

# Trilinos: Enabling Scientific Computing across Diverse Hardware Architectures at Scale

**MATTHIAS MAYR**, Institute for Mathematics and Computer-Based Simulation, Universität der Bundeswehr München, Neubiberg, Germany and Data Science & Computing Lab, Universität der Bundeswehr München, Neubiberg, Germany

**ALEXANDER HEINLEIN**, Delft Institute of Applied Mathematics (DIAM), Delft University of Technology, Delft, Netherlands

**CHRISTIAN GLUSA** and **SIVASANKARAN RAJAMANICKAM**, Sandia National Laboratories, Albuquerque, New Mexico, USA

**MAARTEN ARNST**, University of Liege, Liège, Belgium

**ROSCOE A. BARTLETT**, **LUC BERGER-VERGIAT**, and **ERIK G. BOWMAN**, Sandia National Laboratories, Albuquerque, New Mexico, USA

**KAREN DEVINE**, Retired, Starkville, Mississippi, USA

**GRAHAM HARPER**, Sandia National Laboratories, Albuquerque, New Mexico, USA

**MICHAEL HEROUX**, ParaTools, Avon, Minnesota, USA

**MARK HOEMMEN**, NVIDIA Corp, Santa Clara, California, USA

**JONATHAN J. HU**, Sandia National Laboratories, Livermore, California, USA

**BRIAN KELLEY**, Sandia National Laboratories, Albuquerque, New Mexico, USA

**KYUNGJOO KIM**, NVIDIA Corp, Santa Clara, California, USA

**DREW P. KOURI** and **PAUL KUBERRY**, Sandia National Laboratories, Albuquerque, New Mexico, USA

**KIM LIEGEOIS**, Advanced Micro Devices, Inc., Santa Clara, California, USA

**CURTIS C. OBER**, **ROGER P. PAWLOWSKI**, **CARL PEARSON**, **MAURO PEREGO**, **ERIC PHIPPS**, **DENIS RIDZAL**, **NATHAN V. ROBERTS**, **CHRISTOPHER M. SIEFERT**, and **HEIDI K. THORNQUIST**, Sandia National Laboratories, Albuquerque, New Mexico, USA

**ROMIN TOMASETTI**, University of Liege, Liège, Belgium

**CHRISTIAN R. TROTT**, Sandia National Laboratories, Albuquerque, New Mexico, USA

**RAYMOND S. TUMINARO**, Sandia National Laboratories, Livermore, California, USA

**JAMES M. WILLENBRING**, **MICHAEL M. WOLF**, and **ICHITARO YAMAZAKI**, Sandia National Laboratories, Albuquerque, New Mexico, USA



This work is licensed under [Creative Commons Attribution International 4.0](https://creativecommons.org/licenses/by/4.0/).

© 2026 Copyright held by the owner/author(s).

ACM 1557-7295/2026/6-ART12

<https://doi.org/10.1145/3802822>

Christian Glusa, Alexander Heinlein, Matthias Mayr, and Siva Rajamanickam contributed equally to this article and, thus, share the first authorship.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. This article describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the article do not necessarily represent the views of the U.S. Department of Energy or the United States Government. This work was supported in part by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Scientific Discovery through

Trilinos is a community-developed, open-source software framework that facilitates building large-scale, complex, multiscale, multiphysics simulation code bases for scientific and engineering problems. Since the Trilinos framework has undergone substantial changes to support new applications and new hardware architectures, this document is an update to “An Overview of the Trilinos project” by Heroux et al. (ACM Transactions on Mathematical Software, 31(3):397–423, 2005). It describes the design of Trilinos, introduces its new organization in product areas, and highlights established and new features available in Trilinos. Particular focus is put on the modernized software stack based on the Kokkos ecosystem to deliver performance portability across heterogeneous hardware architectures. This article also outlines the organization of the Trilinos community and the contribution model to help onboard interested users and contributors.

CCS Concepts: • **Mathematics of computing** → **Solvers**; **Mathematical software performance**; • **Computing methodologies** → **Massively parallel algorithms**; • **Software and its engineering** → **Software libraries and repositories**; **Software performance**;

Additional Key Words and Phrases: Heterogeneous Hardware Architectures, High-Performance Computing, Performance Portability, Scientific Software Frameworks

---

Advanced Computing (SciDAC) Program through the FASTMath Institute. This work was supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program. This work was supported by the Laboratory Directed Research and Development program at Sandia National Laboratories. This work was supported in part by dtec.bw—Digitalization and Technology Research Center of the Bundeswehr [project hpc.bw]. dtec.bw is funded by the European Union—NextGenerationEU.

Authors’ Contact Information: Matthias Mayr, Institute for Mathematics and Computer-Based Simulation, Universität der Bundeswehr München, Neubiberg, Germany and Data Science & Computing Lab, Universität der Bundeswehr München, Neubiberg, Germany; e-mail: matthias.mayr@unibw.de; Alexander Heinlein (corresponding author), Delft Institute of Applied Mathematics (DIAM), Delft University of Technology, Delft, Netherlands; e-mail: a.heinlein@tudelft.nl; Christian Glusa, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: caglusa@sandia.gov; Sivasankaran Rajamanickam, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: srajama@sandia.gov; Maarten Arnst, University of Liege, Liège, Belgium; e-mail: maarten.arnst@uliege.be; Roscoe A. Bartlett, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: rabartl@sandia.gov; Luc Berger-Vergiat, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: lberge@sandia.gov; Erik G. Bowman, Center for Computing Research, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: egboman@sandia.gov; Karen Devine, Retired, Starkville, Mississippi, USA; e-mail: devine.hpc@gmail.com; Graham Harper, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: gbharpe@sandia.gov; Michael Heroux, ParaTools, Avon, Minnesota, USA; e-mail: mheroux@paratools.com; Mark Hoemmen, NVIDIA Corp, Santa Clara, California, USA; e-mail: mhoemmen@nvidia.com; Jonathan J. Hu, Sandia National Laboratories, Livermore, California, USA; e-mail: jhu@sandia.gov; Brian Kelley, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: bmkelle@sandia.gov; Kyungjoo Kim, NVIDIA Corp, Santa Clara, California, USA; e-mail: kyungjook@nvidia.com; Drew P. Kouri, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: dpkouri@sandia.gov; Paul Kuberry, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: pakuber@sandia.gov; Kim Liegeois, Advanced Micro Devices, Inc., Santa Clara, California, USA; e-mail: kim.liegeois@amd.com; Curtis C. Ober, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: ccober@sandia.gov; Roger P. Pawlowski, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: rppawlo@sandia.gov; Carl Pearson, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: cwpears@sandia.gov; Mauro Perego, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: mperego@sandia.gov; Eric Phipps, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: etphipp@sandia.gov; Denis Ridzal, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: dridzal@sandia.gov; Nathan V. Roberts, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: nvrober@sandia.gov; Christopher M. Siefert, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: csiefer@sandia.gov; Heidi K. Thornquist, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: hkthorn@sandia.gov; Romin Tomasetti, University of Liege, Liège, Belgium; e-mail: romin.tomasetti@uliege.be; Christian R. Trott, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: crtrott@sandia.gov; Raymond S. Tuminaro, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: rstumin@sandia.gov; James M. Willenbring, Scalable Algorithms, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: jmwill@sandia.gov; Michael M. Wolf, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: mmwolf@sandia.gov; Ichitaro Yamazaki, Sandia National Laboratories, Albuquerque, New Mexico, USA; e-mail: iyamaza@sandia.gov.

**ACM Reference format:**

Matthias Mayr, Alexander Heinlein, Christian Glusa, Sivasankaran Rajamanickam, Maarten Arnst, Roscoe A. Bartlett, Luc Berger-Vergiat, Erik G. Bowman, Karen Devine, Graham Harper, Michael Heroux, Mark Hoemmen, Jonathan J. Hu, Brian Kelley, Kyungjoo Kim, Drew P. Kouri, Paul Kuberry, Kim Liegeois, Curtis C. Ober, Roger P. Pawlowski, Carl Pearson, Mauro Perego, Eric Phipps, Denis Ridzal, Nathan V. Roberts, Christopher M. Siefert, Heidi K. Thornquist, Romin Tomasetti, Christian R. Trott, Raymond S. Tuminaro, James M. Willenbring, Michael M. Wolf, and Ichitaro Yamazaki. 2026. Trilinos: Enabling Scientific Computing across Diverse Hardware Architectures at Scale. *ACM Trans. Math. Softw.* 52, 2, Article 12 (June 2026), 34 pages. <https://doi.org/10.1145/3802822>

---

## 1 Introduction

Trilinos is a community-driven open-source C++ software framework and collection of reusable scientific libraries designed to enable the development of scalable, high-performance algorithms for solving complex, multiscale, and multiphysics engineering and scientific problems on advanced computing architectures. While Trilinos can run on a variety of hardware platforms ranging from small workstations to large supercomputers, the typical use of Trilinos is on leadership-class systems with new or emerging hardware architectures.

Trilinos was originally conceived as a framework of three packages for distributed memory systems. The original Trilinos publication [69] described the motivation, philosophy, and capabilities of Trilinos at that time. Seven years later, a two-part collection of 16 papers provided a new overview of Trilinos and many of its packages [12, 22, 26, 29, 44, 49, 70–72, 76, 94, 100, 106, 109, 110, 137] and described the intervening expansion of Trilinos' capabilities and strategic goals for Trilinos. Trilinos today hews to the overall vision from two decades ago, but has evolved alongside changes in programming models, application needs, hardware architectures, and algorithms. Trilinos has grown from a library of three packages to more than 35 packages with functionality and features supporting a wide range of applications.

For much of its history, packages have relied heavily on the linear algebra classes provided by Epetra. Epetra did not support larger problems (2B+ unknowns), nor did it allow for accelerator offload. Since the packages of the old Epetra stack were archived and removed from the main Trilinos repository on GitHub by the end of 2025, this article describes only the modern Trilinos software stack, which builds upon Tpetra and the Kokkos ecosystem to achieve performance portability across hardware architectures.

This article attempts to capture a snapshot the current state of Trilinos, as an update to the articles published 14 and 21 years ago. Therefore it focuses on the major developments within Trilinos in the last decade as well as new features and functionality that have been added to advance scientific and engineering applications. It will give only an overview of the features, and we refer to the extensive list of references for the details of these features. We are also cognizant of the fact that this article describes software that is being actively developed and constantly evolving. Hence, we will focus on the high-level features and concepts that we expect to remain stable for several years. We also briefly touch upon the Trilinos community and software engineering issues with respect to Trilinos.

### 1.1 Modern Trilinos: Performance Portability through Kokkos

A key goal of modern Trilinos is to offer a performance-portable collection of reusable scientific libraries, allowing users to develop applications that achieve high efficiency across all modern **High Performance Computing (HPC)** hardware architectures. This goal emerged when the HPC architecture landscape started diversifying with the introduction of GPU acceleration for scientific

software. Today, nine of the top 10 systems in the Top500<sup>1</sup> use GPUs, with only a single system using CPUs only. A challenge in this shift is the diversification of vendor-provided programming models: the aforementioned systems have four different preferred programming models: CUDA, OpenACC, HIP, and SYCL for the GPU-based systems and OpenMP for the CPU system.

To avoid the necessity of diverging code paths, Trilinos leverages the Kokkos ecosystem<sup>2</sup> [140] to write performance portable code. The Kokkos ecosystem originated from the native Trilinos packages Kokkos [68] and Kokkos Kernels but split off a decade ago into a stand-alone project hosted and developed independently of Trilinos.<sup>3</sup> Now, Trilinos provides snapshots of the two primary Kokkos subprojects (Core and Kernels) that reflect the latest release of Kokkos. Trilinos can also be built against version-compatible external installations of Kokkos—a capability required for interoperability with other Kokkos-based libraries.

Kokkos enables Trilinos developers to write single source implementations of their packages that perform well on all major HPC hardware architectures. Some of its design principles are also reflected in Trilinos API designs. In particular, the use of `execution spaces` (where does code run) and `memory spaces` (where does data reside) as parameters of data structures and algorithms is a common thread throughout Trilinos. This design characteristic enables Trilinos users to manage the heterogeneity of modern architectures and its impact on Trilinos data structures.

## 1.2 Trilinos Functionality

The features and capabilities of Trilinos are divided into several software units called *packages*. Each Trilinos package has a well-defined set of unique capabilities that is important for scientific or engineering applications. Packages are semi-autonomous, often having their own development team, users, and set of requirements and development principles. However, packages also follow a general set of Trilinos expectations such as having a designated point-of-contact and following software quality expectations (e.g., sufficient documentation, continuous integration testing, clearly defined dependencies, and using the Trilinos infrastructure for building and installation).

While most packages are developed natively in the Trilinos repository, some packages are developed externally and “snapshotted” into Trilinos, i.e., incorporated as fixed copies of their source code at a specific point in time. These are (in alphabetical order): Compadre,<sup>4</sup> Kokkos and Kokkos Kernels, Krino, Rapid Optimization Library (ROL),<sup>5</sup> SEACAS,<sup>6</sup> Sierra ToolKit (STK), and Tribal Build, Integration, and Test System (TriBITS).<sup>7</sup> Some reasons for snapshotting external packages into Trilinos include: external packages and repositories facilitate easy and independent usage by other software projects without having to pull in Trilinos; snapshotting provides an integration process for external packages into Trilinos, ensuring continual compatibility between all packages; and snapshotting allows easy integration into Trilinos-based application codes, alleviating users from managing these dependencies themselves. We refer to the Trilinos website (<https://trilinos.github.io>) for a brief description of these packages.

In the following sections of this article, we group the Trilinos packages that share common objectives (e.g., solving linear systems) together into Trilinos’ five *product areas*: Core, Linear Solvers and Preconditioners, Nonlinear Solvers and Analysis Tools, Discretization Tools, and Framework. Trilinos’ packages and their assignment to a product area are illustrated in Figure 1.

<sup>1</sup><https://top500.org>, 66th edition from 17 November 2025.

<sup>2</sup><https://kokkos.org>.

<sup>3</sup><https://github.com/kokkos>.

<sup>4</sup><https://github.com/sandialabs/compadre>.

<sup>5</sup><https://github.com/sandialabs/rol>.

<sup>6</sup><https://github.com/sandialabs/seacas>.

<sup>7</sup><https://github.com/TriBITSPub/TriBITS>.

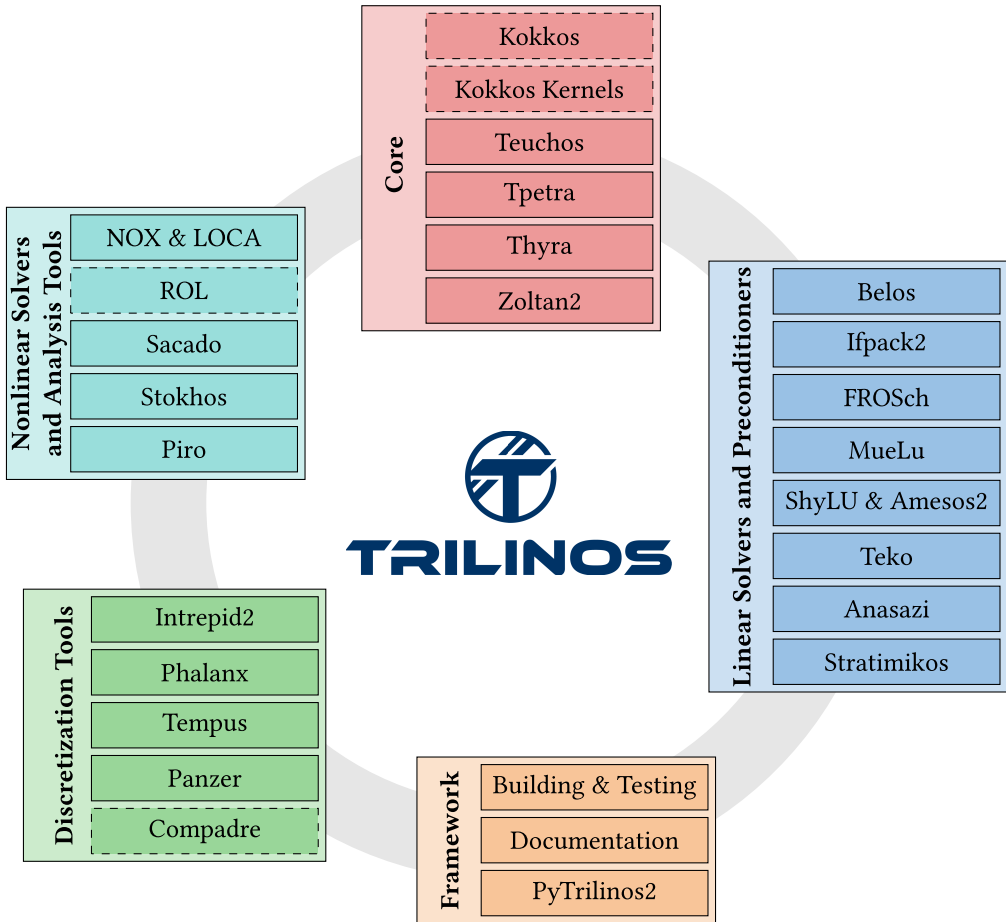


Fig. 1. Organization of the Trilinos library into five product areas and their packages: Product areas are labeled in bold font with their respective packages. Snapshotted packages described in this article are marked by dashed boxes.

