusion equations, Vol. 146 of *Memoirs A.M.S.*, American Mathematical Society, 2000.
- [316] R. Becker, R. Rannacher, An optimal control approach to a posteriori error estimation in finite element methods, *Acta Numer.* (2001) 1–102.
- [317] J. T. Oden, S. Prudhomme, Estimation of modeling error in computational mechanics, *J. Comput. Phys.* 182 (2002) 496–515.
- [318] D. Estep, S. Tavener, T. Wildey, A posteriori analysis and improved accuracy for an operator decomposition solution of a conjugate heat transfer problem, *SIAM J. Numer. Anal.* 46 (2008) 2068–2089. doi:10.1137/060678737.
URL <http://portal.acm.org/citation.cfm?id=1404777.1404796>
- [319] D. Estep, S. Tavener, T. Wildey, A posteriori error analysis for a transient conjugate heat transfer problem, *Finite Elem. Anal. Des.* 45 (2009) 263–271. doi:10.1016/j.finel.2008.10.011.
URL <http://portal.acm.org/citation.cfm?id=1513008.1513409>
- [320] V. Carey, D. Estep, S. Tavener, A posteriori analysis and adaptive error control for multiscale operator decomposition solution of elliptic systems I: triangular systems, *SIAM J. Numer. Anal.* 47 (1) (2009) 740–761. doi:10.1137/070689917.
- [321] R. Ghanem, P. Spanos, *Stochastic Finite Elements: A Spectral Approach*, Springer-Verlag, 1991.
- [322] D. Xiu, *Numerical Methods for Stochastic Computations: A Spectral Method Approach*, Princeton University Press, 2010.
- [323] S. Jobs, *Apple World Wide Developers Conference, Closing Keynote Q & A* (1997).
URL <http://www.youtube.com/watch?v=3LEXae1j6EY#t=41m26s>
- [324] W. Elwasif, D. E. Bernholdt, A. G. Shet, S. S. Foley, R. Bramley, D. B. Batchelor, L. A. Berry, The design and implementation of the SWIM Integrated Plasma Simulator, in: *Parallel, Distributed and Network-Based Processing (PDP), 2010 18th Euromicro International Conference on*, 2010, pp. 419–427. doi:10.1109/PDP.2010.63.
- [325] K. K. Chand, L. F. Diachin, C. O.-G. X. Li, E. S. Seol, M. S. Shephard, T. Tautges, H. Trease, Toward interoperable mesh, geometry and field components for PDE simulation development, *Engineering with Computers* 24 (2) (2008) 165–182.
- [326] K. Devine, L. Diachin, J. Kraftcheck, K. E. Jansen, V. Leung, X. Luo, M. Miller, C. Ollivier-Gooch, A. Ovcharenko, O. Sahni, M. Shephard, T. Tautges, T. Xie, M. Zhou, Interoperable mesh components for large-scale, distributed-memory simulations, *J. Physics: Conf. Series* 180 (2009) 012011.
- [327] C. Ollivier-Gooch, L. F. Diachin, M. S. Shephard, T. Tautges, J. Kraftcheck, V. Leung, X.-J. Luo, M. Miller, An interoperable, data-structure-neutral component for mesh query and manipulation, *Transactions on Mathematical Software* 37 (3) (2010) 29:1–29:28.
- [328] Multi-Physics Coupling Code Interface, <http://www.mppcci.de>.
- [329] J. W. Larson, Ten organising principles for coupling in multiphysics and multiscale models, *ANZIAM* 48 (2009) C1090–C1111.
- [330] J. W. Larson, Graphical notation for diagramming coupled systems, in: *Proceedings of the Ninth International Conference on Computational Science (ICCS2009)*, Lecture Notes in Computer Science, Vol. 5544, Springer, 2009, pp. 745–754.
- [331] M. Glass et al., *STK: Sierra Toolkit Mesh*, <http://trilinos.sandia.gov/packages/stk> (2011).
- [332] T. Tautges, J. Kraftcheck, B. Smith, H.-H. Kim, MOAB: Mesh-oriented database, version 4.0, <http://trac.mcs.anl.gov/projects/ITAPS/wiki/MOAB> (2011).
- [333] B. S. Kirk, J. W. Peterson, R. H. Stogner, G. F. Carey, *libMesh: A C++ Library for Parallel Adaptive Mesh Refinement/Coarsening Simulations*, *Engineering with Computers* 22 (3–4) (2006) 237–254. doi:10.1007/s00366-006-0049-3.
