Batched Linear Algebra for High-Order Finite Elements

SC24
Atlanta, GA | hpc creates.

Tzanio Kolev
Lawrence Livermore National Laboratory

LLNL-PRES-863290
This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Lawrence Livermore National Laboratory

Finite Element Simulations on Unstructured Grids

The discretization engine at the core of many HPC applications

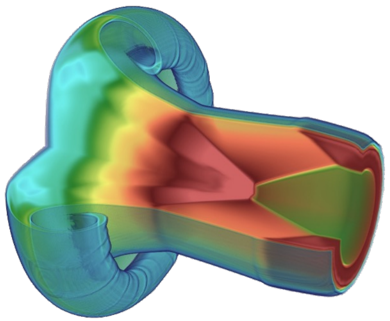

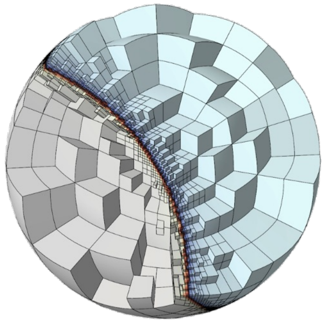

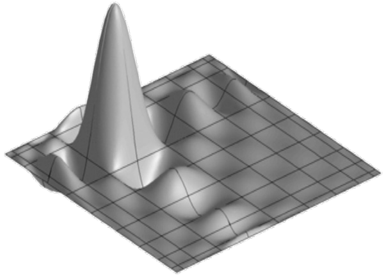

A x = b

Continuous PDE → **Meshing** → **Finite Elements** → **Solvers**

$$-\Delta u = f$$

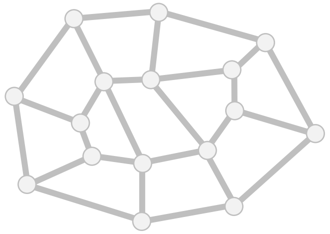

$$\Omega_h$$

$$u_h(x) = \sum_i x_i \varphi_i(x)$$

$$\int_{\Omega_h} \nabla u_h \nabla \varphi_j = \int_{\Omega_h} f \varphi_j$$