We briefly describe these product areas below.

**Core.** Core packages cover all aspects of creating, distributing, and mapping data to processing elements (cores, threads, nodes), load balancing, and redistributing data. This includes Trilinos’ abstractions for linear algebra data structures and algorithms as well as concrete implementations such as Tpetra linear algebra data structures. On a modern accelerator-based compute node, the abstractions provided by the Kokkos library become critical for Tpetra. These capabilities are described in detail in Section 2.

**Linear Solvers and Preconditioners.** The wide variety of applications that use Trilinos need a diverse set of linear solvers. Trilinos has support for both iterative and direct linear solvers, including interfaces to external solver packages. There are a number of preconditioner options from multithreaded or performance portable node-level preconditioners to scalable multilevel domain decomposition or multigrid preconditioners. The preconditioners and solvers use the data abstractions from the core packages. Section 3 provides a detailed description of these features.

*Nonlinear Solvers and Analysis Tools.* These packages provide algorithms for computational simulations and design. Capabilities include solvers for nonlinear equations, parameter continuation, bifurcation tracking, optimization, and uncertainty quantification. Trilinos also provides utility packages to evaluate quantities of interest required by the analysis algorithms. Additional capabilities include **Automatic Differentiation (AD)** technology to evaluate derivatives and embedded ensemble propagation for uncertainty quantification. These packages and their capabilities will be discussed further in Section 4.

*Discretization Tools.* This collection of packages provides functionality for the discretization of differential equations. In particular, it supports mesh-free and mesh-based spatial discretizations, with a focus on high-order finite elements, and time integration. Discretization tools also include support for algorithmic differentiation and for managing **Directed Acyclic Graphs (DAGs)** of evaluation kernels. These capabilities are described further in Section 5 in detail.

*Framework.* The Framework product is different than the other Trilinos products in that most of the resources and services are not associated with Trilinos packages. The Framework product area is primarily focused on activities such as developing and maintaining infrastructure for automated testing and documentation, as well as associated workflows. A small number of infrastructure and cross-cutting packages are also associated with the Framework, for example, the Python wrapper package PyTrilinos2. Section 6 provides further details.

### 1.3 Outline of This Article

Sections 2–6 provide details on Trilinos’ product areas. Section 7 briefly touches upon the Trilinos community, software engineering aspects as well as the integration of Trilinos into the landscape of scientific software frameworks. Finally, Section 8 concludes the article.

In comparison to the first Trilinos publication from 2005 [69] and its updates from 2012 [12, 22, 26, 29, 44, 49, 70–72, 76, 94, 100, 106, 109, 110, 137], the present article not only provides updates on newly added capabilities, but more importantly reflects the overhaul of Trilinos towards performance portability through the full integration of the Kokkos ecosystem [45, 140, 141]. Many packages described in the original Trilinos publication [69] have been replaced by Kokkos-capable successor implementations, among them (in alphabetical order) Amesos2 replacing Amesos, Belos replacing AztecOO, Ifpack2 replacing Ifpack, MueLu replacing ML, and Tpetra replacing Epetra. The present article introduces these new packages and their capabilities and outlines their interplay with Kokkos.

## 2 Trilinos Core

The Core product area provides essential tools for managing and distributing data across processing elements in parallel computing environments, addressing challenges such as load balancing and data redistribution. Key components discussed here include the Kokkos performance portability framework, Teuchos tools, Tpetra distributed-memory linear algebra, Thyra abstractions for high-level formulations of numerical algorithms, and Zoltan2 for combinatorial algorithms in parallel computing.

### 2.1 Performance Portability: Kokkos Core

Kokkos Core<sup>8</sup> is a programming model designed for performance portability across hardware architectures such as CPUs and GPUs. While the development of Kokkos Core has originated

<sup>8</sup>For historical reasons the package name is “Kokkos” since it precedes the creation of other Kokkos subprojects such as Kokkos Kernels.

as a package of Trilinos, Kokkos has become an independent project in 2015 and is nowadays snapshotted into Trilinos for the pure convenience of dependency management. Two primary insights into achieving performance portability shaped Kokkos' development: (i) abstractions for both parallel execution and data structures are required, and (ii) the fundamental programming model must be descriptive rather than prescriptive.

The Kokkos programming model includes execution spaces to specify where an algorithm should run, execution policies to indicate how the algorithm should be run in terms of hardware resources, and execution patterns to designate whether to execute a `parallel_for`, `parallel_reduce`, or `parallel_scan`. The data structure abstractions are built around the concepts of memory spaces to indicate where memory is allocated, memory layouts to specify how data is laid out in memory, and memory traits that enable customizations of how data is accessed. The primary data structure incorporating these concepts is `Kokkos::View`, a multidimensional array abstraction that inspired the ISO C++23 `mdspan` class.<sup>9</sup> Detailed descriptions of Kokkos Core can be found in [45] and [141]. Development of Kokkos Core itself is carried out in its own GitHub repository.<sup>10</sup>

## 2.2 Performance Portable Kernels: Kokkos Kernels

Kokkos Kernels [123] is part of the Kokkos ecosystem [140] and provides node-local implementations of mathematical kernels widely used across packages in Trilinos. As a member of the Kokkos ecosystem, Kokkos Kernels tightly integrates Kokkos features and aims at delivering performance-portable algorithms across major CPU- and GPU-based HPC systems. Due to its node-local nature, Kokkos Kernels does not rely on MPI or other communication libraries, unlike numerous other packages in Trilinos.

The implementations of Kokkos Kernels algorithms leverage the hierarchical parallelism exposed by the Kokkos library [82]. To ensure flexibility for the distributed libraries that might call its algorithms, Kokkos Kernels provides thread-safe implementations for most of its kernels, with increasing coverage for asynchronous implementations that allow an execution space instance to be passed when calling a kernel. Kokkos Kernels also serves as a major point of integration for vendor optimized libraries such as cuBLAS, cuSPARSE, rocBLAS, rocSPARSE, MKL, ARMpl, and others.

The capabilities provided by Kokkos Kernels can be divided into four major categories: (1) BLAS algorithms, (2) sparse linear algebra and preconditioners, (3) graph algorithms, and (4) batched dense and sparse linear algebra [3, 92]. The main points of integration of Kokkos Kernels in Trilinos are Tpetra for the dense and sparse linear algebra capabilities, Ifpack2 for the preconditioners and batched algorithms, and the multigrid package MueLu which relies directly and indirectly on BLAS and graph algorithms and specialized implementations of graph coloring/coarsening [80] and fused Jacobi-SpGEMM (generalized sparse matrix-matrix multiply) kernels. Additionally, batched dense factorization algorithms are used in the construction of grid transfers.

Similar to the Kokkos library, Kokkos Kernels is developed in its own GitHub repository<sup>11</sup> outside of the Trilinos GitHub repository. Every version of the library is integrated and tested in Trilinos as part of the Kokkos ecosystem release process. Additional information on Kokkos Kernels' capabilities can be found in [42, 148].

<sup>9</sup><https://www.open-std.org/jtc1/sc22/wg21/docs/papers/2022/p0009r18.html>.

<sup>10</sup><https://github.com/kokkos/kokkos>.

<sup>11</sup><https://github.com/kokkos/kokkos-kernels>.

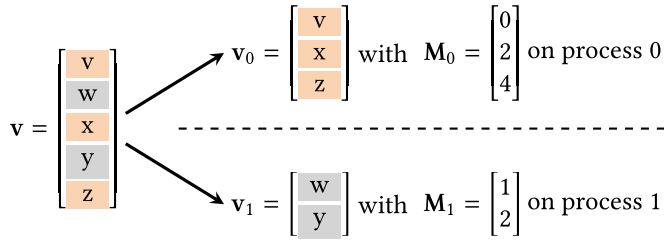


Fig. 2. Illustration of a `Tpetra::Map` for an exemplary vector  $v \in \mathbb{R}^5$  distributed to two processes  $p \in \{0, 1\}$ : Entries  $v$ ,  $x$ , and  $z$  are stored on process 0, entries  $w$  and  $y$  are stored on process 1. Local-to-global mapping of indices is given by  $M_p$ , where the  $i$ th element of  $M_p$  denotes the unique global index corresponding to the local index  $i$  on rank  $p$ . In more complicated scenarios, entries can also be shared between processes.

### 2.3 Tools: Teuchos

Teuchos provides a collection of helper and utility components designed to simplify common programming tasks and to ensure a consistent coding style across Trilinos packages. These tools include memory management classes [19] such as smart pointers and arrays, parameter lists for communicating hierarchical lists of parameters between library or application layers, templated wrappers for BLAS and LAPACK, XML parsers, MPI, and other supporting functionality. Teuchos thereby provides a unified “look and feel” across Trilinos packages and helps to avoid common programming mistakes. It offers reusable building blocks for such recurring tasks implemented in a consistent and coherent manner to be used throughout the entire code base.

### 2.4 Distributed-Memory Linear Algebra: Tpetra

Tpetra [12, 74] provides the distributed-memory infrastructure for linear algebra objects such as sparse graphs, sparse matrices, and dense vectors. These linear algebra objects are implemented using MPI-based abstractions for the inter-process communication, and Kokkos-based data structures for the portions of the object owned by each MPI process. Kokkos Kernels is used for performing local, parallel computations within each MPI process. Due to the fundamental nature of MPI distributed data structures, the majority of Trilinos packages depend directly or indirectly on Tpetra.

The objects provided are templated on scalar type (e.g., double, float, or a Sacado or Stokhos type), local index type, global index type, and Kokkos device (pair of execution and memory spaces). The distribution of the linear algebra objects is described by a Map data structure, which defines which elements are assigned to which MPI ranks. Here, “element” could be a row or column of a sparse matrix or an entry of a vector. Global indices number elements across an entire distributed object, while local indices number elements only within a rank. The Map encodes the relation between local indices and global indices. An illustration for a vector, that is distributed to two MPI processes, is given in Figure 2. Since memory capacity limits the sizes of local objects, local indices can be safely represented with a more compact type than global indices (e.g., 32-bit `int` instead of 64-bit `long`).

`DistObject` is Tpetra’s abstract base class for general distributed objects. Subclasses define how to pack elements into a buffer for sending with MPI, and how to unpack received elements. The `Import` and `Export` classes describe how one `DistObject` communicates data to another. They are both created from a pair of Maps, a source and a target. For example, halo exchange on a vector is an `Import` where the source is the vector’s original distribution, and the target is the same except that each rank also “owns” the degrees-of-freedom in the halo around its subdomain.

The following are the primary linear algebra objects in Tpetra (all are subclasses of `DistObject`):

- *(Multi)Vectors* enable storage of dense vectors or collections of vectors (multivectors) and associated BLAS-1 like kernels (e.g., dot products, norms, scaling, vector addition, pointwise vector multiplication) as well as tall skinny QR (TSQR) factorization for multivectors.
- *CrsGraph* implements a sparse graph in **Compressed Row Storage (CRS)** format. Graphs also include import/export objects for use in halo/boundary exchanges associated with the graph.
- *CrsMatrix* and *BlockCrsMatrix* are sparse matrices in CRS and **Block Compressed Sparse Row (BSR)** formats. Associated format-specific computational kernels include sparse matrix-vector product (SPMV), sparse matrix-matrix multiplication (SPGEMM) and triple-product, sparse matrix-matrix addition, sparse matrix transpose, diagonal extraction, Frobenius norm calculation, and row/column scaling.

## 2.5 Abstract Numerical Algorithms (ANAs): Thyra

The Thyra package provides abstractions for implementing high-level ANAs. Thyra is used by other Trilinos packages for implementing ANAs such as solvers for linear and nonlinear equations, bifurcation and stability analysis, optimization, and uncertainty quantification. The abstractions hide concrete linear algebra and solver implementations so that algorithm development is separated from the implementation of the underlying kernels, allowing the same code to be written once and reused across many application codes. For example, one could write a Krylov solver and, without modifications to computational kernels that are hidden behind the abstraction layer, use it with multiple concrete linear algebra implementations such as Tpetra in Trilinos or the vector/matrix objects in PETSc [14].

Thyra is organized into three sets of abstractions depending on the type of operation required. All are templated on the scalar type. These sets are:

- *Vector/Operator Interfaces*: The first level is an abstraction for vectors and operators [18]. This set includes abstractions for vector spaces, vectors, multivectors, and linear operators. Thyra provides a default set of vector-vector and vector-operator operations for standard linear algebra tasks such as AXPY from BLAS. Defining an extensible interface such that users can add arbitrary vector/matrix operations and maintain performance can be difficult. A general tool to achieve this is the Trilinos utility package RTOp [21], which provides an interface for users to add new Thyra operations. RTOp accepts a vector of Thyra input objects and a vector of Thyra output objects along with a functor. RTOp will then apply the functor to each element of the data.
- *Operator Solve Interfaces*: The second set of abstractions supplies abstract linear solvers that can be attached to operators. The interfaces include base classes for preconditioners, preconditioner factories, linear operators with a solver, and factories for linear operators with a solver. “Factories” build and update the preconditioner and linear solver objects as needed. The solver interfaces allow one to wrap any concrete linear solver or preconditioner for use with the ANAs. As a concrete implementation of this concept, the Stratimikos package wraps all of the Trilinos preconditioners and linear solvers in Thyra abstractions. Run-time configuration of Stratimikos is driven by  `Teuchos::ParameterList` options. Details on Stratimikos are given in Section 3.8.
- *Nonlinear Interfaces*: The final set of abstractions supports nonlinear operations. Thyra provides an interface that wraps nonlinear solvers. Implicit time integration and optimization solvers typically require nonlinear solves and can abstract away the concrete implementation via Thyra. The **Nonlinear Object-Oriented Solutions (NOX)** package in Trilinos (described in

Section 4.1) provides a concrete implementation of this interface. The final interface in Thyra is for querying a physics application for common objects required by Newton-based solvers. The `Thyra::ModelEvaluator` interface allows codes to ask the application to evaluate residuals, Jacobians, parametric sensitivities, Hessians, and post-processed quantities/derivatives of interest. The interface is designed to be stateless (i.e., consecutive calls with the same inputs should produce the same outputs). This design was chosen to avoid pitfalls in complex application codes interacting with general nonlinear algorithms. The design does not preclude its use in stateful applications that have a “memory” between timesteps, but rather forces the explicit use of the stateful information in the interface.

## 2.6 Combinatorial Algorithms with Hybrid Parallelism: Zoltan2

Zoltan2 is a package for partitioning, load balancing, and other combinatorial scientific computing algorithms. It can be viewed as a successor to Zoltan [29] that provides both tighter integration with other Trilinos packages and support for hybrid parallelism. It is used primarily to partition and distribute data objects (such as sparse matrices) across processors.

Zoltan2 has a highly modular design. The user provides input via input adapters. Adapters for sparse graphs and matrices are provided for Xpetra, which can also be used to interface Tpetra data structures. The models create a graph, hypergraph, or coordinate model from the user input, which is typically a sparse matrix, mesh, or vector. The algorithms then use this model to solve the particular problem (e.g., partitioning or coloring) requested by the user.

Zoltan2 supports the use of both distributed-memory parallelism and on-node parallelism on CPUs and GPUs by building on Kokkos. Algorithms that support such hybrid parallelism include Multi-jagged [41] (geometric partitioning), Sphynx [4] (spectral graph partitioning), graph coloring [28], and MPI task placement [40].

## 3 Linear Solvers and Preconditioners

Trilinos offers many linear solver capabilities: dense and sparse direct solvers, iterative solvers, shared-memory preconditioners local to a compute node, and scalable distributed memory domain decomposition and multigrid methods. Furthermore, interfaces to several third-party direct solvers are provided. All native solver capabilities are built on top of Kokkos and are GPU capable to varying degrees; any exceptions are noted in the detailed descriptions below.