URL <http://libmesh.sourceforge.net>
- [334] R. W. Ford, G. D. Riley, M. K. Bane, C. W. Armstrong, T. L. Freeman, GCF: a general coupling framework, *Concurrency and Computation: Practice and Experience* 18. doi:10.1002/cpe.910.
- [335] R. Sahu, M. J. Panthaki, W. H. Gerstle, An object-oriented framework for multidisciplinary, multi-physics computational mechanics, *Engineering with Computers* 15 (1) (1999) 105–125. doi:10.1007/s003660050008.
- [336] L. C. McInnes, B. Allan, R. Armstrong, S. Benson, D. Bernholdt, T. Dahlgren, L. Diachin, M. Krishnan, J. Kohl, J. Larson, S. Lefantzi, J. Nieplocha, B. Norris, S. Parker, J. Ray, S. Zhou, *Numerical Solution of Partial Differential Equations on Parallel Computers*, no. 51 in *Lecture Notes in Computational Science and Engineering*, Springer-Verlag, 2006, Ch. Parallel PDE-Based Simulations Using the Common Component Architecture, pp. 327–384.
- [337] T. Bulatewicz, J. Cuny, A domain-specific language for model coupling, in: *Proceedings of the 2006 Winter Simulation Conference*, 2006.
- [338] Z. DeVito, N. Joubert, F. Palacios, S. Oakley, M. Medina, M. Barrientos, E. Elsen, F. Ham, A. Aiken, K. Duraisamy, E. Darve, J. A. P. Hanrahan, Lizst: A domain specific language for building portable mesh-based PDE solvers, in: *SC'11: Proceedings of the 2011 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, 2011.
- [339] T. Goodale, G. Allen, G. Lanfermann, J. Massó, T. Radke, E. Seidel, J. Shalf, The Cactus framework and toolkit: Design and applications, in: *Vector and Parallel Processing – VECPAR'2002, 5th International Conference, Lecture Notes in Computer Science*, Springer, Berlin, 2003.
URL <http://edoc.mpg.de/3341>
- [340] F. Liu, M. Sosonkina, A multilevel parallelism support for multi-physics coupling, in: *Procedia Computer Science*, Vol. 4, 2011, pp. 261–270.
- [341] J. Michopoulos, P. Tsompanopoulou, E. Houstis, C. Farhat, M. Lesoinne, F. Lechenault, DDEMA: A data driven environment for multiphysics applications, in: *Proceedings of International Conference of Computational Science - ICCS03*, Sloot, P.M.A., et al. (Eds.) Melbourne Australia, June 2-4, LNCS 2660, Part IV, Springer-Verlag, Heidelberg, 2003, pp. 309–318.
- [342] A. Dubey, K. Antypas, M. K. Ganapathy, L. B. Reid, K. Riley, D. Sheeler, A. Siegel, K. Weide, Extensible component-based architecture for FLASH, a massively parallel, multiphysics simulation code, *Parallel Computing* 35 (10-11) (2009) 512 – 522. doi:10.1016/j.parco.2009.08.001.
- [343] A. Sussman, Building complex coupled physical simulations on the grid with InterComm, *Engineering with Computers* 22 (2006) 311–323. doi:10.1007/s00366-006-0037-7.
- [344] J. Hegewald, M. Krafczyk, J. Tölke, A. Hoekstra, B. Chopard, An agent-based coupling platform for complex automata, in: *Proceedings of the 8th International Conference on Computational Science, Part II, ICCS '08*, Springer-Verlag, Berlin, Heidelberg, 2008, pp. 227–233. doi:10.1007/978-3-540-69387-1_25.

REFERENCES

- [345] V. Mahadevan, J. C. Ragusa, V. Mousseau, Verification of multiphysics software: Space and time convergence studies for nonlinearly coupled applications, in: *Proceedings of the International Conference on Mathematics, Computational Methods and Reactor Physics*, Saratoga Springs, New York, May 3-7, 2009, American Nuclear Society, 2009.
- [346] J. Larson, R. Jacob, E. Ong, The Model Coupling Toolkit: A new Fortran90 toolkit for building multiphysics parallel coupled models, *Int. J. High Perf. Comp. App.* 19 (3) (2005) 277–292.
