Combinatorial Scientific Computing and Petascale Simulations (CSCAPES)

SAND2007-2094C

SciDAC Applied Math Institute

Alex Pothen, Florin Dobrian and Assefaw Gebremedhin, Old Dominion University

Erik Boman, Karen Devine and Bruce Hendrickson, Sandia National Laboratories

Paul Hovland, Sanjukta Bhowmick, Boyana Norris and Jean Utke, Argonne National Laboratory

Umit Catalyurek, Ohio State University

Michelle Strout, Colorado State University

www.cscapes.org

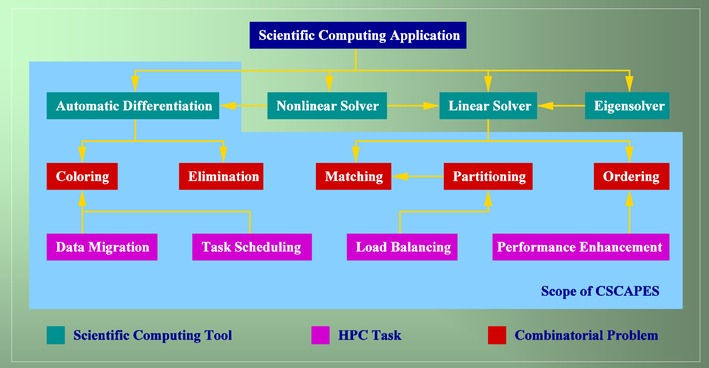

Combinatorial Scientific Computing **(CSC)** is concerned with the development, analysis and utilization of discrete algorithms in scientific and engineering applications. Graph and hypergraph algorithms are among the fundamental tools of CSC. They play a crucial enabling role in applications requiring parallelization, solution of differential equations, mesh generation, optimization, etc. The CSCAPES Institute is founded on the recognition of this critical role. The diagram above highlights some of the relationships between the research areas in CSCAPES and a typical SciDAC application; an arrow from A to B indicates that A in some sense uses B.

Load balancing: In parallel computations, data and tasks have to be mapped to processors to balance workload and reduce communication costs. In dynamic applications, these goals have to be met as the computations and communications change at each iteration. Zoltan is a toolkit for parallelization and load balancing. Recently a hypergraph based partitioner that has superior performance relative to earlier methods has been included in Zoltan. CSCAPES plans to extend the capabilities of Zoltan in several ways to support petascale applications.

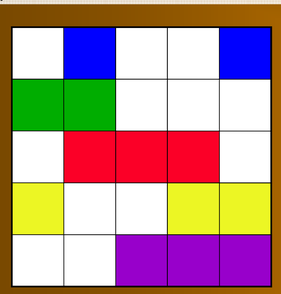

Application matrix

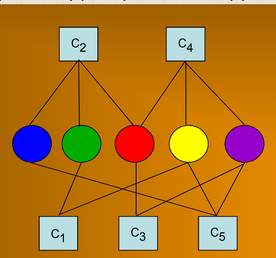

Hypergraph representation used by partitioner

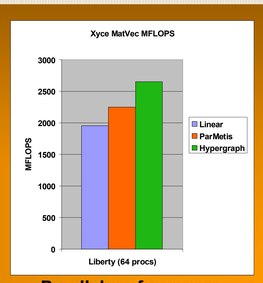

Parallel performance improvement

Performance enhancement: Modern microprocessors are highly sensitive to the spatial and temporal locality of data. **Reordering** the vertices and elements in a mesh can have significant impact on performance. We have developed several reordering algorithms that use the hypergraph representation of a matrix. These algorithms can improve the performance of a mesh smoothing application by nearly 50%. We plan to further develop and enhance these algorithms.

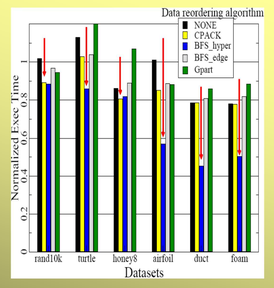

CSCAPES Mission

- **Research and development**
  - Provide load balancing and parallelization toolkits for petascale computation
  - Develop advanced automatic differentiation capabilities for complex applications
  - Advance the state-of-the art in large-scale graph and sparse matrix computations
- **Training and outreach**
  - In collaboration with SciDAC SAPs and CETs, support scientific applications with underlying graph and hypergraph algorithms
  - Organize workshops, tutorials and short courses in CSC
  - Train researchers in CSC skills at pre-doctoral and post-doctoral levels

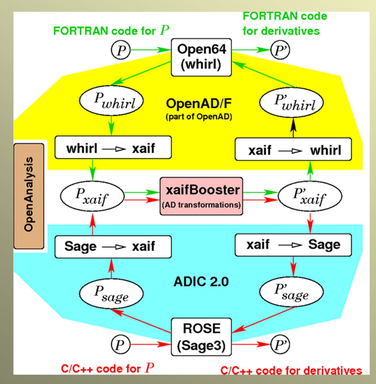

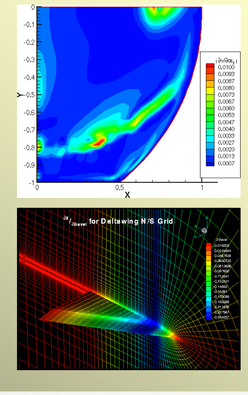

Automatic Differentiation (AD) computes analytic derivatives of functions specified by programs. Efficient derivative accumulation can be posed as graph transformation (**elimination**) problems. Checkpointing for irregular problems is another source of combinatorial problems in AD. CSCAPES will develop efficient algorithms for these problems, extend the capabilities of the AD tools ADIC and ADIFOR using the OpenAD framework, and provide users with integrated load balancing and coloring capabilities.

Graph coloring: Assorted variants of graph coloring are key tools for reducing the work required to compute sparse derivative matrices using automatic differentiation. Coloring is also used in discovering concurrency in parallel computation. We have developed novel coloring algorithms to make Jacobian and Hessian computations efficient. Preliminary parallel versions have been developed for some of these variants. We plan to extend the parallel coloring software and integrate it with automatic differentiation tools.

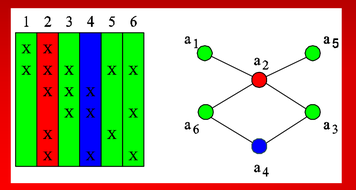

A symmetrically orthogonal partition of a Hessian and its representation as a star coloring of the adjacency graph

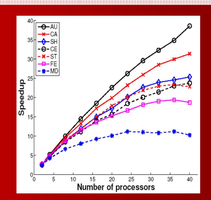

Performance of parallel distance-1 coloring algorithm

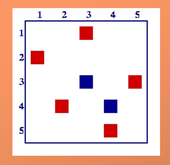

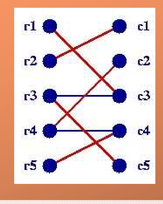

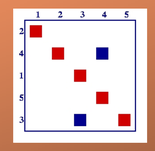

Graph Matching: Various kinds of graph matching are critical components of computational science applications. Traditional matching algorithms compute optimal solutions in superlinear time, but current trends are toward algorithms that find approximate solutions faster. CSCAPES plans to develop petascale parallel matching algorithms based on approximation techniques.

Office of Science U.S. Department of Energy

Old Dominion University

Sandia National Laboratories

Argonne National Laboratory

The Ohio State University

Colorado State University