### 3.1 Iterative Linear Solvers and Krylov Methods: Belos

Belos [22] provides next-generation iterative linear solvers and a powerful framework to develop solvers for linear systems and least-squares problems. This framework is powerful because it provides several algorithmic abstractions that facilitate code reuse and extensibility. These abstractions are at every level of the iterative solver hierarchy, from the foundational linear algebra to the solver managers that embody the overall solution strategy, facilitating flexible solver development.

Belos fundamentally decouples the algorithms from the implementation of the underlying linear algebra object using abstract traits classes that prescribe the required matrix and vector operations. Adapters are necessary to concrete linear algebra implementations, which are provided for Tpetra and Thyra. Users can also implement their own interfaces to leverage existing matrix and vector representations. Next, there are abstractions to orthogonalization to ease the integration of application or architecture-specific orthogonalization methods. Implementations of iterated classical Gram-Schmidt, DGKS-corrected (Daniel, Gragg, Kaufman, and Stewart-corrected) classical Gram-Schmidt, and iterated modified Gram-Schmidt algorithms are provided. Finally, solver managers encapsulate the strategy for solving a linear system or least-squares problem using abstract interfaces to iteration kernels. As a result, the algorithms developed using Belos abstractions

can be relatively agnostic to data layout in memory or distributed over processors and parallel matrix/vector operations.

Belos supersedes the AztecOO package [67] by providing solver managers for individual iterative methods. A solver manager wraps a specific iteration and provides the high-level control logic and supporting functionality for one method, such as CG or GMRES, rather than for a collection of methods. Belos includes solver managers for single-vector iterative solvers, as well as extensions of several of these methods to block iterative methods for solving linear systems with multiple right-hand sides. Single-vector solver managers include conjugate gradient (CG), minimum residual (MINRES), generalized minimal residual (GMRES), stabilized biconjugate gradient (BiCGStab), and transpose-free QMR (TFQMR) as well as “seed” solvers (hybrid GMRES, PCPG), subspace recycling solvers (GCRO-DR and RCG), and least-squares solvers (LSQR). There are versions of CG, GMRES, and GCRO-DR that construct a block Krylov subspace to solve for one or more right-hand sides. There are also “pseudo-block” versions of CG, GMRES, and TFQMR that apply the single-vector iteration simultaneously on multiple right-hand sides, where matrix and preconditioner applications are aggregated to achieve better performance. Any single-vector iteration can be used to solve multiple right-hand sides as well but, if a block or pseudo-block version is not used, the solver will sequentially solve the right-hand sides.

In recent years, application or architecture-specific variants of the classic iterative methods have been integrated into Belos. Some of these variations are so minor that they are enabled through input parameters on the classic method. For instance, flexible GMRES [125] is an option for the block GMRES solver, and pipelined CG [53] is an option for both the block CG and pseudo-block CG solver. Additionally, there are stand-alone solvers for a pseudo-block stochastic CG method [108] and a fixed-point iteration method.

### 3.2 One-Level Domain Decomposition and Basic Iterative Methods: Ifpack2

Trilinos provides domain decomposition approaches in two different packages: Ifpack2 and ShyLU (specifically the ShyLU\_DD and FROSch subpackages). Ifpack2 implements overlapping additive Schwarz approaches with several options for the local subdomain solves. The local subdomain solvers may be either CPU-only versions of incomplete factorization preconditioners implemented in Ifpack2 itself, such as ILU(k) and ILUt (thresholded ILU), or architecture-portable algorithms for incomplete factorizations and triangular solvers implemented in Kokkos Kernels. It is also possible to use direct solvers as subdomain solvers, and ShyLU also supports shared-memory inexact incomplete factorization preconditioners. Such one-level preconditioners can be used as smoothers within multigrid methods, for solving “simpler” problems where the setup cost of more robust and scalable multilevel methods is prohibitive in comparison to the reduction in the number of iterations, or when the underlying problem is simply not amenable to multilevel methods.

Ifpack2 also supplies classic iterative methods based on matrix-splitting techniques, such as Jacobi iteration, Gauss-Seidel, and an MPI-oriented hybrid of Jacobi and Gauss-Seidel (e.g., Jacobi between ranks and Gauss-Seidel on them). Ifpack2 also provides preconditioners based on Chebyshev iterations. The aforementioned preconditioners are available both in point and block forms and can operate on CRS and BSR matrices. In the block case, line relaxation is also supported, while in the point case, techniques like Vanka relaxation [143] are possible. Auxiliary-space smoothing for Maxwell and Darcy type problems in the style of Hiptmair [73] are also supported. Local kernels are either implemented in Ifpack2 itself but can also be called from Kokkos Kernels for performance-portable shared-memory algorithmic variants of Gauss-Seidel and Jacobi iterations.

### 3.3 Multilevel Domain Decomposition Methods: FROSch

**FROSch (Fast and Robust Overlapping Schwarz)** is a framework for the construction of multilevel Schwarz domain decomposition preconditioners. Besides *parallel scalability*, FROSch emphasizes *applicability and robustness across a wide range of challenging problems* while supporting *algebraic construction* of the preconditioners, that is, solely from the fully assembled system matrix.

Most FROSch preconditioners can be constructed algebraically, though some variants can take advantage of additional geometric input. The algebraic construction relies on two key algorithmic components: first, the creation of an overlapping domain decomposition at the initial level based on the sparsity pattern of the system matrix, similar to Ifpack2; and second, the integration of extension-based coarse spaces, as in the classical two-level **Generalized Dryja–Smith–Widlund (GDSW)** preconditioner [43] and its variants. While the initial version of FROSch [63] was based on the outdated Epetra linear algebra framework, the current implementation [62] leverages Xpetra. Originally designed as a lightweight wrapper around Epetra and Tpetra, the Xpetra package facilitated over the years compatibility with both the Epetra and Tpetra stacks. With the deprecation of Epetra, it now exclusively supports the Tpetra stack. Algorithmic variants of Schwarz methods implemented in FROSch include:

- *Extension-based coarse spaces based on a partition of unity on the interface*, such as classical GDSW, reduced dimension GDSW coarse spaces, and **Multiscale Finite Element Method (MsFEM)** coarse spaces; cf. [63, 64];
- *Monolithic Schwarz preconditioners* for block systems; cf. [58].
- *Multilevel Schwarz preconditioners*, which are obtained from two-level Schwarz preconditioners by recursively applying Schwarz preconditioners as an inexact solver for the coarse problems; cf. [66].

The wide applicability of FROSch has been demonstrated for various challenging application problems, including: scalar elliptic and elasticity problems [63], possibly with heterogeneities [5]; **Computational Fluid Dynamics (CFD)** problems [58, 61]; time-harmonic Maxwell's and fluid-structure interaction problems [59]; pharmaco-mechanical interactions in arterial walls [15]; and coupled multiphysics problems for land ice simulations [65]. The latter three have been solved using monolithic preconditioning techniques. To extend robustness for heterogeneous model problems, an implementation of spectral coarse spaces [60] is currently under development. FROSch preconditioners have scaled to more than 200,000 cores on the Theta Cray XC40 supercomputer at the Argonne Leadership Computing Facility; cf. [66].

In its current implementation, FROSch assumes a one-to-one correspondence of subdomains and MPI ranks. Using an interface to the other solver packages in Trilinos, a variety of inexact solvers can be employed for the solution of the local subdomain problems. An extension to multiple subdomains per MPI rank is currently being implemented. Using Kokkos and Kokkos Kernels, FROSch has recently also been ported to GPUs [150] with performance gains for the triangular solve or inexact solves with ILU on GPUs; speedups of up to 4.8× are observed when comparing ILU solve times on CPU and GPU.

A demo/tutorial for FROSch can be found at the GitHub repository [1].

### 3.4 Multigrid Methods: MueLu

MueLu is a flexible and scalable high-performance multigrid solver library. It provides a variety of multigrid algorithms for problems ranging from Poisson-like operators, over elasticity, convection-diffusion, Navier-Stokes, and Maxwell's equations, all the way to multigrid methods for coupled multiphysics systems. Besides its strong focus on aggregation-based algebraic multigrid (AMG)

methods, MueLu comes with specialized capabilities for (semi-)structured grids to perform semi-coarsening along grid lines while forming the coarse operator via a Galerkin product (in contrast to classical geometric multigrid methods). MueLu is easily extensible which makes it a good framework for the research and development of new multigrid preconditioning methods, e.g., different types of strength-of-connection or drop schemes. Its weak and strong scalability up to 131,000 cores of a Cray XC40 and one million cores of a Blue Gene/Q system, even for vector-valued **Partial Differential Equations (PDEs)** on unstructured meshes, have been shown in [93, 138].

MueLu provides several approaches to constructing and solving the multilevel problem:

- *Algebraic Smoothed Aggregation Approach* [144]: The matrix graph is colored to create aggregates (groups) of nodes. These aggregates define a tentative projection operator. A final projection operator is created by applying a smoother to the tentative operator. There are a variety of deterministic and non-deterministic coloring algorithms implemented directly in MueLu and in Kokkos Kernels. For a full description, see [24].
- *Algebraic Multigrid for Maxwell's Equations*: MueLu implements specialized solvers for the solution of curl-curl problems [27]. Scaling results on Haswell, KNL, ARM, and NVIDIA V100 GPUs at full machine scale can be found in [25].
- *Multigrid for Multiphysics Problems*: MueLu implements a toolbox to compile multi-level block preconditioners for block matrices arising from coupled multiphysics problems. Applications range from Navier-Stokes equations over surface-coupling (as in fluid-structure interaction or contact mechanics [147]) to volume-coupled problems (e.g., in magneto-hydro dynamics [105]).
- *Semi-Coarsening Algebraic Multigrid Approach* [142]: Specialized aggregation procedures for three-dimensional meshes generated by extrusion of a two-dimensional unstructured grid allow to first coarsen in the direction of extrusion to reduce the system to a two-dimensional representation and then perform classical aggregation-based AMG for the remaining coarsenings.
- *AMG for (Semi-)Structured Grids*: MueLu has a structured aggregation capability in which the user may specify the coarsening rate used to compute interpolation operators. The coarse operators are then formed via a Galerkin product to avoid remeshing on the coarse levels. This work has been extended to semi-structured grids to leverage structured-grid computational performance also for globally unstructured grids [96].
- *Geometric Multigrid/Matrix-Free Capabilities*: MueLu treats the objects in its hierarchy as operators instead of matrices whenever possible. This enables one to use MueLu in a matrix-free fashion provided the user can supply their own grid transfers and coarse operators.

MueLu has a required dependency on Kokkos, and the main code paths in MueLu (e.g., smoothed aggregation, semi-coarsening,  $p$ -coarsening, structured aggregation) utilize Kokkos and have been shown to scale on device [25]. Some of the earlier algorithms such as energy minimization [126] and the variable degree-of-freedom Laplacian have not been adapted to fully utilize Kokkos yet, however, they are part of ongoing efforts to use Kokkos throughout MueLu completely and to merge serial and device-capable algorithms.

There are a variety of resources to learn more about MueLu. An overview is given on the MueLu website.<sup>12</sup> The MueLu User's Guide [24] summarizes installation instructions and a reference to most of MueLu's configuration parameters. The MueLu Tutorial [97] introduces beginners and experts to various topics in MueLu and shows how to solve or precondition different linear systems

---

<sup>12</sup><https://trilinos.github.io/muelu.html>.

using MueLu. Details on the compatibility of MueLu and its predecessor ML [50, 69] can be found in the MueLu User's Guide [24].

Besides its C++ API, MueLu offers a MATLAB interface, MueMex, to provide access to MueLu's aggregation and solver routines from MATLAB. MueMex allows users to setup and solve arbitrarily many problems with either MueLu as a preconditioner, Belos as a solver, and Tpetra for data structures.

MueLu requires the packages Teuchos, Tpetra, Xpetra, Kokkos, Kokkos Kernels, Amesos2, Ifpack2, and Zoltan2. Additionally, it supports interfaces to abstraction layer packages such as Stratimikos and Thyra through the MueLu adapters library. These interfaces are required to use MueLu with the Teko block preconditioning package. For more details on using MueLu with Teko, see the MueLu examples directory in Trilinos.

### 3.5 Direct Linear Solvers: Amesos2 and ShyLU

Amesos2 [22] serves as the central collection of direct solvers for any type of linear systems: it natively implements the direct solver KLU2 [38], it offers a uniform interface to third-party direct linear solvers such as CHOLMOD, MUMPS, Pardiso\_MKL, SuperLU, SuperLU\_MT, SuperLU\_Dist, and STRUMPACK, and it provides an interface to the on-node sparse direct solvers implemented in ShyLU.

ShyLU [139] provides an implementation of three solvers: Basker, Tacho, and a distributed-memory solver based on the Schur complement method, which all have been developed for specific target applications.

Basker [31] is a shared-memory sparse direct solver based on LU factorization for the problems that have the **Block Triangular Form (BTF)** typically seen in circuit simulation applications. Basker uses this structure to factor and solve the diagonal blocks in parallel. The larger diagonal blocks can themselves be factored in parallel by discovering the parallelism available using a nested-dissection reordering. Basker focuses on exploiting thread-parallelism on the multi-core CPU architectures. Amesos2 also has a templated implementation of a sequential KLU solver called KLU2, which also exploits the BTF structure.

Tacho is a shared-memory sparse direct solver that exploits supernodal block structure commonly found in sparse direct factorizations of matrices from mechanics applications. Tacho exploits this supernodal structure for both factorization and triangular-solve phases. It is based on Kokkos. Originally, Tacho implemented a task-parallel Cholesky factorization of sparse symmetric positive definite matrices [83]. However, to improve its portability, it has been extended to compute the sparse factorization based on level-set scheduling. Moreover, its functionality has been extended to the computation of an LDLT factorization of symmetric indefinite matrices, as well as the computation of an LU factorization of general matrices with a symmetric sparsity pattern.

In addition to their stand-alone use, the aforementioned node-level solvers may be used as the local solvers for domain decomposition preconditioners (Ifpack2 or FROSch) or as the coarse solvers for multilevel preconditioners (MueLu or FROSch).

ShyLU also implements a distributed-memory linear solver based on the Schur complement method [139]. This hybrid approach combines direct and iterative solvers: each subdomain problem is solved in parallel using a direct solver, while the Schur complement is handled iteratively. The preconditioner for the Schur complement solver is computed using a probing approach or using a threshold-based dropping strategy. It has been developed to address the requirements of circuit simulation applications. The solver is also hybrid in terms of parallel computing, utilizing MPI+threads. However, it is not designed for GPU architectures, and therefore, it has not been implemented using Kokkos.

### 3.6 Physics Block Operators and Preconditioners: Teko

Many multi-physics problems are conveniently described by a block-structured matrix which allows easy access to the underlying matrices representing single-field physics or coupling between different physical fields. The Teko library [37] provides interfaces for operators and preconditioners that are constructed from such large physics-based sub-blocks. These sub-blocks are Thyra operators which are themselves often implemented using Tpetra matrices.

Teko provides generic block preconditioning strategies such as  $2 \times 2$  block-LU factorizations as well as block versions of Jacobi- and Gauss-Seidel-type methods for any number of blocks, as well as commonly used approximate inverse strategies for the Navier-Stokes equation such as SIMPLEC, LSC, and PCD preconditioners [36]. The Block-LU class offers the possibility to implement user-specific inverse approximation strategies, for example using a sparse approximate inverse to construct the Schur complement [46]. More complicated multilevel hierarchies of block solvers can be generated via `Teuchos::ParameterList` objects. Block preconditioners for first-order formulations of Maxwell's equations and Darcy flow are implemented in the Panzer package. Teko's solvers can be registered in the Stratimikos interface for usage entirely driven by `Teuchos::ParameterList` objects.

### 3.7 Eigensolvers: Anasazi

Anasazi [13] is an extensible and interoperable framework for large-scale eigenvalue algorithms. This framework provides a generic interface to a collection of algorithms that are built upon abstract interfaces that facilitate code reuse and extensibility. Similar to Belos, Anasazi decouples the algorithms from the implementation of the underlying linear algebra objects using traits mechanisms. Concrete linear algebra adapters are provided for Tpetra and Thyra, while users can also implement their own interfaces to leverage any existing matrix and vector representations. Any libraries that understand Tpetra and Thyra matrices and vectors, like Belos and Ifpack2, may also be used in conjunction with Anasazi. The suite of eigensolvers provided by Anasazi includes locally optimal block preconditioned conjugate gradient, block Davidson, Riemannian Trust-Region, and block Krylov-Schur methods. Recently, a family of trace minimization (TraceMin) methods and a generalized Davidson method have been added to the suite of eigensolvers in Anasazi.