- [347] S. P. Zwart, S. McMillan, B. O. Nualláin, D. Heggie, J. Lombardi, P. Hut, S. Banerjee, H. Belkus, T. Fragos, J. Fregeau, M. Fuji, E. Gaburov, E. Glebbeek, D. Groen, S. Harfst, R. Izzard, P. T. M. Jurić and, S. Justha and, O. Y. J. van Beve and, M. Zem, A multiphysics and multiscale software environment for modeling astrophysical systems, in: *Proceedings of the 8th International Conference on Computational Science, Part II, ICCS '08*, Springer-Verlag, Berlin, Heidelberg, 2008, pp. 206–216. doi:10.1007/978-3-540-69387-1_23.
- [348] J. M. Grandy, Simulations on multiple meshes: Algorithms, Tech. Rep. Report U-COPJ-2004-0545, Lawrence Livermore National Laboratory (2004).
- [349] W. D. Henshaw, Overture: An object-oriented framework for overlapping grid applications, in: *2002 AIAA conference on Applied Aerodynamics*, 2002.
- [350] B. Gatzhammer, M. Mehl, T. Neckel, A coupling environment for partitioned multiphysics simulations applied to fluid-structure interaction scenarios, in: *Procedia Computer Science*, Vol. 1, 2010, pp. 681–689.
- [351] J. R. Stewart, H. C. Edwards, A framework approach for developing parallel adaptive multiphysics applications, *Finite Elements in Analysis and Design* 40 (2004) 1599–1617. doi:10.1016/j.finel.2003.10.006.
- [352] F. Zhang, C. Docan, M. Parashar, S. Klasky, Enabling multi-physics coupled simulations within the PGAS programming framework, in: *Cluster, Cloud and Grid Computing (CCGrid)*, 2011 11th IEEE/ACM International Symposium on, 2011, pp. 84–93. doi:10.1109/CCGrid.2011.73.
- [353] J. Bell et al., *BoxLib User's Guide*, Tech. rep., CCSE, Lawrence Berkeley National Laboratory, <https://ccse.lbl.gov/BoxLib/BoxLibUsersGuide.pdf> (2011).
- [354] M. S. Day, J. B. Bell, Numerical simulation of laminar reacting flows with complex chemistry, *Combust. Theory Modelling* 4 (4) (2000) 535–556.
- [355] A. S. Almgren, V. E. Beckner, J. B. Bell, M. S. Day, L. H. Howell, C. C. Joggerst, M. J. Lijewski, A. Nonaka, M. Singer, M. Zingale, CASTRO: A New Compressible Astrophysical Solver. I. Hydrodynamics and Self-gravity, *Astrophysical Journal* 715 (2010) 1221–1238.
- [356] W. Zhang, L. Howell, A. Almgren, A. Burrows, J. Bell, CASTRO: A New Compressible Astrophysical Solver. II. Gray Radiation Hydrodynamics, *Astrophysical Journal Supplement* 20 (2011) 1962.
- [357] A. S. Almgren, J. B. Bell, C. A. Rendleman, M. Zingale, Low mach number modeling of type ia supernovae. i. hydrodynamics, *Astrophysical Journal* 637 (2006) 922–936.
- [358] A. S. Almgren, J. B. Bell, C. A. Rendleman, M. Zingale, Low mach number modeling of type ia supernovae. ii. energy evolution, *Astrophysical Journal* 649 (2006) 927–938.
- [359] A. S. Almgren, J. B. Bell, A. Nonaka, M. Zingale, Low Mach Number Modeling of Type Ia Supernovae. III. Reactions, *Astrophysical Journal* 684 (2008) 449–470. doi:10.1086/590321.
- [360] M. Zingale, A. S. Almgren, J. B. Bell, A. Nonaka, S. E. Woosley, Low Mach Number Modeling of Type Ia Supernovae. IV. White Dwarf Convection, *Astrophysical Journal* 704 (2009) 196–210.
- [361] A. Nonaka, A. S. Almgren, J. B. Bell, M. J. Lijewski, C. M. Malone, M. Zingale, MAESTRO: an adaptive low mach number hydrodynamics algorithm for stellar flows, *Astrophysical Journal Supplement* 188 (2010) 358–383.