$$A_{ij} = \int_{\Omega_h} \nabla \varphi_j \nabla \varphi_i$$

$$b_j = \int_{\Omega_h} f \varphi_i$$

Hierarchical Data Structures Expose Parallelism

Finite elements at global, device/rank, element, and quadrature point level

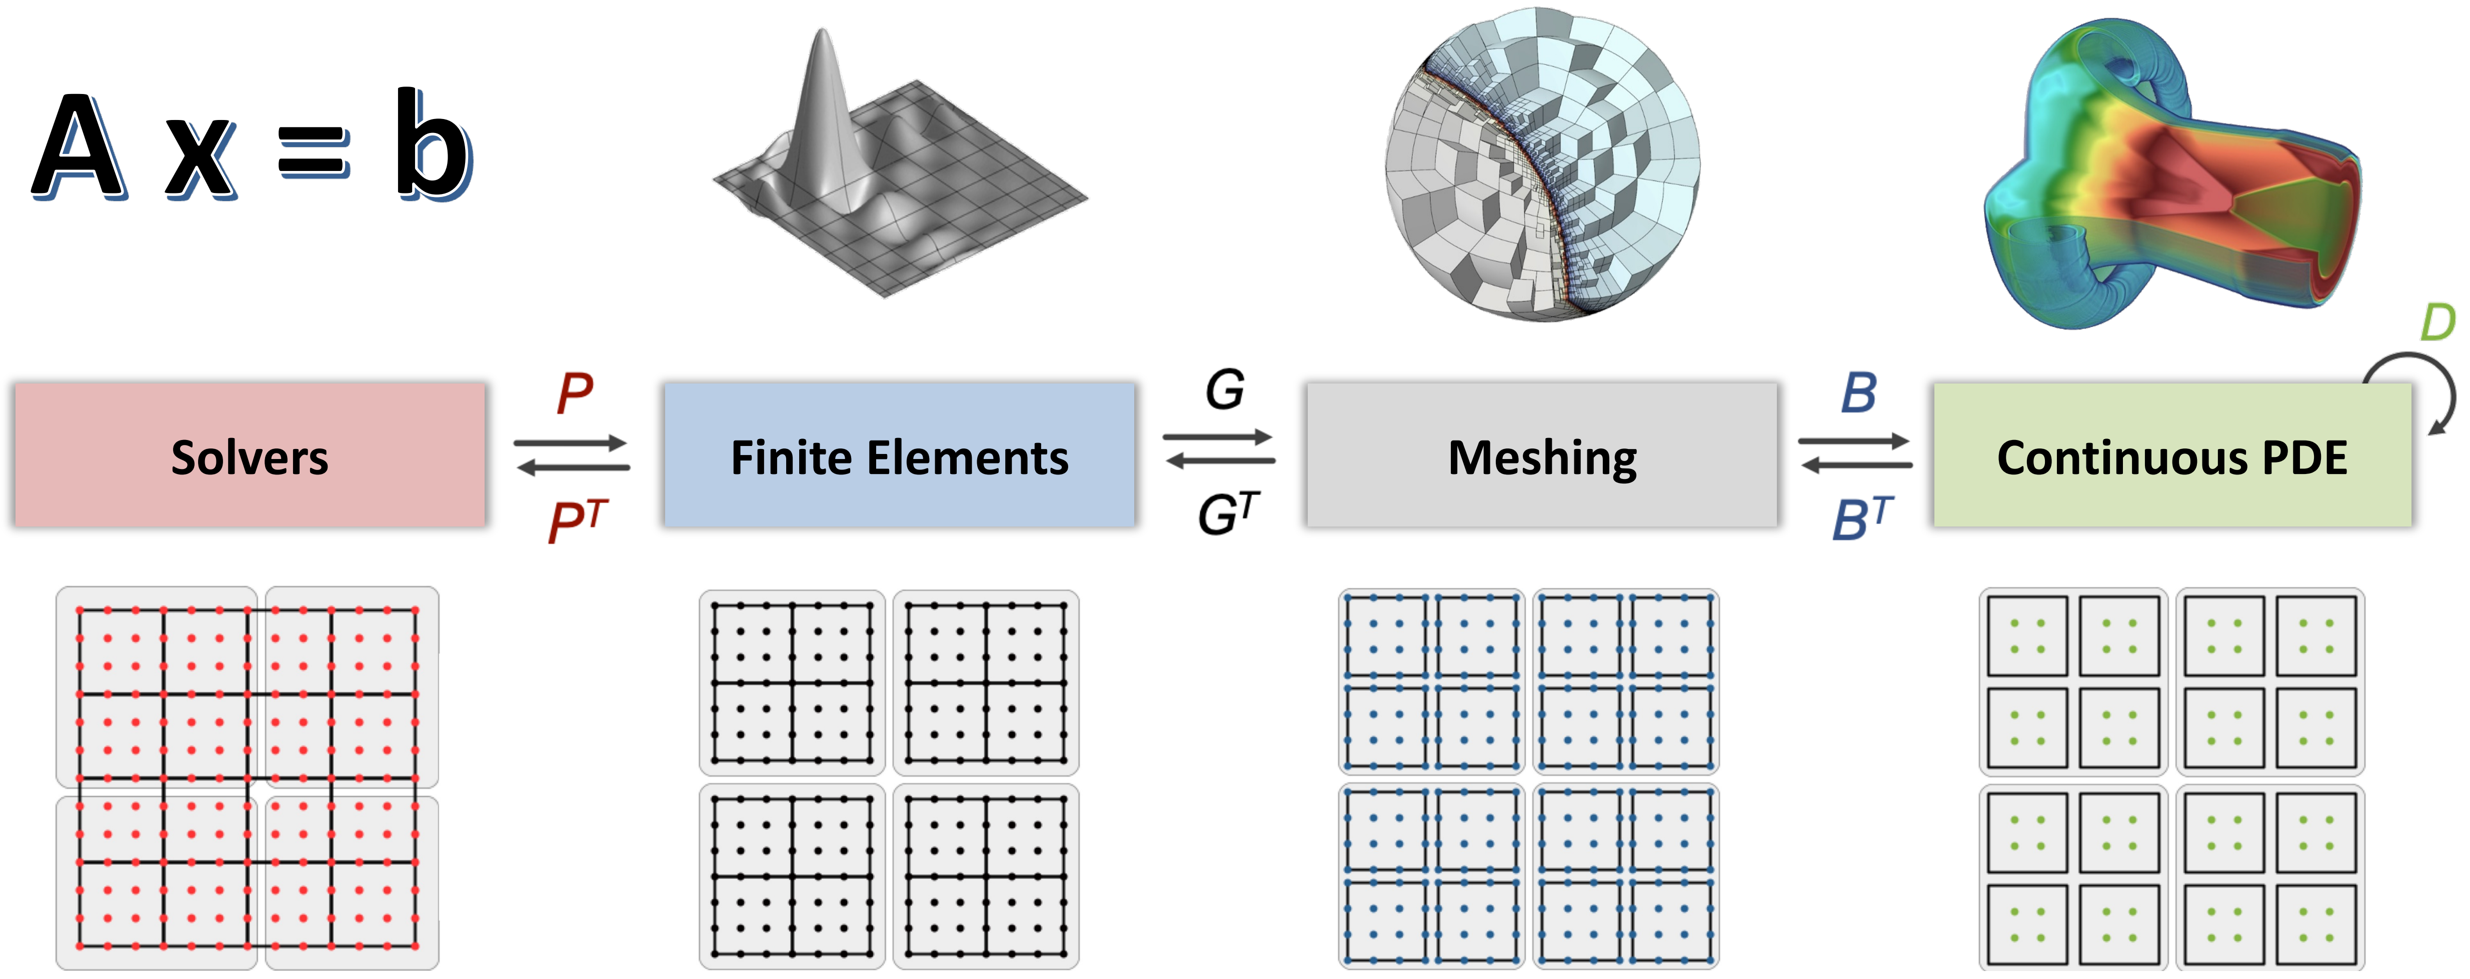

FEM Decomposition the Foundation for HPC Algorithms

Decompose A into parallel, mesh, finite element, and geometry/physics components

$$A = P^T G^T B^T D B G P$$