### 3.8 Unified Solver Interface: Stratimikos

The Stratimikos package offers a cohesive interface to various linear solvers and preconditioners within Trilinos, including Amesos2, Belos, FROSch, Ifpack2, MueLu, and Teko. It requires that the matrix, right-hand side, and the solution vectors all support the Thyra interface. Thyra also provides wrappers for Tpetra linear algebra. Users can specify solver and preconditioner parameters via a `Teuchos::ParameterList`, which can be easily populated from an XML file. Moreover, the package allows for the easy addition of new solver and preconditioner factories via adapters, with examples available in the `packages/stratimikos/adapters` directory. Stratimikos does not support eigenvalue problems. Anasazi has, however, a Thyra interface and can therefore be used with different linear algebra backends.

## 4 Nonlinear Solvers and Analysis Tools

The packages included in this section provide the top-level algorithms for computational simulations and design studies. These include nonlinear solvers, bifurcation tracking, stability analysis, parameter continuation, optimization, and uncertainty quantification. A common theme of this collection is the philosophy of "analysis beyond simulation," which aims to automate many computational tasks that are often performed by application code users by trial-and-error or repeated

simulation. Tasks that can be automated include performing parameter studies, sensitivity analysis, calibration, optimization, and the task of locating instabilities. Additional utilities for nonlinear analysis include abstraction layers and interfaces for application callbacks, as well as AD tools that can provide the derivatives critical to analysis algorithms.

#### 4.1 Nonlinear Solvers: NOX, Library of Continuation Algorithms (LOCA)

NOX provides robust and efficient algorithms for solving systems of nonlinear equations. NOX implements a number of Newton-based globalization techniques including line search [113], trust region [113, 114], and homotopy algorithms [33]. Additionally, it provides lower- and higher-order models including Broyden, Anderson acceleration [145], and tensor methods [10]. The algorithms have been designed for large-scale parallel inexact linear solvers using Krylov methods and support Jacobian-free Newton-Krylov variants. The library interacts with application codes through the Thyra model evaluator interface. NOX provides linear algebra abstractions for applications to use custom implementations with the algorithms. All algorithms are implemented under base-class abstractions for the solver, direction, and line search objects. Users can provide derived implementations for each of these base classes to insert new algorithms into the library. NOX provides various stopping criteria, including absolute, relative, and weighted root mean square norms, as well as stagnation and NaN detection. Applications can build custom stopping criteria within a tree-based logical structure and provide additional criteria through the `StatusTest` abstraction. NOX also provides capabilities for continuation and stability analysis through the subpackage LOCA.

LOCA [127] provides techniques for computing families of solutions to nonlinear equations as well as methods for investigating their stability when these nonlinear equations define equilibria of dynamical systems. It builds on the NOX nonlinear solver package to track solutions to sets of nonlinear equations as a function of one or more parameters (continuation). Given an interface to NOX that defines the nonlinear equations, users only need to provide a method for setting the parameter values being varied. LOCA provides several continuation methods, including pseudo-arclength continuation which allows tracking solution curves around turning points/folds. Furthermore, LOCA has hooks to call the Anasazi eigensolver package to estimate leading eigenvalues of the linearization at each point along the continuation curve for linear stability analysis, including various spectral transformations highlighting eigenvalues in different regimes (e.g., largest magnitude, largest real). Finally, LOCA implements equations augmenting the original nonlinear equations to locate and track bifurcation points where linear stability changes (e.g., turning point, pitchfork, and Hopf bifurcations) as a function of additional parameters. Examples of using LOCA for stability and bifurcation analysis include stability of a differentially heated cavity [130], design of chemical vapor deposition reactors [111], flow instabilities in counterflowing jets [112], and ignition in laminar diffusion flames [132].

#### 4.2 Numerical Optimization: ROL

ROL [2, 77] is a package for numerical optimization. ROL brings an extensive collection of state-of-the-art optimization algorithms to virtually any application. Its programming interface supports modern computational hardware, including heterogeneous multi-core systems. ROL has been used with great success for optimal control, optimal design, inverse problems, image processing, and mesh optimization. Application areas include geophysics, climate science [115], structural dynamics [8, 32], fluid dynamics [7], electromagnetics, quantum computing, hypersonics, and geospatial imaging [88].

The key design, performance, and algorithmic features of ROL can be summarized as follows:

- *Vector Abstractions and Matrix-Free Interface for Universal Applicability*: Similar to other Trilinos packages that implement ANAs, ROL’s design is centered around an abstract linear algebra interface, through the `ROL::Vector` class, which enables the use of any ROL algorithm with any data type, such as the C++ `std::vector`, MPI-parallel data structures (e.g., Tpetra, PETSc, HYPRE vectors), GPU-supporting data structures (e.g., Kokkos, ArrayFire [149]), and so on. In contrast to other vector implementations, ROL’s abstract linear algebra interface also enables the definition of custom vector dot products with duality pairings, which are critical for performance in large-scale optimization with differential equations. Finally, all ROL algorithms are matrix-free and rely on user-defined applications of linear and nonlinear operators. ROL’s vector class design is inspired by the Hilbert Class Library [54], the Rice Vector Library [107], and the RTop framework [21], while ROL’s operator interfaces for simulation-based optimization, known as the SimOpt middleware, are inspired by [57].
- *Modern Algorithms for Unconstrained and Constrained Optimization*: ROL features a large collection of well-established algorithms for smooth optimization as well as several novel algorithms developed by the ROL team. ROL categorizes optimization problems into four types based on the presence of certain constraints: Type U (unconstrained), Type B (bound constraints), Type E (equality constraints), and Type G (general constraints). Additionally, each problem type can include linear constraints, which ROL can reduce or eliminate algorithmically. The novel algorithms include augmented Lagrangian methods for infinite-dimensional problems with general constraints [6] and matrix-free trust-region methods for convex-constrained optimization [84].
- *Easy-to-Use Methods for Stochastic and Risk-Aware Optimization*: ROL contains specialized interfaces and algorithms for stochastic optimization. Optimization problems with stochastic and uncertain problem data pose challenges in problem formulation and efficient numerical solution methods. ROL’s interface for stochastic optimization enables the use of numerous modern risk measures, probabilistic functions, and robust problem formulations. To approximate stochastic problems [135], ROL includes sample-based approaches such as sample average approximation, stochastic approximation, and adaptive sparse-grid quadrature [85, 86]. ROL also includes specialized algorithms such as progressive hedging [124] and the primal-dual risk minimization algorithm [89].
- *Fast and Robust Algorithms for Nonsmooth Optimization*: ROL includes specialized algorithms to minimize the sum of a smooth function and a nonsmooth, convex function. This class of problems is ubiquitous in computational science. For example, nonsmooth sparsifying regularizers such as the  $\ell^1$ -norm are common in data science and machine learning [48]. These algorithms can be used to solve convex constrained optimization problems. ROL’s methods for this class of problems are built on the proximity operator and include proximal gradient and Newton methods [23, 78, 104]. As discussed in the subsequent bullet, ROL can also incorporate inexact values and gradients for the smooth objective function term using trust regions [16]. ROL also include an implementation of the trust bundle method [131] for solving general nonconvex, nonsmooth optimization problems.
- *Trust-Region Methods for Inexact and Adaptive Computations*: ROL controls inexact function evaluations and exploits adaptive computations and multi-fidelity models through its trust-region methods. Several of these methods have been pioneered by the ROL team. Their implementations typically require some knowledge of error in the computations (i.e., an *error estimate*). For a comprehensive overview of problem formulations, algorithms with error control, and their ROL implementations, see [87].

- *PDE-OPT Application Development Kit for PDE-Constrained Optimization*: To build and optimize models based on PDEs, ROL offers an application development kit for finite element models, called PDE-OPT. PDE-OPT is based on a modular design with three layers: local finite element computations, which encode the physical equations governing the optimization problem; global computations, which utilize Tpetra data structures; and an interface to ROL, enabling efficient optimization. To solve PDE-constrained optimization problems, users only need to implement the local finite element layer of PDE-OPT (i.e., the evaluations of weak forms on a mesh element); PDE-OPT automates all remaining steps.

In 2025, the development of ROL has moved to its own GitHub repository.<sup>13</sup> The ROL code is snapshotted into the Trilinos repository for convenience of Trilinos users.

### 4.3 Automatic Differentiation: Sacado

Sacado [117, 119, 121] provides forward and reverse-mode operator-overloading-based AD tools within Trilinos. Sacado's forward AD tools have been integrated with Kokkos and have demonstrated good performance on GPU architectures [118]. Sacado, along with its Kokkos integration, provides high-performance derivative capabilities to numerous Office of Science and NNSA (National Nuclear Security Administration) extreme scale applications, including: Albany for solid mechanics and land ice modeling [75, 129]; Charon for semiconductor device modeling [101] and multiphase chemically reacting flows [102]; Drekar for CFD [134, 136], magnetohydrodynamics [133], and plasma physics [35, 98]; Xyce for electronic circuit simulation [11, 79]; and SPARC for hypersonic fluid flows [81].

### 4.4 Uncertainty Quantification: Stokhos

Stokhos [116, 122] provides implementations of two intrusive uncertainty quantification strategies: the intrusive stochastic Galerkin uncertainty quantification method [51, 52] and the embedded ensemble propagation method [120].

Stokhos' implementation of the intrusive stochastic Galerkin uncertainty quantification method allows stochastic projections to be computed, such as polynomial chaos and generalized polynomial chaos expansions. The implementation includes interfaces for forming the linear and/or nonlinear systems that follow from the stochastic Galerkin projection, as well as linear solver methods that can exploit the block structure of these systems. In particular, stochastic Galerkin methods applied to PDEs employ a tensor product discretization between the spatial and stochastic discretizations that generates a two-level sparsity pattern in the resulting stochastic Galerkin operator. This operator is traditionally organized in a block fashion where the outer sparsity pattern is dictated by the stochastic discretization while each block follows the sparsity pattern of the PDE spatial discretization. To improve performance on architectures such as GPUs, Stokhos commutes this layout to obtain an outer sparsity pattern dictated by the spatial discretization where each block follows the stochastic discretization and leverages Kokkos hierarchical parallelism to implement calculations corresponding to the stochastic discretization in a scalable and efficient manner [116]. The implementation has been used in [34] to efficiently propagate uncertainty in multiphysics systems by reducing the full system with a nonlinear elimination method.

The embedded ensemble propagation consists of propagating a subset of samples gathered into a so-called ensemble through the forward simulation at once. It builds on [109] for automating embedded analysis capabilities. Stokhos defines an ensemble type: a data type that is able to store the values of input, output, and state variables for every sample of an ensemble and can take advantage of single instruction, multiple data operations. This type can then be used within the

<sup>13</sup><https://github.com/sandialabs/rol>.

Tpetra solver stack as an argument for its scalar type template parameter, allowing instantiation of Tpetra data structures on the ensemble type to propagate ensembles through supported linear solvers and preconditioners. This approach results in reduced computation time in four ways: the sample-independent data and computation can be reused for every sample of an ensemble, memory access patterns can be improved, operations on the ensemble type can be vectorized efficiently, and message passing costs can be reduced by sending fewer but larger messages. However, the approach requires solvers and BLAS functions to be aware of the extra dimension associated to the ensemble; for example, GMRES for ensemble types [91] needs to monitor the convergence of the individual samples in order to decide when to terminate based on whether every sample within the ensemble has satisfied the chosen stopping criteria.

#### 4.5 Nonlinear Analysis Tools: Piro

Piro [128] provides a simple and unified interface to nonlinear analysis tools in Trilinos, exposing capabilities for computing sensitivities, performing parameter continuation and bifurcation analysis, and solving constrained optimization problems. Piro wraps many of the common functionalities of Trilinos packages such as Belos, NOX, LOCA, Tempus and ROL, into a simple interface. In particular, Piro implements main driver classes for:

- *Linear/nonlinear solvers and sensitivity analysis* for transient and nontransient problems. As an example, this capability can be used to compute the discrete solution of a PDE depending on a parameter and the sensitivity of a quantity of interest with respect to this parameter. Sensitivities can be computed in a forward or adjoint fashion, the latter being preferable in the presence of high-dimensional parameters. To this end, Piro performs forward and adjoint solves, relying on Belos, NOX and Tempus for the solution of linear, nonlinear and transient problems.
- *Constrained optimization problems* with linear/nonlinear equality constraints: Piro provides tools for transient and nontransient (snapshot) optimizations, featuring gradient-based reduced-space and full-space methods. This capability is used by applications such as Albany to solve large-scale PDE-constrained optimization problems [115]. Piro interfaces with ROL for algorithms to solve constrained optimization problems.
- *Parameter continuation and bifurcation analysis*. The core capabilities for continuation and bifurcation analysis are provided through the package LOCA. The Piro driver integrates these functionalities with linear, nonlinear, and time integration solvers, enhancing flexibility and usability.

These driver classes share a similar interface based on the `Thyra::ModelEvaluator` and the `Teuchos::ParameterList` classes. Applications define the problems to be targeted (e.g., the equations to be solved, the parameters, the quantities of interests) by providing a concrete implementations of the `Thyra::ModelEvaluator`; see [109, 110].

## 5 Discretization Tools

This section describes Trilinos packages that provide tools for spatial and temporal discretization of integro-differential equations. Most discretization efforts in Trilinos have been devoted to implementing tools for mesh-based discretizations of PDEs with a focus on high-order finite elements. A notable exception is the research package Compadre (Section 5.5), which provides tools for meshless approximation of linear operators that can be used for the discretization of differential equations and for data transfer. Trilinos' discretization packages have been adopted by many applications addressing a wide range of physics problems, including solid mechanics, earth system modeling, semiconductor devices modeling, and electro-magnetics. These applications

have taken different approaches in adopting Trilinos mesh-based discretization tools. The less intrusive approach is the adoption of Intrepid2 tools to perform local finite element assembly. In this case, the application has to manage the global assembly, possibly using the DoFManager provided by Panzer, and FECrs matrix and vector structures provided by Tpetra. A more intrusive approach is to additionally use the Phalanx package for managing dependencies of field evaluations in conjunction with the Thyra Model Evaluator and Sacado AD, and possibly Tempus for time integration. This approach is particularly useful when developing solvers for complex multiphysics problems, because it allows easy re-use of computational kernels and automates the computation of Jacobians and sensitivities. The most intrusive approach is to build the application around the Panzer package, which provides all of the above, plus the handling of linear and nonlinear solvers and integrated constrained optimization capabilities. In the following, we neither include a description of the snapshot packages STK and Krino, which provide mesh and level-set tools, nor a description of the Shards package, which provides tools for mesh cell topology.

### 5.1 Local Assembly: Intrepid2

Intrepid2 provides interoperable tools for compatible discretizations of PDEs; it is a performance-portable re-implementation and extension of the legacy Intrepid package [26]. Intrepid2 mainly focuses on local assembly of continuous and discontinuous finite elements. It also provides limited capabilities for finite volume discretization. Intrepid2 works on batches of elements (cells), and provides tools to efficiently compute discretized linear functionals (e.g., right-hand-side vectors) and differential operators (e.g., stiffness matrices) at the element level. Intrepid2 implements compatible finite element spaces of various polynomial orders for  $H(\text{grad})$ ,  $H(\text{curl})$ ,  $H(\text{div})$ , and  $L^2$  function spaces on triangles, quadrilaterals, tetrahedrons, hexahedrons, wedges, and pyramids. It provides both Lagrangian basis functions and hierarchical basis functions [47], and it implements performance optimizations (e.g., sum factorizations) exploiting the underlying structure of the problem (e.g., tensor-product elements or other symmetries). The **Degrees of Freedom (DOF)** of  $H(\text{div})$  and  $H(\text{curl})$  finite elements as well as high-order  $H(\text{grad})$  finite elements depend on the global orientation of edges and faces and Intrepid2 provides orientation tools for matching the DOF on shared edges and faces. It also provides interpolation-based projection tools for projecting functions in  $H(\text{grad})$ ,  $H(\text{curl})$ ,  $H(\text{div})$ , and  $L^2$  to the respective discrete spaces. Intrepid2 implements these capabilities through the following classes:

- **CellTools**: This class provides geometric operations in the reference and physical frames. This includes the computation of tangents and normals to edges and faces in the physical frame, the computation of the Jacobian of the reference-to-physical frame map, and other geometric computations.
- **CubatureFactory**: This class provides quadrature rules (called *cubatures* in Intrepid2) of various degrees of accuracy for approximating integrals over elements and their boundaries.
- **Basis**: This is the base class that provides a common interface for functionalities related to finite element bases. Intrepid2 provides derived classes that implement this common interface for a variety of compatible finite element spaces. Each derived class implements the `getValues()` method that computes the values taken by the basis functions or their derivatives (e.g., the gradient for  $H(\text{grad})$  functions, the curl for  $H(\text{curl})$  functions) at a set of input points. The implementation of `getValues()` can be very different depending on the basis. Specific optimizations are available for tensor-product elements. Additionally, there is a `BasisFamily` class with a convenience method, `getBasis()`, which constructs a basis depending on a template argument specifying the type of basis (e.g., hierarchical or nodal), the cell topology, the function space on which it is defined, and its polynomial degree.