- [362] G. S. H. Pau, A. S. Almgren, J. B. Bell, M. J. Lijewski, A parallel second-order adaptive mesh algorithm for incompressible flow in porous media, *Phil. Trans. R. Soc. A* 367 (2009) 4633–4654.
- [363] A. Almgren, J. Bell, M. Lijewski, Z. Lukić, NYX: Scalable adaptive mesh refinement of hydrodynamics and dark matter in an expanding universe (in preparation).
- [364] P. Colella, D. T. Graves, D. Modiano, D. B. Serafini, B. V. Straalen, Chombo software package for AMR applications, Tech. rep., Lawrence Berkeley National Laboratory, <http://seesar.lbl.gov/anag/chombo/> (2000).
- [365] B. Van Straalen, P. Colella, D. Graves, N. Keen, Petascale block-structured AMR applications without distributed meta-data, in: E. Jeannot, R. Namyst, J. Roman (Eds.), *Euro-Par 2011 Parallel Processing*, Vol. 6853 of *Lecture Notes in Computer Science*, Springer Berlin / Heidelberg, 2011, pp. 377–386. doi:10.1007/978-3-642-23397-5_37.
- [366] E. Ateljevich, P. Colella, D. T. Graves, T. J. Ligocki, J. Percelay, P. O. Schwartz, Q. Shu, CFD modeling in the San Francisco Bay and Delta, in: *Proceedings of the Fourth SIAM Conference on Mathematics for Industry (MI09)*, 2009, 2010, pp. 99–107.
- [367] S. L. Cornford, D. F. Martin, D. T. Graves, D. F. Ranken, A. M. LeBrocq, R. M. Gladstone, A. J. Payne, E. G. Ng, W. H. Lipscomb, Adaptive mesh, finite-volume modeling of marine ice sheets, *Journal of Computational Physics* submitted.
- [368] M. R. Dorr, R. H. Cohen, P. Colella, M. A. Dorf, J. A. F. Hittinger, D. F. Martin, Numerical simulation of phase space advection in gyrokinetic models of fusion plasmas, in: *Proceedings of the 2010 Scientific Discovery through Advanced Computing (SciDAC) Conference*, Chattanooga, Tennessee, July 11-15, 2010, Oak Ridge National Laboratory, 2010, pp. 42–52. URL http://computing.ornl.gov/workshops/scidac2010/2010_SciDAC_Proceedings.pdf
- [369] Illinois rocstar, Website (2011). URL www.illinoisrocstar.com
- [370] M. Heath (PI), Center for Simulation of Advanced Rockets, Website. URL <http://www.csar.illinois.edu/>
- [371] Escript/Finley, Website (2011). URL <https://launchpad.net/escript-finley/>
- [372] L. Gross, L. Bourgooin, A. J. Hale, H.-B. Muhlhaus, Interface modeling in incompressible media using level sets in escript, *Physics of the Earth and Planetary Interiors* 163 (2007) 23–34. doi:10.1016/j.pepi.2007.04.004.
- [373] R. Schmidt, K. Belcourt, K. Clarno, R. Hooper, L. Humphries, A. Lorber, R. Pryor, W. Spatz, Foundational development of an advanced nuclear reactor integrated safety code, Sandia Report SAND2010-0878, Sandia National Laboratories (February 2010).

REFERENCES

- [374] R. Pawlowski, R. Bartlett, N. Belcourt, R. Hooper, R. Schmidt, A theory manual for multi-physics code coupling in LIME, Sandia Report SAND2011-2195, Sandia National Laboratories (March 2011). doi:10.2172/1011710.
URL http://www.cs.sandia.gov/~rabart1/LIME_TheoryManual.pdf
- [375] D. Kothe, Casl, Website (September 2011).
URL <http://www.casl.gov>
- [376] T. J. Tautges, A. Caceres, Scalable parallel solution coupling for multiphysics reactor simulation, Journal of Physics: Conference Series 180. URL http://www.iop.org/EJ/article/1742-6596/180/1/012017/jpconf9_180_012017.pdf
- [377] F. Brown, Mcnp website, Website (2004).
URL <http://mcnp-green.lanl.gov/index.html>
- [378] M. Sawan, P. Wilson, T. Tautges, L. El-Guebaly, D. Henderson, T. Bohm, E. Marriott, B. Kiedrowski, B. Smith, A. Ibrahim, et al., Application of cad-neutronics coupling to geometrically complex fusion systems, in: Fusion Engineering, 2009. SOFE 2009. 23rd IEEE/NPSS Symposium on, IEEE, 2009, pp. 1–6.