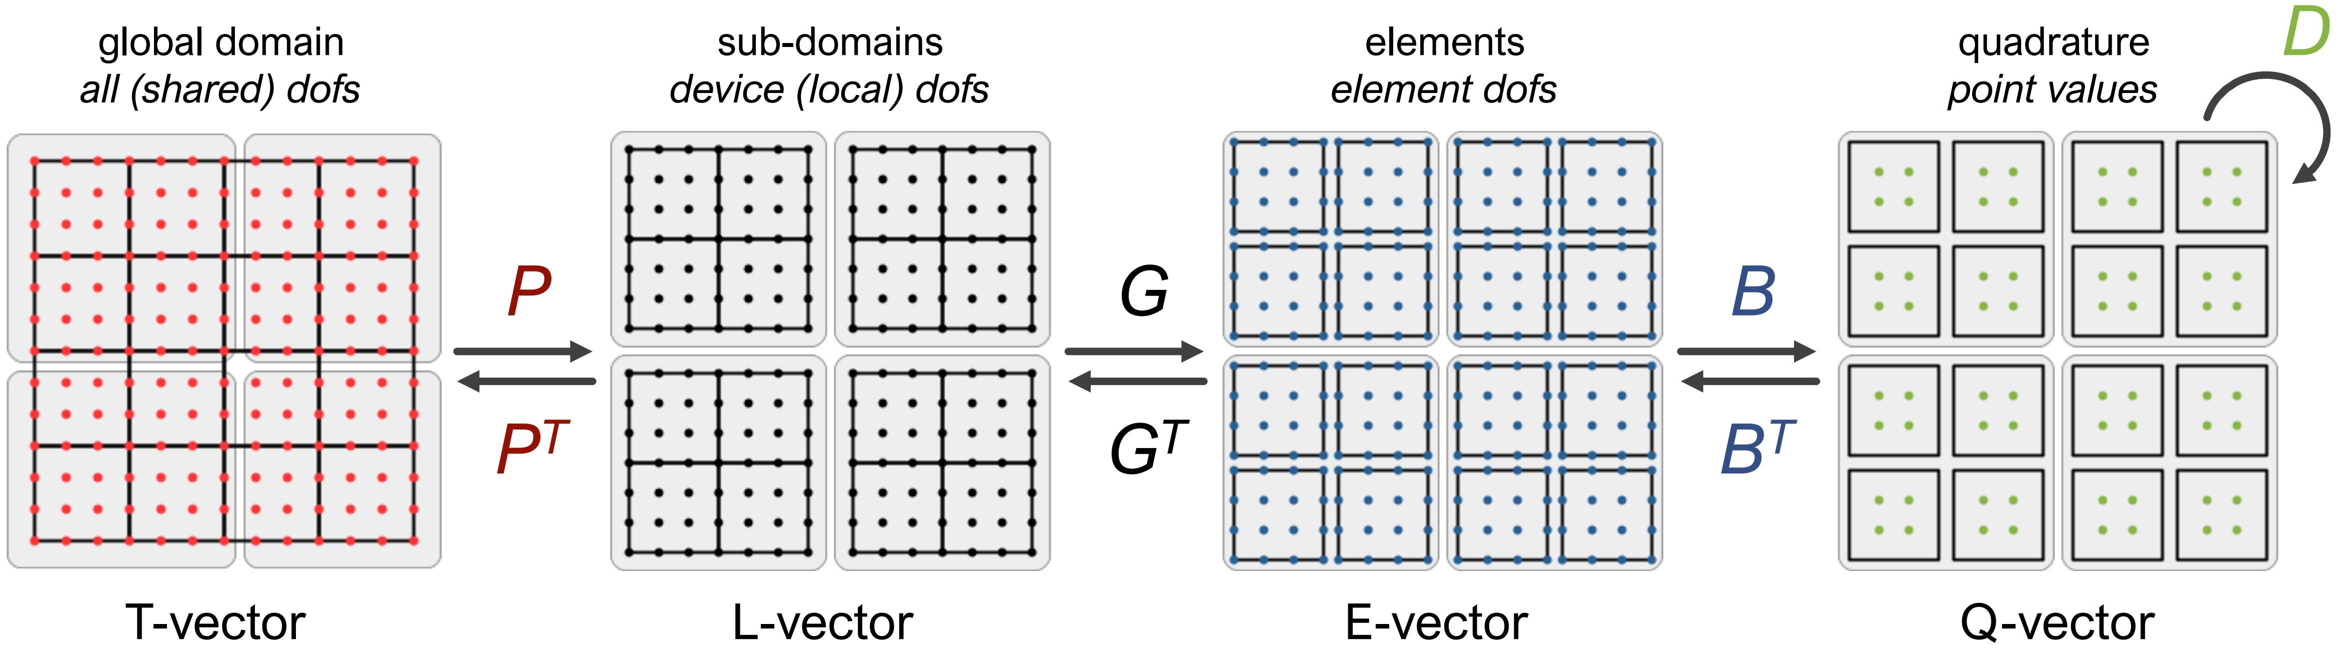

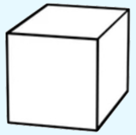

- **Store only D, evaluate action of B via B_{1d}**
- **Batched tensor contractions (B_{1d} the same)**

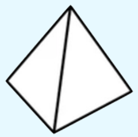

- **Compute and store B^T D B**
- **Batched GEMMs (B the same)**

CEED Diffusion Benchmark in MFEM

Using libCEED's cuda/gen and cuda/magma-det backends