- `OrientationTools`: This class provides methods to orient the basis functions based on the global orientation of edges and faces, determined by the global numbering of the cell vertices. This is achieved by building a linear operator (a permutation for tensor-product elements) that encodes the orientation of a particular cell and applying that operator to the reference basis functions.
- `ProjectionTools`: This class provides methods for interpolation-based projections of a given function into a compatible finite element space or between compatible finite element spaces [39]. The provided projections commute with the corresponding differential operators if the quadrature rules can exactly integrate the functions being projected. As an example, projecting an  $H(\text{grad})$  function into the  $H(\text{grad})$  finite element space and then taking its gradient gives the same result as taking the gradient of the function first and then projecting the gradient into the  $H(\text{curl})$  finite element space.
- `FunctionSpaceTools`: This class provides transformations of fields from the reference to the physical frame and back, computation of measures on edges, faces, and cells, scalar/vector/tensor multiplications and contractions for computing integrals.
- `IntegrationTools`: This class provides integration methods that can take advantage of tensor product structures in basis values, providing mechanisms for performance-portable, *sum-factorized* assembly across  $H(\text{grad})$ ,  $H(\text{curl})$ ,  $H(\text{div})$ , and  $L^2$  function spaces. In the future, we plan to provide similar interfaces to support matrix-free discretizations.

Intrepid2 makes use of Kokkos containers to enable memory layouts that are adapted to the computational platform. Intrepid2 also uses Kokkos for its core computational kernels, enabling threaded execution across a variety of architectures. The data types used by Intrepid2 are templated; it is therefore possible to propagate Sacado types through Intrepid2 to perform AD. Current development of Intrepid2 focuses on providing efficient matrix-free discretizations to enhance efficiency on GPU architectures.

## 5.2 Local Field Evaluation: Phalanx

Phalanx is a local field evaluation library designed for equation assembly in PDE applications. The goal of Phalanx is to decompose a complex problem into a number of simpler problems with managed dependencies to support rapid development and extensibility of PDE codes [103, 109, 110]. The data structures use Kokkos for performance portability. Through the use of template metaprogramming, Phalanx supports arbitrary user defined data types and evaluation types. This feature allows for simple integration of AD tools via operator-overloaded scalar types from the Sacado package (Section 4.3). From a simple definition of equations, quantities such as Jacobians, Jacobian-vector products, Hessians, and parameter sensitivities can be evaluated to machine precision. These quantities can be used by other Trilinos packages for operations including Newton-based nonlinear solves, gradient-based optimization, constraint enforcement, and bifurcation analysis [109, 110].

Phalanx uses a graph-based design to manage data dependencies. The runtime-defined DAG allows for rapid prototyping in a production environment where simple interfaces for analysts, flexible models/data structures, and integration of non-trivial third-party libraries are paramount. Phalanx is used in a number of large-scale parallel codes including Albany [129], Charon [101], Drekar [35, 98, 133], and EMPIRE [25].

In recent years, Phalanx has been extended to provide utilities for performance portability under AD. For example, Phalanx provides tools for building and managing a Kokkos view-of-views on device without the use of unified virtual memory and provides utilities for running virtual functions on device. In the future, these utilities may be separated into a separate package.

### 5.3 Time Integration: Tempus

Tempus (Latin, meaning time as in “tempus fugit” → “time flies”) is Trilinos’ time-integration package for advanced transient analysis. It includes various time integrators and embedded sensitivity analysis for next-generation code architectures. Tempus provides “out-of-the-box” time-integration capabilities, which allows users to quickly and easily incorporate time-integration capabilities in their applications and switch between various time integrators depending on the simulation needs. Additionally, Tempus provides “build-your-own” capabilities, which allows applications to incorporate various Tempus components to augment or replace an application’s transient capabilities. Other capabilities include embedded error analysis, sensitivity analysis, and transient optimization with ROL.

Tempus offers a general infrastructure for the time evolution of solutions through a variety of general integration schemes. Tempus provides time integrators for explicit and implicit methods and for first- and second-order **Ordinary Differential Equations (ODEs)**. It can be used from small systems of equations (e.g., single ODEs for the time evolution of plasticity models and multiple ODEs for coupled chemical reactions) to large-scale transient simulations requiring exascale computing (e.g., flow fields around reentry vehicles and magneto-hydrodynamics).

Tempus has several components that can be used in concert or individually, depending on the needs of the application.

- Integrators are the time-loop structure for time integration and provide several features, e.g., controlling the advancement of the solution, selecting the next time step size and handling the solution output.
- Time steppers are individual methods that advance the solution from one step to the next. A variety of time steppers are available:
  - *Classic one-step methods* (e.g., forward Euler and trapezoidal methods);
  - *Explicit Runge-Kutta (RK) methods* (e.g., RK explicit 4 Stage);
  - *Diagonally Implicit Runge-Kutta (DIRK) methods* (e.g., general tableau DIRK and many specific DIRK/SDIRK methods);
  - *Implicit-Explicit (IMEX) Runge-Kutta methods* (e.g., IMEX RK SSP2, IMEX RK SSP3, and general tableau IMEX RK methods);
  - *Multi-Step methods* (i.e., BDF2);
  - *Second-order ODE methods* (e.g., Leapfrog, Newmark methods, and HHT-Alpha);
  - *Steppers with subSteppers* (e.g., operator-split and subcycling methods);
- Solution history is used to maintain the solution during time step failure, for solution restart/output, for interpolation of the solution between time steps, and to provide the solution for transient adjoint sensitivities.
- Time step control and strategies provide methods to select the time step size based on user input and/or temporal error control (e.g., bounding min/max time-step size, relative/absolute maximum error, and time step adjustments for output and checkpointing).

Additionally, Tempus has several mechanisms which allow users to insert application-specific algorithms into Tempus components (e.g., through observers and creating derived classes).

### 5.4 Finite Element Analysis: Panzer

The Panzer package provides global tools for finite element analysis. The package is targeted for large-scale parallel implicit PDEs using continuous and discontinuous compatible finite elements as well as control volume finite elements. Panzer enables the solution of nonlinear problems by interfacing with several Trilinos linear and nonlinear solvers. It uses the Intrepid2 package for

arbitrary-order finite element bases. It computes derivatives and sensitivities with the Sacado AD package. Panzer achieves performance portability through the Kokkos programming model via Tpetra data structures.

Panzer is designed for multiphysics systems. The assembly engine allows for different equation sets in each element block of the mesh and allows for mixed bases for DOF within each element block (e.g., mixed  $H(\text{div})$  and  $H(\text{curl})$  system for electro-magnetics). The assembly tools can create a single fully coupled Jacobian matrix with all off-block dependencies (as a single `Tpetra::CrMatrix`), or it can create a blocked system, grouping sets of DOFs into separate explicit matrices. This allows for physics-based block preconditioning strategies using the Teko package in Trilinos [30, 37] or MueLu's fully coupled AMG hierarchies for multiphysics systems [105]. The assembly process relies on the Phalanx package for efficient assembly of the multiphysics systems. Panzer additionally can wrap the assembly in a `Thyra::ModelEvaluator`, providing direct interfaces to the linear and nonlinear analysis packages in Trilinos.

Panzer can provide utilities for application codes to build on, or can be used as a high-level application framework. Important capabilities include the following.

- `DOFManager`: Panzer provides a stand-alone DOF manager class in the `dof-mgr` subpackage. Given a list of DOFs and their corresponding basis and element blocks, the DOF manager can provide the mapping from DOFs on the mesh entities to the entries in linear algebra objects such as residual vectors and Jacobian matrices. The DOF manager can return the objects required to build distributed Tpetra maps and graphs for both uniquely owned global indices and ghosted indices used during assembly. It provides the local indexing used during assembly as well. The DOF manager contains a mesh abstraction called a connection manager that provides information about mesh connectivity, global numbering of mesh topological entities, and element block groups. It is designed to support any underlying mesh database, allowing applications to use the DOF manager with any finite element application.
- `STK_Interface`: Panzer contains a concrete implementation of the connection manager API for the STK mesh database package in Trilinos. The `STK_Interface` object wraps an STK mesh database. It provides a simple interface for accessing global indices of the mesh and can be used to read/write associated solution data to the database. It additionally supports SEACAS for writing mesh data to disk. The `STK_Interface` also includes support for periodic boundary conditions. This capability can match topological entities on periodic parallel distributed faces. Once matched, the DOFs are unified to enforce periodicity for the DOF manager. This capability is in the `adapters-stk` subpackage.
- *Linear Object Factory*: Panzer provides a linear object factory and linear object container designed to support parallel distributed assembly. The returned containers hold either the linear algebra objects for a uniquely owned DOF map used for solving the linear system or a ghosted version of the linear objects that can be used for assembly. The containers support export and import operations between the unique and ghosted containers. These are used to simplify the assembly process and abstract the underlying Tpetra linear algebra objects. The linear object factory can also create DOF gather and scatter functors for the corresponding (and possibly blocked) matrices. The gather and scatter operations are specialized on the assembly types, including residuals, Jacobians, and tangents. Support for Hessians will be implemented as needed. The capability is found in the `disc-fe` subpackage.
- *Worksets*: The Panzer assembly process provides workset containers to hold static data that can be reused in finite element computations. For example, when using a static mesh, these containers could hold the Jacobian, Jacobian inverse, and all finite element basis values at the quadrature points. The workset concept breaks the local elements on an MPI process into

smaller blocks of uniform computations that can be dispatched to a CPU or GPU. The workset helps to control total memory use for temporaries in the Phalanx DAG. This capability is part of the disc-fe subpackage.

- *Assembly Kernels*: Panzer provides a number of assembly kernels that are optimized for use with Sacado AD data types. When assembling on GPUs, the assembly kernels were found to be an order of magnitude faster if the AD evaluation was parallelized over the internal derivative dimension [118]. This required using kernels with a Kokkos : : TeamPolicy for finite element assembly, where the lowest level of parallelism is used internally by the Sacado AD type. When building Panzer for GPUs, the DFad hierarchic parallelism flag in Sacado should be enabled for kernel performance. The kernels can be found in the disc-fe subpackage.
- Panzer provides an example miniapp for implicit electro-magnetics that is used for benchmarking linear solver performance and for acceptance testing of new HPC systems. It demonstrates an  $H(\text{curl})$ - $H(\text{div})$  formulation for the electric and magnetic fields. It can be found in the MiniEM subpackage.

The Panzer package is intended to provide both low- and high-level tools for implicit finite elements discretizations. The high-level tools aggregate many Trilinos discretization and solver packages. While the high-level tools can be used as a rapid prototyping environment, the front-end is fairly complex to set up, as opposed to true rapid prototyping frameworks such as deal.ii [9], FEniCSx [17], or Firedrake [56]. Panzer is not user friendly in this regard, however, the examples and miniapps are a good starting point and can be quickly adapted to other physics applications. A number of performance portable applications are using Panzer tools at different levels of adoption; these include Albany [129], Charon [101], Drekar [35, 98, 133], and EMPIRE [25].

## 5.5 Approximation of Linear Operators: Compadre

The Compadre package provides tools for the approximation of linear operators (including point evaluation and derivatives), given the location of samples of a function over an unstructured cloud of points. The resulting stencils, when applied directly to samples of the function at these locations, provide an approximation of the linear operator acting on the function at the point(s) queried. This is useful for meshed and meshless data transfer applications (remap). Samples of the function at the specified locations can also be viewed as unknowns, in which case the solution returned by Compadre can be used as a stencil for meshless discretization of PDEs [90].

The package uses **Generalized Moving Least Squares (GMLS)** for approximating functionals. GMLS generalizes classic **Moving Least Squares (MLS)** in the sense that selecting point evaluations for the sampling functionals and the target operator in GMLS provides a traditional MLS reconstruction. However, GMLS offers significantly greater flexibility than MLS by allowing the evaluation of the target function at points to be replaced with *sampling* functionals of a function (e.g., integrals of the function over local domains), and instead of approximating the target function, to approximate a target *operator* (e.g., gradient, curl, divergence) of a function. A detailed description of GMLS can be found in [99, 146].

This increased flexibility enables more advanced and unconventional remapping approaches. For instance, it is possible to use an average vector normal integral over edges (Raviart-Thomas type representation) or a cell-average integral (finite-volume representation) as the sampling functionals or the target operator [95]. Compadre supports full space reconstruction in 1D, 2D, and 3D, and it also supports select sampling functionals and target operators on manifolds [55].

Compadre’s stencil generation involves independent problems to be solved in parallel at the team level with loops over the thread and vector level within each problem. This hierarchical parallelism

is achieved with performance portability by using the Kokkos programming model and leveraging the batched QR with pivoting algorithm implemented in Kokkos Kernels.

Future plans include the implementation of other meshless methods like locally-supported radial basis functions.

## 6 Trilinos Framework

The Trilinos Framework product area provides supporting infrastructure for Trilinos users and developers. This infrastructure is largely not what Trilinos is well-known for but is essential for supporting all Trilinos capabilities.

### 6.1 Build and Test Infrastructure

The primary focus of the Trilinos Framework product relates to building, testing, and releases of Trilinos. This effort includes setting up and maintaining the general infrastructure for structured building of subsets of packages, automated testing, and maintaining and adding to the various testing configurations.

The Trilinos Framework infrastructure is built on top of the TriBITS [20], which is built on top of the open-source tools CMake and CTest.<sup>14</sup> The TriBITS framework allows building arbitrary subgraphs of dependent (Trilinos) CMake packages in one or more individual aggregated CMake projects (in any arrangement desired). Each Trilinos/TriBITS package lists its direct (required and optional) dependent upstream packages, thus forming a package dependency graph. The TriBITS framework uses this package dependency graph to automatically determine what indirect dependent internal packages must be enabled and processed (and built) and what external packages must be found. TriBITS then orchestrates the processing of all of the required CMake code to find the needed external packages and configure and build (and optionally test and install) the selected set of internal packages. This allows a large number of (Trilinos) CMake packages to be configured, built, and tested in a flexible and efficient manner. In addition, TriBITS provides support for a number of advanced features that are not available in raw CMake/CTest including: eliminating a large amount of boiler-plate CMake code and avoiding common mistakes; enabling and testing all downstream packages given a set of enabled i.e., modified) upstream packages; managing the enabling and disabling of tests based on various criteria; producing build and test results submitted to a CDash site on a package-by-package basis; producing reduced source tarballs for only a desired subset of enabled packages. As of version 14.4 of Trilinos, TriBITS and Trilinos have been updated to allow packages using raw CMake to be integrated with just a few well-defined integration requirements. The TriBITS framework has allowed Trilinos to scalably grow in the number of packages and the complexity without undue burden on individual Trilinos developers and users. An always up-to-date overview of the configurations used for automated testing of pull requests against Trilinos can be found in the source code repository.<sup>15</sup>

### 6.2 Documentation Infrastructure

The Trilinos Framework product is also responsible for maintaining the general infrastructure for the Trilinos website,<sup>16</sup> wiki,<sup>17</sup> and Doxygen documentation.<sup>18</sup> It owns some of the documented workflows on these sites, such as the process for reproducing testing failures, but the Trilinos package developers own all of the documentation associated with their associated packages.

<sup>14</sup><https://cmake.org>.

<sup>15</sup><https://github.com/trilinos/Trilinos/blob/master/.github/workflows/AT2.yml>.

<sup>16</sup><https://trilinos.github.io>.

<sup>17</sup><https://github.com/trilinos/Trilinos/wiki>.

<sup>18</sup><https://trilinos.github.io/documentation.html>.

### 6.3 Python Wrappers: PyTrilinos2

PyTrilinos2 is a set of automatically generated Python wrappers for selected Trilinos packages including Tpetra, Teuchos and Thyra, and for exposing solver capabilities from Amesos2, Belos, Ifpack2, and MueLu through Stratimikos. In the future, the list of wrapped packages will be enlarged to provide users with more features and to enable efficient prototyping of new algorithms for developers.

## 7 Trilinos Community

### 7.1 Contributing to Trilinos

Contributions to Trilinos can be offered through the standard GitHub pull request model. Proposed code changes are required to pass a set of tests as well as review and approval prior to be merged. Detailed instructions can be found in the contributing guidelines<sup>19</sup> in the source code repository. Questions regarding the inclusion of new packages should be discussed with the Trilinos leadership team.<sup>20</sup>

### 7.2 Platforms for Exchange among Users or Developers

The community of Trilinos users and developers operates several forums for exchange and discussion. Technical discussions about the source code and its development happen within the Trilinos GitHub repository.<sup>21</sup> The #trilinos channel within the Kokkos slack workspace<sup>22</sup> provides a quick and accessible forum to ask questions.