- [379] T. Evans, A. Stafford, R. Slaybaugh, K. Clarno, Denovo: A new three-dimensional parallel discrete ordinates code in SCALE, Nuclear technology 171 (2) (2010) 171–200.
- [380] R. Biezuner, J. Brown, G. Ercole, E. Martins, Computing the first eigenpair of the p -laplacian via inverse iteration of sublinear supersolutions, Journal of Scientific Computing (2011) 1–22doi:10.1007/s10915-011-9540-0.
- [381] J. Brown, Efficient nonlinear solvers for nodal high-order finite elements in 3D, Journal of Scientific Computing 45 (2010) 48–63. doi:10.1007/s10915-010-9396-8.
- [382] D. Gaston, C. Newman, G. Hansen, MOOSE: a parallel computational framework for coupled systems of nonlinear equations, in: American Nuclear Society 2009 International Conference on Advances in Mathematics, Computational Methods, and Reactor Physics, Saratoga Springs, NY, 2009.
- [383] B. S. Kirk, J. W. Peterson, R. H. Stogner, G. F. Carey, libMesh: a C++ library for parallel adaptive mesh refinement/coarsening simulations, Eng Comput-Germany 22 (3-4) (2006) 237–254. doi:10.1007/s00366-006-0049-3.
- [384] G. Hansen, R. Martineau, C. Newman, D. Gaston, Framework for simulation of pellet cladding thermal interaction (PCTI) for fuel performance calculations, in: American Nuclear Society 2009 International Conference on Advances in Mathematics, Computational Methods, and Reactor Physics, Saratoga Springs, NY, 2009.
- [385] D. Knoll, R. Park, K. Smith, Application of the Jacobian-free Newton-Krylov method in computational reactor physics, in: American Nuclear Society 2009 International Conference on Advances in Mathematics, Computational Methods, and Reactor Physics, Saratoga Springs, NY, 2009.
- [386] D. Knoll, H. Park, C. Newman, Acceleration of k -Eigenvalue / Criticality calculations using the Jacobian-free Newton-Krylov method, Nuclear Science and Engineering Submitted.
- [387] H. Park, D. Gaston, S. Kadioglu, D. Knoll, R. Martineau, W. Taitano, D. Lebrun-Grandie, Tightly coupled multiphysics simulations for pebble bed reactors, in: American Nuclear Society 2009 International Conference on Advances in Mathematics, Computational Methods, and Reactor Physics, Saratoga Springs, NY, 2009.
- [388] H. Park, D. A. Knoll, D. R. Gaston, R. C. Martineau, Tightly coupled multiphysics algorithms for pebble bed reactors, Nuclear Science and Engineering 166 (2) (2010) 118–133.
- [389] L. Guo, H. Huang, D. Gaston, G. Redden, Modeling of calcite precipitation driven by bacteria-facilitated urea hydrolysis in a flow column using a fully coupled, fully implicit parallel reactive transport simulator, in: Eos Transactions American Geophysical Union, 90(52), Fall Meeting Supplement, AGU 90(52), San Francisco, CA, 2010.
- [390] D. Gaston, L. Guo, G. Hansen, H. Huang, R. Johnson, D. Knoll, C. Newman, H. Park, R. Podgorney, M. Tonks, R. Williamson, Parallel algorithms and software for nuclear, energy, and environmental applications. Part II: Multiphysics software, Communications in Computational Physics Accepted.
- [391] S. Buis, A. Piacentini, D. Déclat, PALM: A computational framework for assembling high performance computing applications, Concurrency Computat.: Pract. Exper. 18 (2) (2006) 247–262.
- [392] A. Piacentini, T. Morel, A. Thévenin, F. Duchaine, Open-PALM: an open source dynamic parallel coupler, in: Coupled Problems 2011, June 20-22, 2011, Kos Island, Greece, 2011.
- [393] OpenPALM, http://www.cerfacs.fr/globc/PALM_WEB/.
- [394] H. Elman, V. E. Howle, J. Shadid, R. Shuttleworth, R. Tuminaro, A taxonomy and comparison of parallel block multi-level preconditioners for the incompressible navier-stokes equations, J. Comput. Phys. 227 (2008) 1790–1808. doi:10.1016/j.jcp.2007.09.026.