For in-person exchange, the *Trilinos User-Developer Group (TUG) Meeting* takes place at Sandia National Laboratories in Albuquerque every year. At TUG, all Trilinos users and developers can gather to inform themselves about recent progress and advances, and discuss current challenges and upcoming topics relevant to the entire Trilinos community.

The *European Trilinos User Group (EuroTUG) Meeting series*<sup>23</sup> offers a platform for Europe-based users and developers of the Trilinos project. EuroTUG offers Europe-based researchers and application engineers interested in the Trilinos project easy access to the Trilinos community, minimizing travel burdens by removing the need to travel to Albuquerque. It includes tutorials, user presentations showcasing Trilinos applications, and developer updates on news and ongoing work.

### 7.3 Embedding into Other Software Initiatives

Trilinos is a founding member of the **High Performance Software Foundation (HPSF)**.<sup>24</sup> Established in 2024 under the Linux Foundation, the HPSF aims to build, promote, and advance a portable software stack for HPC by fostering collaboration among industry, academia, and government entities. As an initial technical project within the HPSF, Trilinos contributes its expertise in data structures for parallel computing, linear, nonlinear, and transient solvers, as well as optimization and uncertainty quantification in support of HPSF's mission to enhance the HPC software ecosystem as a whole.

<sup>19</sup><https://github.com/trilinos/Trilinos/blob/master/CONTRIBUTING.md>.

<sup>20</sup><https://trilinos.github.io/support.html>.

<sup>21</sup><https://github.com/trilinos/Trilinos>.

<sup>22</sup>See <https://kokkos.org> for details.

<sup>23</sup><https://eurotug.github.io>.

<sup>24</sup><https://hpsf.io>.

## 8 Concluding Remarks

The evolution of the Trilinos project demonstrates its critical role in scientific and HPC. This update underscores Trilinos' dedication to maintaining relevance amidst the rapid development of HPC software frameworks and hardware architectures, and the demands of increasingly complicated multiscale multiphysics simulation code bases for tackling scientific and engineering problems. By adopting the Kokkos ecosystem, expanding package functionalities, and embracing a modular structure, Trilinos ensures both performance portability and adaptability to existing and emerging application domains.

The integration of data structures for distributed-memory paradigms, advanced linear and nonlinear solvers, discretization technology, and optimization tools highlights its versatility and impact across scientific and engineering disciplines. Its community-driven development model fosters innovation while ensuring robust support for users and contributors alike. Moving forward, the library's commitment to collaboration, scalability, and innovation positions it as a cornerstone of next-generation computational frameworks.

As the computational landscape evolves, Trilinos remains poised to address emerging challenges, leveraging its rich feature set and strong community foundation to drive progress in scientific discovery and technological innovation.

## References

- [1] GitHub. 2024. FROSch Demo. Retrieved December 06, 2024 from <https://github.com/searhein/frosch-demo>
- [2] Sandia National Laboratories. 2024. Rapid Optimization Library (ROL). Retrieved December 06, 2024 from <https://rol.sandia.gov>
- [3] Ahmad Abdelfattah, Timothy Costa, Jack Dongarra, Mark Gates, Azzam Haidar, Sven Hammarling, Nicholas J. Higham, Jakub Kurzak, Piotr Luszczek, Stanimire Tomov, et al. 2021. A set of batched basic linear algebra subprograms and LAPACK routines. *ACM Transactions on Mathematical Software* 47, 3, Article 21 (June 2021), 23 pages. DOI : <https://doi.org/10.1145/3431921>
- [4] Seher Acer, Erik G. Boman, Christian A. Glusa, and Sivasankaran Rajamanickam. 2021. Sphynx: A parallel multi-GPU graph partitioner for distributed-memory systems. *Parallel Computing* 106 (2021), 102769. DOI : <https://doi.org/10.1016/j.parco.2021.102769>
- [5] Filipe A. C. S. Alves, Alexander Heinlein, and Hadi Hajibeygi. 2024. A computational study of algebraic coarse spaces for two-level overlapping additive Schwarz preconditioners. arXiv:2408.08187. Retrieved from <https://arxiv.org/abs/2408.08187>
- [6] Harbir Antil, Drew P. Kouri, and Denis Ridzal. 2023. ALESQP: An augmented Lagrangian equality-constrained SQP method for optimization with general constraints. *SIAM Journal on Optimization* 33, 1 (2023), 237–266. DOI : <https://doi.org/10.1137/20M1378399>
- [7] Harbir Antil, Drew P. Kouri, Denis Ridzal, David B. Robinson, and Maher Salloum. 2024. Uniform flow in axisymmetric devices through permeability optimization. *Optimization and Engineering* 25, 2 (2024), 669–697. DOI : <https://doi.org/10.1007/s11081-023-09820-0>
- [8] Wilkins Aquino, Gregory Bunting, Scott T. Miller, and Timothy F. Walsh. 2019. A gradient-based optimization approach for the detection of partially connected surfaces using vibration tests. *Computer Methods in Applied Mechanics and Engineering* 345 (2019), 323–335. DOI : <https://doi.org/10.1016/j.cma.2018.11.002>
- [9] Daniel Arndt, Wolfgang Bangerth, Maximilian Bergbauer, Marco Feder, Marc Fehling, Johannes Heinz, Timo Heister, Luca Heltai, Martin Kronbichler, Matthias Maier, et al. 2023. The deal.II library, version 9.5. *Journal of Numerical Mathematics* 31, 3 (2023), 231–246. DOI : <https://doi.org/10.1515/jnma-2023-0089>
- [10] Brett W. Bader. 2005. Tensor-Krylov methods for solving large-scale systems of nonlinear equations. *SIAM Journal on Numerical Analysis* 43, 3 (2005), 1321–1347. DOI : <https://doi.org/10.1137/040607095>
- [11] Chris Baker, Erik Boman, Mike Heroux, Eric Keiter, Siva Rajamanickam, Rich Schiek, and Heidi Thornquist. 2012. Enabling next-generation parallel circuit simulation with Trilinos. In *Euro-Par 2011: Parallel Processing Workshops*. Michael Alexander, Pasqua D'Ambra, Adam Belloum, George Bosilca, Mario Cannataro, Marco Danelutto, Beniamino Di Martino, Michael Gerndt, Emmanuel Jeannot, Raymond Namyst, et al. (Eds.), Springer, Berlin, 315–323. DOI : [https://doi.org/10.1007/978-3-642-29737-3\\_36](https://doi.org/10.1007/978-3-642-29737-3_36)
- [12] C. G. Baker and M. A. Heroux. 2012. Tpetra, and the use of generic programming in scientific computing. *Scientific Programming* 20, 2 (2012), 115–128. DOI : <https://doi.org/10.3233/SPR-2012-0349>

- [13] C. G. Baker, U. L. Hetmaniuk, R. B. Lehoucq, and H. K. Thornquist. 2009. Anasazi software for the numerical solution of large-scale eigenvalue problems. *ACM Transactions on Mathematical Software* 36, 3, Article 13 (July 2009), 23 pages. DOI : <https://doi.org/10.1145/1527286.1527287>
- [14] Satish Balay, Shrirang Abhyankar, Mark F. Adams, Jed Brown, Peter Brune, Kris Buschelman, Lisandro Dalcin, Alp Dener, Victor Eijkhout, William D. Gropp, et al. 2021. PETSc Web Page. Retrieved from <https://www.mcs.anl.gov/petsc>
- [15] Daniel Balzani, Alexander Heinlein, Axel Klawonn, Jascha Knepper, Sharan Nurani Ramesh, Oliver Rheinbach, Lea Saßmannshausen, and Klemens Uhlmann. 2024. A computational framework for pharmaco-mechanical interactions in arterial walls using parallel monolithic domain decomposition methods. *GAMM-Mitteilungen* 47, 1 (2024), e202370002. DOI : <https://doi.org/10.1002/gamm.202370002>
- [16] Robert J. Baraldi and Drew P. Kouri. 2023. A proximal trust-region method for nonsmooth optimization with inexact function and gradient evaluations. *Mathematical Programming* 201, 1–2 (2023), 559–598.
- [17] Igor A. Baratta, Joseph P. Dean, Jørgen S. Dokken, Michal Habera, Jack S. Hale, Chris N. Richardson, Marie E. Rognes, Matthew W. Scroggs, Nathan Sime, and Garth N. Wells. 2023. DOLFINx: The next generation FEniCS problem solving environment. Zenodo. DOI : <https://doi.org/10.5281/zenodo.10447666>
- [18] Roscoe A. Bartlett. 2007. *Thyra Linear Operators and Vectors*. Technical Report SAND2007-5984. Sandia National Laboratories.
- [19] Roscoe A. Bartlett. 2010. *Teuchos C++ Memory Management Classes, Idioms, and Related Topics, the Complete Reference: A Comprehensive Strategy for Safe and Efficient Memory Management in C++ for High Performance Computing*. Technical Report SAND2023-12265. Sandia National Laboratories. DOI : <https://doi.org/10.2172/992335>
- [20] Roscoe A. Bartlett. 2014. *TriBITS Developers Guide and Reference*. Technical Report CASL-U-2014-0075000b. Oak Ridge National Laboratory.
- [21] Roscoe A. Bartlett, Bart G. Van Bloemen Waanders, and Michael A. Heroux. 2004. Vector reduction/transformation operators. *ACM Transactions on Mathematical Software* 30, 1 (2004), 62–85. DOI : <https://doi.org/10.1145/974781.974785>
- [22] Eric Bavier, Mark Hoemmen, Sivasankaran Rajamanickam, and Heidi Thornquist. 2012. Amesos2 and Belos: Direct and iterative solvers for large sparse linear systems. *Scientific Programming* 20, 3 (2012), 241–255. DOI : <https://doi.org/10.3233/SPR-2012-0352>
- [23] Amir Beck. 2017. *First-Order Methods in Optimization*. SIAM.
- [24] Luc Berger-Vergiat, Christian A. Glusa, Graham Harper, Jonathan J. Hu, Matthias Mayr, Andrey Prokopenko, Christopher M. Siefert, Raymond S. Tuminaro, and Tobias A. Wiesner. 2023. *MueLu User’s Guide*. Technical Report SAND2023-12265. Sandia National Laboratories, Albuquerque, NM.
- [25] Matthew T. Bettencourt, Dominic A. S. Brown, Keith L. Cartwright, Eric C. Cyr, Christian A. Glusa, Paul T. Lin, Stan G. Moore, Duncan A. O. McGregor, Roger Pawlowski, Edward G. Phillips, et al. 2021. EMPIRE-PIC: A performance portable unstructured particle-in-cell code. *Communications in Computational Physics* 30, 4 (2021), 1232–1268. DOI : <https://doi.org/10.4208/cicp.OA-2020-0261>
- [26] P. Bochev, H. C. Edwards, R. C. Kirby, K. Peterson, and D. Ridzal. 2012. Solving PDEs with intrepid. *Scientific Programming* 20 (2012), 151–180. DOI : <https://doi.org/10.3233/SPR-2012-0340>
- [27] Pavel B. Bochev, Jonathan J. Hu, Christopher M. Siefert, and Raymond S. Tuminaro. 2008. An algebraic multigrid approach based on a compatible gauge reformulation of Maxwell’s equations. *SIAM Journal on Scientific Computing* 31, 1 (2008), 557–583.
- [28] Ian Bogle, Erik G. Boman, Karen Devine, Sivasankaran Rajamanickam, and George M. Slota. 2020. Distributed memory graph coloring algorithms for multiple GPUs. In *2020 IEEE/ACM 10th Workshop on Irregular Applications: Architectures and Algorithms (IA3)*, 54–62. DOI : <https://doi.org/10.1109/IA351965.2020.00013>
- [29] Erik G. Boman, Ümit V. Çatalyürek, Cédric Chevalier, and Karen D. Devine. 2012. The Zoltan and Isorropia parallel toolkits for combinatorial scientific computing: Partitioning, ordering and coloring. *Scientific Programming* 20, 2 (2012), 129–150. DOI : <https://doi.org/10.3233/SPR-2012-0342>
- [30] J. Bonilla, J. N. Shadid, X.-Z. Tang, M. M. Crockatt, P. Ohm, E. G. Phillips, R. P. Pawlowski, S. Conde, and O. Beznosov. 2023. On a fully-implicit VMS-stabilized FE formulation for low Mach number compressible resistive MHD with application to MCF. *Computer Methods in Applied Mechanics and Engineering* 417 (2023), 116359. DOI : <https://doi.org/10.1016/j.cma.2023.116359>
- [31] Joshua D. Booth, Nathan D. Ellingwood, Heidi K. Thornquist, and Sivasankaran Rajamanickam. 2017. Basker: Parallel sparse LU factorization utilizing hierarchical parallelism and data layouts. *Parallel Computing* 68 (2017), 17–31. DOI : <https://doi.org/10.1016/j.parco.2017.06.003>
- [32] Gregory Bunting, Scott T. Miller, Timothy F. Walsh, Clark R. Dohrmann, and Wilkins Aquino. 2021. Novel strategies for modal-based structural material identification. *Mechanical Systems and Signal Processing* 149 (2021), 107295. DOI : <https://doi.org/10.1016/j.ymssp.2020.107295>
- [33] Todd S. Coffey, C. T. Kelley, and David E. Keyes. 2003. Pseudotransient continuation and differential-algebraic equations. *SIAM Journal on Scientific Computing* 25, 2 (2003), 553–569. DOI : <https://doi.org/10.1137/S106482750241044X>