- [395] B. Aagaard, S. Kientz, M. G. Knepley, S. Somala, L. Strand, C. Williams, Pylith user manual version 1.6.0.
URL http://geodynamics.org/cig/software/pylith/pylith_manual-1.6.0.pdf
- [396] R. F. Katz, M. G. Knepley, B. Smith, M. Spiegelman, E. Coon, Numerical simulation of geodynamic processes with the Portable Extensible Toolkit for Scientific Computation, Phys. Earth Planet. In. 163 (2007) 52–68.
- [397] D. A. May, L. Moresi, Preconditioned iterative methods for stokes flow problems arising in computational geodynamics, Physics of the Earth and Planetary Interiors 171 (1-4) (2008) 33 – 47, recent Advances in Computational Geodynamics: Theory, Numerics and Applications. doi:10.1016/j.pepi.2008.07.036.
- [398] R. F. Katz, M. Spiegelman, B. Holtzman, The dynamics of melt and shear localization in partially molten aggregates, Nature 442 (2006) 676–679.
URL <http://www.nature.com/nature/journal/v442/n7103/abs/nature05039.html>
- [399] R. Katz, M. Worster, The stability of ice-sheet grounding lines, Proc. Roy. Soc. A 466 (2010) 1597–1620. doi:10.1098/rspa.2009.0434.
URL http://www.earth.ox.ac.uk/~richardk/publications/Katz_Worster_PRSA10.pdf
- [400] T. Tautges et al., SISIPHUS: Scalable ice-sheet solvers and infrastructure for petascale, high-resolution, unstructured simulations, <http://trac.mcs.anl.gov/projects/sisiphus/wiki>.
- [401] P. Lichtner (PI), Modeling Multiscale-Multiphase-Multicomponent Subsurface Reactive Flows Using Advanced Computing, sciDAC2

REFERENCES

- project, see <http://www.scidac.gov/groundwater/gwflow.html>.
- [402] M. McCourt, T. D. Rognlien, L. C. McInnes, H. Zhang, Improving parallel scalability for edge plasma transport simulations with neutral gas species, in: Twenty Second International Conference on Numerical Simulations of Plasmas, Sept. 7-9, 2011, Long Branch, NJ, 2012, to appear.
- [403] S. Abhyankar, Development of an implicitly coupled electromechanical and electromagnetic transients simulator for power systems, Ph.D. thesis, Illinois Institute of Technology (2011).
- [404] J. N. Shadid, E. C. Cyr, R. P. Pawlowski, R. S. Tuminaro, L. Chacón, P. T. Lin, Initial performance of fully-coupled AMG and approximate block factorization preconditioners for solution of implicit FE resistive MHD, in: J. C. F. Pereira, A. Sequeira (Eds.), V European Conference on Computational Fluid Dynamics, ECCOMAS CFD 2010, Lisbon Portugal, 2010.
- [405] R. P. Pawlowski, Phalanx Multiphysics Assembly Package, <http://trilinos.sandia.gov/packages/phalanx/> (2011).
- [406] P. K. Notz, R. P. Pawlowski, J. C. Sutherland, Graph-based software design for managing complexity and enabling concurrency in multi-physics pde software, ACM Trans. Math. Softw. Submitted.
- [407] J. N. Shadid, R. P. Pawlowski, J. W. Banks, L. Chacón, P. T. Lin, R. S. Tuminaro, Towards a scalable fully-implicit fully-coupled resistive MHD formulation with stabilized FE methods, J. Comput. Phys. 229 (20) (2010) 7649–7671.
- [408] L. C. Musson, R. P. Pawlowski, A. G. Salinger, T. J. Madden, K. B. Hewett, Multiphase reacting flow modeling of singlet oxygen generators for chemical oxygen iodine lasers, Proc. of SPIE 7131 (2009) 71310J:1–71310J:8.
- [409] P. T. Lin, J. N. Shadid, M. Sala, R. S. Tuminaro, G. L. Hennigan, R. J. Hoekstra, Performance of a parallel algebraic multilevel preconditioner for stabilized finite element semiconductor device modeling, Journal Computational Physics 228 (2009) 6250–6267.
- [410] S. G. Parker, A component-based architecture for parallel multi-physics PDE simulation, Future Generation Comput. Sys. 22 (1) (2006) 204–216.