- [34] Paul G. Constantine, Eric T. Phipps, and Timothy M. Wildey. 2014. Efficient uncertainty propagation for network multiphysics systems. *International Journal for Numerical Methods in Engineering* 99, 3 (2014), 183–202.
- [35] Michael M. Crockett, Sibusiso Mabuza, John N. Shadid, Sidafa Conde, Thomas M. Smith, and Roger P. Pawlowski. 2022. An implicit monolithic AFC stabilization method for the CG finite element discretization of the fully-ionized ideal multifluid electromagnetic plasma system. *Journal of Computational Physics* 464 (2022), 111228. DOI : <https://doi.org/10.1016/j.jcp.2022.111228>
- [36] Eric C. Cyr, John N. Shadid, and Raymond S. Tuminaro. 2012. Stabilization and scalable block preconditioning for the Navier–Stokes equations. *Journal of Computational Physics* 231, 2 (2012), 345–363. DOI : <https://doi.org/10.1016/j.jcp.2011.09.001>
- [37] Eric C. Cyr, John N. Shadid, and Raymond S. Tuminaro. 2016. Teko: A block preconditioning capability with concrete example applications in Navier–Stokes and MHD. *SIAM Journal on Scientific Computing* 38, 5 (2016), S307–S331. DOI : <https://doi.org/10.1137/15M1017946>
- [38] Timothy A. Davis and Ekanathan Palamadai Natarajan. 2010. Algorithm 907: KLU, a direct sparse solver for circuit simulation problems. *ACM Transactions on Mathematical Software* 37, 3 (2010), 1–17. DOI : <https://doi.org/10.1145/1824801.1824814>
- [39] Leszek Demkowicz, Jason Kurtz, David Pardo, Maciej Paszynski, Waldemar Rachowicz, and Adam Zdunek. 2007. *Computing with Hp-Adaptive Finite Elements, Vol. 2: Frontiers Three Dimensional Elliptic and Maxwell Problems with Applications* (1st ed.). Chapman & Hall/CRC.
- [40] Mehmet Deveci, Karen D. Devine, Kevin Pedretti, Mark A. Taylor, Sivasankaran Rajamanickam, and Ümit V. Çatalyürek. 2019. Geometric mapping of tasks to processors on parallel computers with mesh or torus networks. *IEEE Transactions on Parallel and Distributed Systems* 30, 9 (2019), 2018–2032. DOI : <https://doi.org/10.1109/TPDS.2019.2900043>
- [41] Mehmet Deveci, Sivasankaran Rajamanickam, Karen D. Devine, and Ümit V. Çatalyürek. 2016. Multi-Jagged: A scalable parallel spatial partitioning algorithm. *IEEE Transactions on Parallel and Distributed Systems* 27, 3 (2016), 803–817. DOI : <https://doi.org/10.1109/TPDS.2015.2412545>
- [42] Mehmet Deveci, Christian Trott, and Sivasankaran Rajamanickam. 2018. Multithreaded sparse matrix-matrix multiplication for many-core and GPU architectures. *Parallel Computing* 78 (2018), 33–46.
- [43] Clark R. Dohrmann, Axel Klawonn, and Olof B. Widlund. 2008. Domain decomposition for less regular subdomains: Overlapping schwarz in two dimensions. *SIAM Journal on Numerical Analysis* 46, 4 (2008), 2153–2168. DOI : <https://doi.org/10.1137/070685841>
- [44] H. Carter Edwards, Daniel Sunderland, Vicki Porter, Chris Amsler, and Sam Mish. 2012. Manycore performance-portability: Kokkos multidimensional array library. *Scientific Programming* 20, 2 (2012), 89–114. DOI : <https://doi.org/10.3233/SPR-2012-0343>
- [45] H. Carter Edwards, Christian R. Trott, and Daniel Sunderland. 2014. Kokkos: Enabling manycore performance portability through polymorphic memory access patterns. *Journal of Parallel and Distributed Computing* 74, 12 (2014), 3202–3216. DOI : <https://doi.org/10.1016/j.jpdc.2014.07.003>
- [46] Max Firmbach, Ivo Steinbrecher, Alexander Popp, and Matthias Mayr. 2024. An approximate block factorization preconditioner for mixed-dimensional beam-solid interaction. *Computer Methods in Applied Mechanics and Engineering* 431 (2024), 117256. DOI : <https://doi.org/10.1016/j.cma.2024.117256>
- [47] Federico Fuentes, Brendan Keith, Leszek Demkowicz, and Sriram Nagaraj. 2015. Orientation embedded high order shape functions for the exact sequence elements of all shapes. *Computers & Mathematics with Applications* 70, 4 (2015), 353–458. DOI : <https://doi.org/10.1016/j.camwa.2015.04.027>
- [48] Jan Niklas Fuhg, Reese Edward Jones, and Nikolaos Bouklas. 2024. Extreme sparsification of physics-augmented neural networks for interpretable model discovery in mechanics. *Computer Methods in Applied Mechanics and Engineering* 426 (2024), 116973.
- [49] Jérémie Gaidamour, Jonathan Hu, Chris Siefert, and Ray Tuminaro. 2012. Design considerations for a flexible multigrid preconditioning library. *Scientific Programming* 20, 3 (2012), 223–239. DOI : <https://doi.org/10.3233/SPR-2012-0344>
- [50] Michael W. Gee, Christopher M. Siefert, Jonathan J. Hu, Ray S. Tuminaro, and Marzio G. Sala. 2006. *ML 5.0 Smoothed Aggregation User’s Guide*. Technical Report SAND2006-2649. Sandia National Laboratories, Albuquerque, NM.
- [51] Roger G. Ghanem and Pol D. Spanos. 1990. Polynomial chaos in stochastic finite elements. *Journal of Applied Mechanics* 57, 1 (1990), 197–202. DOI : <https://doi.org/10.1115/1.2888303>
- [52] Roger G. Ghanem and Pol D. Spanos. 2003. *Stochastic Finite Elements: A Spectral Approach*. Courier Corporation.
- [53] P. Ghysels and W. Vanroose. 2014. Hiding global synchronization latency in the preconditioned conjugate gradient algorithm. *Parallel Computing* 40, 7 (2014), 224–238. DOI : <https://doi.org/10.1016/j.parco.2013.06.0017>
- [54] M. S. Gockenbach and W. W. Symes. 1996. Hilbert class library: A library of abstract C++ classes for optimization and inversion. *Computers & Mathematics with Applications* 32, 6 (1996), 1–13. DOI : [https://doi.org/10.1016/0898-1221\(96\)00140-X](https://doi.org/10.1016/0898-1221(96)00140-X)

- [55] B. J. Gross, N. Trask, P. Kuberry, and P. J. Atzberger. 2020. Meshfree methods on manifolds for hydrodynamic flows on curved surfaces: A Generalized Moving Least-Squares (GMLS) approach. *Journal of Computational Physics* 409 (2020), 109340. DOI : <https://doi.org/10.1016/j.jcp.2020.109340>
- [56] David A. Ham, Paul H. J. Kelly, Lawrence Mitchell, Colin J. Cotter, Robert C. Kirby, Koki Sagiyama, Nacime Bouziani, Thomas J. Gregory, Jack Betteridge, Daniel R. Shapero, et al. 2023. *Firedrake User Manual* (1st ed.). Imperial College London, University of Oxford, Baylor University, and University of Washington. DOI : <https://doi.org/10.25561/104839>
- [57] Matthias Heinkenschloss and Luís N. Vicente. 1999. An interface between optimization and application for the numerical solution of optimal control problems. *ACM Transactions on Mathematical Software* 25, 2 (1999), 157–190. DOI : <https://doi.org/10.1145/317275.317278>
- [58] Alexander Heinlein, Christian Hochmuth, and Axel Klawonn. 2019. Monolithic overlapping schwarz domain decomposition methods with GDSW coarse spaces for incompressible fluid flow problems. *SIAM Journal on Scientific Computing* 41, 4 (2019), C291–C316. DOI : <https://doi.org/10.1137/18M1184047>
- [59] Sebastian Kinnewig, Alexander Heinlein, and Thomas Wick. 2025. Coupling deal.II and FROSch: A sustainable and accessible (O)RAS preconditioner. *ACM Transactions on Mathematical Software* 51, 4, Article 22 (2025), 1–29. DOI: <https://doi.org/10.1145/3766906>
- [60] Alexander Heinlein, Axel Klawonn, Jascha Knepper, and Oliver Rheinbach. 2019. Adaptive GDSW coarse spaces for overlapping schwarz methods in three dimensions. *SIAM Journal on Scientific Computing* 41, 5 (2019), A3045–A3072. DOI : <https://doi.org/10.1137/18M1220613>
- [61] Alexander Heinlein, Axel Klawonn, Jascha Knepper, and Lea Saß Mannshausen. 2026. *Monolithic and Block Overlapping Schwarz Preconditioners for the Incompressible Navier–Stokes Equations*. Electronic Transactions on Numerical Analysis (ETNA).
- [62] Alexander Heinlein, Axel Klawonn, Sivasankaran Rajamanickam, and Oliver Rheinbach. 2020. FROSch: A fast and robust overlapping Schwarz domain decomposition preconditioner based on Xpetra in Trilinos. In *Domain Decomposition Methods in Science and Engineering XXV*. Lecture Notes in Computational Science and Engineering, Vol. 138, Springer, Cham, 176–184. DOI : [https://doi.org/10.1007/978-3-030-56750-7\\_19](https://doi.org/10.1007/978-3-030-56750-7_19)
- [63] Alexander Heinlein, Axel Klawonn, and Oliver Rheinbach. 2016. A parallel implementation of a two-level overlapping schwarz method with energy-minimizing coarse space based on Trilinos. *SIAM Journal on Scientific Computing* 38, 6 (2016), C713–C747. DOI : <https://doi.org/10.1137/16M1062843>
- [64] Alexander Heinlein, Axel Klawonn, Oliver Rheinbach, and Olof B. Widlund. 2018. Improving the parallel performance of overlapping Schwarz methods by using a smaller energy minimizing coarse space. In *Domain Decomposition Methods in Science and Engineering XXIV*. Lecture Notes in Computational Science and Engineering, Vol. 125, Springer, Cham, 383–392. DOI : [https://doi.org/10.1007/978-3-319-93873-8\\_3](https://doi.org/10.1007/978-3-319-93873-8_3)
- [65] Alexander Heinlein, Mauro Perego, and Sivasankaran Rajamanickam. 2022. FROSch preconditioners for land ice simulations of Greenland and Antarctica. *SIAM Journal on Scientific Computing* 44, 2 (2022), B339–B367. DOI : <https://doi.org/10.1137/21M1395260>
- [66] Alexander Heinlein, Oliver Rheinbach, and Friederike Röver. 2022. Parallel scalability of three-level FROSch preconditioners to 220000 cores using the theta supercomputer. *SIAM Journal on Scientific Computing* 45, 3 (Aug. 2022), S173–S198. DOI : <https://doi.org/10.1137/21M1431205>
- [67] M. A. Heroux. 2004. *AztecOO User Guide*. Technical Report SAND2004-3796. Sandia National Laboratories, Albuquerque, NM.
- [68] Michael Allen Heroux. 2011. *Toward Portable Programming of Numerical Linear Algebra on Manycore Nodes*. Technical Report SAND20113556C Sandia National Lab.(SNL-NM), Albuquerque, NM. Retrieved from <https://www.osti.gov/servlets/purl/1109301>
- [69] Michael A. Heroux, Roscoe A. Bartlett, Vicki E. Howle, Robert J. Hoekstra, Jonathan J. Hu, Tamara G. Kolda, Lehoucq, Kevin R. Long, Roger P. Pawlowski, Eric T. Phipps, Andrew G. Salinger, et al. 2005. An overview of the Trilinos project. *ACM Transactions on Mathematical Software* 31, 3 (2005), 397–423. DOI : <https://doi.org/10.1145/1089014.1089021>
- [70] Michael A. Heroux and Damian Rouson. 2012. Special issue on the Trilinos project, part 1 of 2. *Scientific Programming* 20, 2 (2012), 104210. DOI : <https://doi.org/10.3233/SPR-2012-0356>
- [71] Michael A. Heroux and Damian Rouson. 2012. Special issue on the Trilinos project, part 2 of 2. *Scientific Programming* 20, 3 (2012), 513437. DOI : <https://doi.org/10.3233/SPR-2012-0357>
- [72] Michael A. Heroux and James M. Willenbring. 2012. A new overview of the Trilinos project. *Scientific Programming* 20, 2 (2012), 408130. DOI : <https://doi.org/10.3233/SPR-2012-0355>
- [73] Ralf Hiptmair. 1997. Multigrid method for  $H(\text{div})$  in three dimensions. *Electronic Transactions on Numerical Analysis* 6 (1997), 133–152.
- [74] Mark Hoemmen. 2015. Tpetra Project Overview. In *HPCOR Workshop*. Retrieved from <https://www.osti.gov/biblio/1307283>

- [75] M. J. Hoffman, M. Perego, S. F. Price, W. H. Lipscomb, T. Zhang, D. Jacobsen, I. Tezaur, A. G. Salinger, R. Tuminaro, and L. Bertagna. 2018. MPAS-Albany land ice (MALI): A variable-resolution ice sheet model for earth system modeling using Voronoi grids. *Geoscientific Model Development* 11, 9 (2018), 3747–3780. DOI: <https://doi.org/10.5194/gmd-11-3747-2018>
- [76] Victoria E. Howle, Robert C. Kirby, Kevin Long, Brian Brennan, and Kimberly Kennedy. 2012. Playa: High-performance programmable linear algebra. *Scientific Programming* 20, 3 (2012), 257–273. DOI: <https://doi.org/10.3233/SPR-2012-0347>
- [77] A. Javeed, D. P. Kouri, D. Ridzal, and G. Von Winckel. 2022. Get ROL-ing: An introduction to sandia’s rapid optimization library. In *7th International Conference on Continuous Optimization*.
- [78] Christian Kanzow and Theresa Lechner. 2021. Globalized inexact proximal newton-type methods for nonconvex composite functions. *Computational Optimization and Applications* 78 (2021), 377–410.
- [79] Eric R. Keiter, Laura P. Swiler, and Ian Z. Wilcox. 2018. Gradient-enhanced polynomial chaos methods for circuit simulation. In *Scientific Computing in Electrical Engineering*. Ulrich Langer, Wolfgang Amrhein, and Walter Zulehner (Eds.), Springer International Publishing, Cham, 55–68. DOI: [https://doi.org/10.1007/978-3-319-75538-0\\_6](https://doi.org/10.1007/978-3-319-75538-0_6)
- [80] Brian Kelley and Sivasankaran Rajamanickam. 2022. Parallel, portable algorithms for distance-2 maximal independent set and graph coarsening. In *2022 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, IEEE Computer Society, Los Alamitos, CA, 280–290.
- [81] Sarah L. Kieweg, Jaideep Ray, V. G. Weirs, Brian Carnes, Derek Dinzl, Brian Freno, Micah Howard, Eric Phipps, William Rider, and Thomas Smith. 2019. Validation assessment of hypersonic double-cone flow simulations using uncertainty quantification, sensitivity analysis, and validation metrics. In *AIAA Scitech 2019 Forum*. DOI: <https://doi.org/10.2514/6.2019-2278>
- [82] Kyungjoo Kim, Timothy B. Costa, Mehmet Deveci, Andrew M. Bradley, Simon D. Hammond, Murat E. Guney, Sarah Knepper, Shane Story, and Sivasankaran Rajamanickam. 2017. Designing vector-friendly compact BLAS and LAPACK kernels. In *International Conference for High Performance Computing, Networking, Storage and Analysis*. ACM, New York, NY, USA, 1–12.
- [83] K. Kim, H. C. Edwards, and S. Rajamanickam. 2018. Tacho: Memory-scalable task parallel sparse Cholesky factorization. In *2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. DOI: <https://doi.org/10.1109/IPDPSW.2018.00094>
- [84] Drew P. Kouri. 2022. A matrix-free trust-region newton algorithm for convex-constrained optimization. *Optimization Letters* 16, 3 (2022), 983–997. DOI: <https://doi.org/10.1007/s11590-021-01794-1>
- [85] D. P. Kouri, M. Heinkenschloss, D. Ridzal, and B. G. van Bloemen Waanders. 2013. A trust-region algorithm with adaptive stochastic collocation for PDE optimization under uncertainty. *SIAM Journal on Scientific Computing* 35, 4 (2013), A1847–A1879.
- [86] D. P. Kouri, M. Heinkenschloss, D. Ridzal, and B. G. van Bloemen Waanders. 2014. Inexact objective function evaluations in a trust-region algorithm for PDE-constrained optimization under uncertainty. *SIAM Journal on Scientific Computing* 36, 6 (2014), A3011–A3029.
- [87] Drew P. Kouri and Denis Ridzal. 2018. *Inexact Trust-Region Methods for PDE-Constrained Optimization*. Springer, New York, NY, 83–121. DOI: [https://doi.org/10.1007/978-1-4939-8636-1\\_3](https://doi.org/10.1007/978-1-4939-8636-1_3)
- [88] Drew P. Kouri and Eric A. Shields. 2014. Efficient multiframe super-resolution for imagery with lateral shifts. *Applied Optics* 53, 24, (Aug. 2014), F1–F9. DOI: <https://doi.org/10.1364/AO.53.0000F1>
- [89] Drew P. Kouri and Thomas M. Surowiec. 2022. A primal–dual algorithm for risk minimization. *Mathematical Programming* 193, 1 (2022), 337–363.
- [90] Paul Kuberry, Pavel Bochev, Jake Koester, and Nathaniel Trask. 2025. A discontinuous piecewise polynomial generalized moving least squares scheme for robust finite element analysis on arbitrary grids. *Engineering with Computers* 41, 3 (2025), 1435–1554. DOI: <https://doi.org/10.1007/s00366-024-02036-5>
- [91] Kim Liegeois, Romain Boman, Eric T. Phipps, Tobias A. Wiesner, and Maarten Arnst. 2020. GMRES with embedded ensemble propagation for the efficient solution of parametric linear systems in uncertainty quantification of computational models. *Computer Methods in Applied Mechanics and Engineering* 369 (2020), 113188.
- [92] Kim Liegeois, Sivasankaran Rajamanickam, and Luc Berger-Vergiat. 2023. Performance portable batched sparse linear solvers. *IEEE Transactions on Parallel and Distributed Systems* 34, 5 (2023), 1524–1535.
- [93] Paul T. Lin, John N. Shadid, Jonathan J. Hu, Roger P. Pawlowski, and Eric C. Cyr. 2017. Performance of fully-coupled algebraic multigrid preconditioners for large-scale VMS resistive MHD. *Journal of Computational and Applied Mathematics* 344 (2017), 782–793.
- [94] K. Long, Paul T. Boggs, and Bart G. van Bloemen Waanders. 2012. Sundance: High-level software for PDE-constrained optimization. *Scientific Programming* 20, 3 (2012), 293–310. DOI: <https://doi.org/10.3233/SPR-2012-0341>