- [411] S. G. Parker, J. Guilkey, T. Harman, A component-based parallel infrastructure for the simulation of fluid-structure interaction, Engineering with Computers 22 (3-4) (2006) 277–292.
- [412] J. Luitjens, B. Worthen, M. Berzins, T. C. Henderson, Chapter 4: Scalable parallel AMR for the Uintah multiphysics code (2008).
- [413] D. Pershing (PI), Center for the Simulation of Accidental Fires and Explosions (C-SAFE), Website.
URL <http://www.csafe.utah.edu/>
- [414] C. Feichtinger, S. Donath, H. Koestler, J. Goetz, U. Ruede, WaLberla: HPC software design for computational engineering simulations, Journal of Computational Science 2 (2) (2011) 105 – 112. doi:10.1016/j.jocs.2011.01.004.
- [415] K. Iglberger, U. Ruede, Large scale rigid body dynamics, Multibody System Dynamics 25 (1) (2011) 81–95. doi:10.1007/s11044-010-9212-0.
- [416] J. Götz, K. Iglberger, M. Stürmer, U. Ruede, Direct numerical simulation of particulate flows on 294,912 processor cores, in: Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis, SC '10, IEEE Computer Society, Washington, DC, USA, 2010, pp. 1–11. doi:10.1109/SC.2010.20.
- [417] D. E. Post, R. P. Kendall, Software project management and quality engineering practices for complex, coupled multiphysics, massively parallel computational simulations: Lessons learned from ASCI, International Journal of High Performance Computing Applications 18 (4) (2004) 399–416.
- [418] M. C. Miller, J. F. Reus, R. P. Matzke, Q. A. Koziol, A. P. Cheng, Smart libraries: Best SQE practices for libraries with emphasis on scientific computing, Tech. Rep. UCRL-JRNL-28636, Lawrence Livermore National Laboratory (2004).
- [419] M. Gee, U. Küttler, W. Wall, Truly monolithic algebraic multigrid for fluid-structure interaction, International Journal for Numerical Methods in Engineering 85 (8) (2011) 987–1016.
- [420] A. Barker, X.-C. Cai, Two-level Newton and hybrid Schwarz preconditioners for fluid-structure interaction, SIAM Journal on Scientific Computing 32 (4) (2010) 2395–2417.
- [421] J. Hron, S. Turek, A monolithic FEM/multigrid solver for an ALE formulation of fluid-structure interaction with applications in biomechanics, Fluid-Structure Interaction (2006) 146–170.
- [422] M. Adams, Algebraic multigrid methods for constrained linear systems with applications to contact problems in solid mechanics, Numerical Linear Algebra with Applications 11 (23) (2004) 141–153.
- [423] J. Mandel, C. Dohrmann, R. Tezaur, An algebraic theory for primal and dual substructuring methods by constraints, Applied numerical mathematics 54 (2) (2005) 167–193.
- [424] C. Dohrmann, O. Widlund, Hybrid domain decomposition algorithms for compressible and almost incompressible elasticity, International Journal for Numerical Methods in Engineering 82 (2) (2010) 157–183.
- [425] R. P. Pawlowski, E. T. Phipps, A. G. Salinger, Automating embedded analysis capabilities using template-based generic programming, Scientific Programming Submitted.
- [426] R. P. Pawlowski, E. T. Phipps, A. G. Salinger, Applying template-based generic programming to the simulation and analysis of partial differential equations, Scientific Programming Submitted.
- [427] A. Logg, K.-A. Mardal, G. N. Wells et al., Automated Solution of Differential Equations by the Finite Element Method, Springer, 2011.
- [428] K. Long, Solving partial differential equations with Sundance (2003).
- [429] COMSOL, <http://www.comsol.com>.
- [430] G. Baumgartner, A. Auer, D. E. Bernholdt, A. Bibireata, V. Choppella, D. Cociorva, X. Gao, R. J. Harrison, S. Hirata, S. Krishnamoorthy, S. Krishnan, C. Lam, Q. Lu, M. Nooijen, R. M. Pitzer, J. Ramanujam, P. Sadayappan, A. Sibiriyakov, D. E. Bernholdt, A. Bibireata, D. Cociorva, X. Gao, S. Krishnamoorthy, S. Krishnan, Synthesis of high-performance parallel programs for a class of ab initio quantum chemistry models, in: Proceedings of the IEEE, 2005, p. 2005.