- [95] V. S. Mahadevan, J. E. Guerra, X. Jiao, P. Kuberry, Y. Li, P. Ullrich, D. Marsico, R. Jacob, P. Bochev, and P. Jones. 2022. Metrics for Intercomparison of Remapping Algorithms (MIRA) protocol applied to Earth system models. *Geoscientific Model Development* 15, 17 (2022), 6601–6635. DOI : <https://doi.org/10.5194/gmd-15-6601-2022>
- [96] Matthias Mayr, Luc Berger-Vergiat, Peter Ohm, and Raymond S. Tuminaro. 2022. Non-invasive multigrid for semi-structured grids. *SIAM Journal on Scientific Computing* 44, 4 (2022), A2734–A2764. DOI : <https://doi.org/10.1137/20M1375413>
- [97] Matthias Mayr, Christian A. Glusa, Graham B. Harper, Jonathan J. Hu, Andrey Prokopenko, and Tobias A. Wiesner. 2025. The MueLu Tutorial. Retrieved from <https://muelu.github.io/>
- [98] S. T. Miller, E. C. Cyr, J. N. Shadid, R. M. J. Kramer, E. G. Phillips, S. Conde, and R. P. Pawlowski. 2019. IMEX and exact sequence discretization of the multi-fluid plasma model. *Journal of Computational Physics* 397 (2019), 108806. DOI : <https://doi.org/10.1016/j.jcp.2019.05.052>
- [99] Davoud Mirzaei, Robert Schaback, and Mehdi Dehghan. 2012. On generalized moving least squares and diffuse derivatives. *IMA Journal of Numerical Analysis* 32, 3 (2012), 983–1000.
- [100] Karla Morris, Damian W. I. Rouson, M. Nicole Lemaster, and Salvatore Filippone. 2012. Exploring capabilities within ForTrilinos by solving the 3D burgers equation. *Scientific Programming* 20, 3 (2012), 275–292. DOI : <https://doi.org/10.3233/SPR-2012-0353>
- [101] Lawrence C. Musson, Gary L. Hennigan, Xujiao Gao, Richard Humphries, Mihai Negoita, and Andy Huang. 2020. *Charon Users Manual: v. 2.1 (revision1)*. Technical Report SAND2020-5266. Sandia National Laboratories.
- [102] Lawrence C. Musson, Roger P. Pawlowski, Andrew G. Salinger, Timothy J. Madden, and Kevin B. Hewett. 2009. Multiphase reacting flow modeling of singlet oxygen generators for chemical oxygen iodine lasers. In *XVII International Symposium on Gas Flow, Chemical Lasers, and High-Power Lasers*. Rui Vilar, Olinda Conde, Marta Fajardo, Luis O. Silva, Margarida Pires, and Andrei Utkin (Eds.), Vol. 7131, International Society for Optics and Photonics, SPIE, 148–155. DOI : <https://doi.org/10.1117/12.816415>
- [103] Patrick K. Notz, Roger P. Pawlowski, and James C. Sutherland. 2012. Graph-based software design for managing complexity and enabling concurrency in multiphysics PDE software. *ACM Transactions on Mathematical Software* 39, 1, Article 1 (Nov. 2012), 21 pages. DOI : <https://doi.org/10.1145/2382585.2382586>
- [104] Peter Ochs, Yunjin Chen, Thomas Brox, and Thomas Pock. 2014. iPiano: Inertial proximal algorithm for nonconvex optimization. *SIAM Journal on Imaging Sciences* 7, 2 (2014), 1388–1419. DOI : <https://doi.org/10.1137/130942954>
- [105] Peter Ohm, Tobias A. Wiesner, Eric C. Cyr, Jonathan J. Hu, John N. Shadid, and Raymond S. Tuminaro. 2022. A monolithic algebraic multigrid framework for multiphysics applications with examples from resistive MHD. *ETNA—Electronic Transactions on Numerical Analysis* 55 (2022), 365–390. DOI : [https://doi.org/10.1553/etna\\_vol55s365](https://doi.org/10.1553/etna_vol55s365)
- [106] Ron A. Oldfield, Gregory D. Sjaardema, Gerald F. Lofstead, II, and Todd Kordenbrock. 2012. Trilinos I/O support (Trios). *Scientific Programming* 20, 2 (2012), 181–196. DOI : <https://doi.org/10.3233/SPR-2012-0345>
- [107] Anthony D. Padula, Shannon D. Scott, and William W. Symes. 2009. A software framework for abstract expression of coordinate-free linear algebra and optimization algorithms. *ACM Transactions on Mathematical Software* 36, 2 (2009), 1–36. DOI : <https://doi.org/10.1145/1499096.1499097>
- [108] Albert E. Parker and Colin Fox. 2012. Sampling Gaussian distributions in Krylov spaces with conjugate gradients. *SIAM Journal on Scientific Computing* 34, 3 (2012), B312–B334.
- [109] Roger P. Pawlowski, Eric T. Phipps, and Andrew G. Salinger. 2012. Automating embedded analysis capabilities and managing software complexity in multiphysics simulation, part I: Template-based generic programming. *Scientific Programming* 20, 2 (2012), 197–219.
- [110] Roger P. Pawlowski, Eric T. Phipps, Andrew G. Salinger, Steven J. Owen, Christopher M. Siefert, and Matthew L. Staten. 2012. Automating embedded analysis capabilities and managing software complexity in multiphysics simulation, part II: Application to partial differential equations. *Scientific Programming* 20, 3 (2012), 327–345.
- [111] Roger P. Pawlowski, Andrew G. Salinger, Louis R. Romero, and John N. Shadid. 2001. Computational design and analysis of MOVPE reactors. *Journal de Physique IV France* 11, PR3 (2001), Pr3-197–Pr3-204. DOI : <https://doi.org/10.1051/jp4:2001325>
- [112] R. P. Pawlowski, A. G. Salinger, J. N. Shadid, and T. J. Mountziaris. 2006. Bifurcation and stability analysis of laminar isothermal counterflowing jets. *Journal of Fluid Mechanics* 551 (2006), 117–139. DOI : <https://doi.org/10.1017/S0022112005008396>
- [113] Roger P. Pawlowski, John N. Shadid, Joseph P. Simonis, and Homer F. Walker. 2006. Globalization techniques for Newton–Krylov methods and applications to the fully coupled solution of the Navier–Stokes equations. *SIAM Review* 48, 4 (2006), 700–721. DOI : <https://doi.org/10.1137/S0036144504443511>
- [114] Roger P. Pawlowski, Joseph P. Simonis, Homer F. Walker, and John N. Shadid. 2008. Inexact newton dogleg methods. *SIAM Journal on Numerical Analysis* 46, 4 (2008), 2112–2132. DOI : <https://doi.org/10.1137/050632166>

- [115] Mauro Perego. 2022. Large-Scale PDE-Constrained Optimization for Ice Sheet Model Initialization. *SIAM News*. Retrieved from <https://sinews.siam.org/Details-Page/large-scale-pde-constrained-optimization-for-ice-sheet-model-initialization>
- [116] Eric Phipps, H. Carter Edwards, Jonathan Hu, and Jakob T. Ostien. 2014. Exploring emerging manycore architectures for uncertainty quantification through embedded stochastic Galerkin methods. *International Journal of Computer Mathematics* 91, 4 (2014), 707–729.
- [117] Eric Phipps and Roger Pawlowski. 2012. Efficient expression templates for operator overloading-based automatic differentiation. In *Recent Advances in Algorithmic Differentiation*. Springer, Berlin, 309–319.
- [118] Eric Phipps, Roger Pawlowski, and Christian Trott. 2022. Automatic differentiation of C++ codes on emerging manycore architectures with Sacado. *ACM Transactions on Mathematical Software* 48, 4 (2022), 1–29.
- [119] Eric T. Phipps, Roscoe A. Bartlett, David M. Gay, and Robert J. Hoekstra. 2008. Large-scale transient sensitivity analysis of a radiation-damaged bipolar junction transistor via automatic differentiation. *Advances in Automatic Differentiation* 64 (2008), 351–362.
- [120] E. Phipps, M. D’Elia, H. C. Edwards, M. Hoemmen, J. Hu, and S. Rajamanickam. 2017. Embedded ensemble propagation for improving performance, portability, and scalability of uncertainty quantification on emerging computational architectures. *SIAM Journal on Scientific Computing* 39, 2 (2017), C162–C193.
- [121] E. T. Phipps and D. M. Gay. 2011. Sacado Automatic Differentiation Package. Retrieved from <http://trilinos.sandia.gov/packages/sacado/>
- [122] Eric T. Phipps and Andrew G. Salinger. 2016. *Embedded Uncertainty Quantification Methods via Stokhos*. Springer International Publishing, Cham, 1–43. DOI: [https://doi.org/10.1007/978-3-319-11259-6\\_55-1](https://doi.org/10.1007/978-3-319-11259-6_55-1)
- [123] Sivasankaran Rajamanickam, Seher Acer, Luc Berger-Vergiat, Vinh Dang, Nathan Ellingwood, Evan Harvey, Brian Kelley, Christian R. Trott, Jeremiah Wilke, and Ichitaro Yamazaki. 2021. Kokkos kernels: Performance portable sparse/dense linear algebra and graph kernels. arXiv:2103.11991. Retrieved from <https://arxiv.org/abs/2103.11991>
- [124] R. T. Rockafellar and Roger J.-B. Wets. 1991. Scenarios and policy aggregation in optimization under uncertainty. *Mathematics of Operations Research* 16, 1 (1991), 119–147.
- [125] Youcef Saad. 1993. A flexible inner-outer preconditioned GMRES algorithm. *SIAM Journal on Scientific Computing* 14, 2 (1993), 461–469. DOI: <https://doi.org/10.1137/0914028>
- [126] Marzio G. Sala and Ray S. Tuminaro. 2008. A new Petrov–Galerkin smoothed aggregation preconditioner for nonsymmetric linear systems. *SIAM Journal on Scientific Computing* 31, 1 (2008), 143–166. DOI: <https://doi.org/10.1137/060659545>
- [127] A. G. Salinger, E. A. Burroughs, R. P. Pawlowski, E. T. Phipps, and L. A. Romero. 2005. Bifurcation tracking algorithms and software for large scale applications. *International Journal of Bifurcation and Chaos* 15, 3 (2005), 1015–1032.
- [128] Andrew Salinger and Eric T. Phipps. 2010. Piro v. 1.0, Version 00. Retrieved from <https://www.osti.gov/biblio/1231283>
- [129] Andrew G. Salinger, Roscoe A. Bartlett, Andrew M. Bradley, Qiushi Chen, Irina P. Demeshko, Xujiao Gao, Glen A. Hansen, Alejandro Mota, Richard P. Muller, Erik Nielsen, et al. 2016. Albany: Using component-based design to develop a flexible, generic multiphysics analysis code. *International Journal for Multiscale Computational Engineering* 14, 4 (2016), 415–438.
- [130] Andrew G. Salinger, Richard B. Lehoucq, Roger P. Pawlowski, and John N. Shadid. 2002. Computational bifurcation and stability studies of the 8:1 thermal cavity problem. *International Journal for Numerical Methods in Fluids* 40, 8 (2002), 1059–1073. DOI: <https://doi.org/10.1002/flid.392>
- [131] Helga Schramm and Jochem Zowe. 1992. A version of the bundle idea for minimizing a nonsmooth function: Conceptual idea, convergence analysis, numerical results. *SIAM Journal on Optimization* 2, 1 (1992), 121–152. DOI: <https://doi.org/10.1137/0802008>
- [132] J. N. Shadid, A. G. Salinger, R. P. Pawlowski, P. T. Lin, G. L. Hennigan, R. S. Tuminaro, and R. B. Lehoucq. 2006. Large-scale stabilized FE computational analysis of nonlinear steady-state transport/reaction systems. *Computer Methods in Applied Mechanics and Engineering* 195, 13 (2006), 1846–1871. DOI: <https://doi.org/10.1016/j.cma.2005.05.047>
- [133] J. N. Shadid, R. P. Pawlowski, E. C. Cyr, R. S. Tuminaro, L. Chacon, and P. D. Weber. 2016. Scalable implicit incompressible resistive MHD with stabilized FE and fully-coupled Newton-Krylov-AMG. *Computer Methods in Applied Mechanics and Engineering* 304 (2016), 1–25. DOI: <https://doi.org/10.1016/j.cma.2016.01.019>
- [134] J. N. Shadid, T. M. Smith, E. C. Cyr, T. M. Wildey, and R. P. Pawlowski. 2016. Stabilized FE simulation of prototype thermal-hydraulics problems with integrated adjoint-based capabilities. *Journal of Computational Physics* 321 (2016), 321–341. DOI: <https://doi.org/10.1016/j.jcp.2016.04.062>
- [135] Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczyński. 2021. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM.
- [136] David Sondak, Thomas M. Smith, Roger P. Pawlowski, Sidafa Conde, and John N. Shadid. 2021. High Rayleigh number variational multiscale large eddy simulations of Rayleigh–Bénard convection. *Mechanics Research Communications* 112 (2021), 103614. DOI: <https://doi.org/10.1016/j.mechrescom.2020.103614>

- [137] William F. Spitz. 2012. PyTrilinos: Recent advances in the python interface to Trilinos. *Scientific Programming* 20, 3 (2012), 311–325. DOI: <https://doi.org/10.3233/SPR-2012-0346>
- [138] S. J. Thomas, S. Ananthan, S. Yellapantula, Jonathan J. Hu, M. Lawson, and Michael A. Sprague. 2019. A comparison of classical and aggregation-based algebraic multigrid preconditioners for high-fidelity simulations of wind turbine incompressible flows. *SIAM Journal on Scientific Computing* 41, 5 (2019), S196–S219.
- [139] H. K. Thornquist and S. Rajamanickam. 2014. A hybrid approach for parallel transistor-level full-chip circuit simulation. In *International Conference on High Performance Computing for Computational Science*. Springer, Cham, 102–111.
- [140] Christian Trott, Luc Berger-Vergiat, David Poliakoff, Sivasankaran Rajamanickam, Damien Lebrun-Grandie, Jonathan Madsen, Nader Al Awar, Milos Gligoric, Galen Shipman, and Geoff Womeldorff. 2021. The Kokkos ecosystem: Comprehensive performance portability for high performance computing. *Computing in Science & Engineering* 23, 5 (2021), 10–18.
- [141] Christian R. Trott, Damien Lebrun-Grandié, Daniel Arndt, Jan Ciesko, Vinh Dang, Nathan Ellingwood, Rahul Kumar Gayatri, Evan Harvey, Daisy S. Hollman, Dan Ibanez, et al. 2022. Kokkos 3: Programming model extensions for the exascale era. *IEEE Transactions on Parallel and Distributed Systems* 33, 4 (2022), 805–817. DOI: <https://doi.org/10.1109/TPDS.2021.3097283>
- [142] Raymond S. Tuminaro, Mauro Perego, Irina Tezaur, Andrew G. Salinger, and S. Price. 2016. A matrix dependent/algebraic multigrid approach for extruded meshes with applications to ice sheet modeling. *SIAM Journal on Scientific Computing* 38, 5 (2016), C504–C532. DOI: <https://doi.org/10.1137/15M1040839>
- [143] S. P. Vanka. 1986. Block-implicit multigrid solution of Navier-Stokes equations in primitive variables. *Journal of Computational Physics* 65, 1 (1986), 138–158.
- [144] Petr Vaněk, Jan Mandel, and Marian Brezina. 1996. Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems. *Computing* 56 (1996), 179–196.
- [145] Homer F. Walker and Peng Ni. 2011. Anderson acceleration for fixed-point iterations. *SIAM Journal on Numerical Analysis* 49, 4 (2011), 1715–1735. DOI: <https://doi.org/10.1137/10078356X>
- [146] Holger Wendland. 2004. *Data Approximation*, Vol. 17. Cambridge University Press.
- [147] Tobias A. Wiesner, Matthias Mayr, Alexander Popp, Michael W. Gee, and Wolfgang A. Wall. 2021. Algebraic multigrid methods for saddle point systems arising from mortar contact formulations. *International Journal for Numerical Methods in Engineering* 122, 15 (2021), 3749–3779. DOI: <https://doi.org/10.1002/nme.6680>
- [148] Michael M. Wolf, Mehmet Deveci, Jonathan W. Berry, Simon D. Hammond, and Sivasankaran Rajamanickam. 2017. Fast linear algebra-based triangle counting with KokkosKernels. In *2017 IEEE High Performance Extreme Computing Conference (HPEC)*. IEEE, Los Alamitos, CA, USA, 1–7.
- [149] Pavan Yalamanchili, Umar Arshad, Zakiuddin Mohammed, Pradeep Garigipati, Peter Entschew, Brian Kloppenborg, James Malcolm, and John Melonakos. 2015. ArrayFire—A high performance software library for parallel computing with an easy-to-use API. Retrieved from <https://github.com/arrayfire/arrayfire>
- [150] Ichitaro Yamazaki, Alexander Heinlein, and Sivasankaran Rajamanickam. 2023. An experimental study of two-level Schwarz domain-decomposition preconditioners on GPUs. *IEEE Computer Society*, 680–689. DOI: <https://doi.org/10.1109/IPDPS54959.2023.00073>

Received 11 March 2025; revised 1 November 2025; accepted 1 March 2026