- [431] G. Hansen, S. Owen, Mesh generation technology for nuclear reactor simulation; barriers and opportunities, Nuclear Engineering and Design 238 (10) (2008) 2590–2605.
- [432] L. Dubcova, P. Solin, G. Hansen, H. Park, Comparison of multimesh *hp*-FEM to interpolation and projection methods for spatial coupling of thermal and neutron diffusion calculations, Journal of Computational Physics 230 (4) (2011) 1182–1197. doi:10.1016/j.jcp.2010.10.034.

REFERENCES

- [433] R. W. Johnson, G. Hansen, C. Newman, The role of data transfer on the selection of a single vs. multiple mesh architecture for tightly coupled multiphysics applications, *Applied Mathematics and Computation* 217 (2011) 8943–8962. doi:10.1016/j.amc.2011.03.101.
- [434] P. Solin, et al., Hermes project, http://en.wikipedia.org/wiki/Hermes_Project.
- [435] R. M. O'Bara, M. W. Beall, M. S. Shephard, Attribute management system for engineering analysis, *Engineering with Computers* 18 (2002) 339–351.
- [436] M. W. Beall, M. S. Shephard, A general topology-based mesh data structure, *Int. J. Num. Method. Engng* 40 (9) (1997) 1573–1596.
- [437] J.-F. Remacle, M. S. Shephard, An algorithm oriented mesh database, *International Journal for Numerical Methods in Engineering* 58 (2003) 349–374.
- [438] E. S. Seol, M. S. Shephard, Efficient distributed mesh data structure for parallel automated adaptive analysis, *Engineering with Computers* 22 (3-4) (2006) 197–213.
- [439] A. Alexandrescu, *Modern C++ Design: Generic Programming and Design Patterns Applied*, C++ In-Depth, Addison-Wesley, 2001.
- [440] Kitware, Cmake: the cross-platform, open-source build system, <http://www.cmake.org/> (2011).
- [441] R. A. Bartlett, Thyra coding and documentation guidelines, Sandia Report SAND2010-2051, Sandia National Laboratories (May 2010).
- [442] A. S. Eddington. [link].
URL http://www.brainyquote.com/quotes/authors/a/arthur_eddington.html
- [443] L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, R. C. Whaley, *ScaLAPACK Users' Guide*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 1997.
- [444] C. T. Kelley, D. E. Keyes, Convergence analysis of pseudo-transient continuation, *SIAM J. Numer. Anal.* 35 (1998) 508–523.
- [445] J.-L. Lions, Y. Maday, G. Turinici, A parareal in time discretization of PDEs, *C. R. Acad. Sci. Paris Series I* 332 (2001) 661–668.
- [446] J. Dongarra, P. Beckman, et al., The international exascale software roadmap, *Int. J. High Perf. Comput. Applics.* 25 (2011) 3–60.
- [447] R. Stevens, Director, Institute for Computing in Science (ICiS), <http://www.icis.anl.gov/> (2011).
- [448] F. J. Blom, A monolithic fluid-structure interaction algorithm applied to the piston problem, *Comput Methods Appl Mech Eng* 167 (1998) 369–391.
- [449] H. G. Matthies, J. Steindorf, Partitioned but strongly coupled iteration schemes for nonlinear fluid-structure interaction, *Computers and Structures* 80 (2002) 1991–1999.
- [450] D. E. Keyes, L. C. McInnes, C. Woodward, Workshop Attendees, ICiS Multiphysics Workshop Website, <https://sites.google.com/site/icismultiphysics2011/> (2011).
- [451] J. R. Kamm, J. S. Brock, S. T. Brandon, D. L. Cotrell, B. Johnson, P. Knupp, W. Rider, T. Trucano, V. Weirs, Enhanced verification test suite for physics simulation codes, Tech. Rep. LA-14379, Los Alamos National Laboratory (Sept 2008).
URL <http://library.lanl.gov/cgi-bin/getfile?LA-14379.pdf>



Mathematics and Computer Science Division

Argonne National Laboratory
9700 South Cass Avenue, Bldg. 240
Argonne, IL 60439-4847

www.anl.gov



Argonne National Laboratory is a U.S. Department of Energy
laboratory managed by UChicago Argonne, LLC