



Tools and Primitives for High Performance Graph Computation

John R. Gilbert

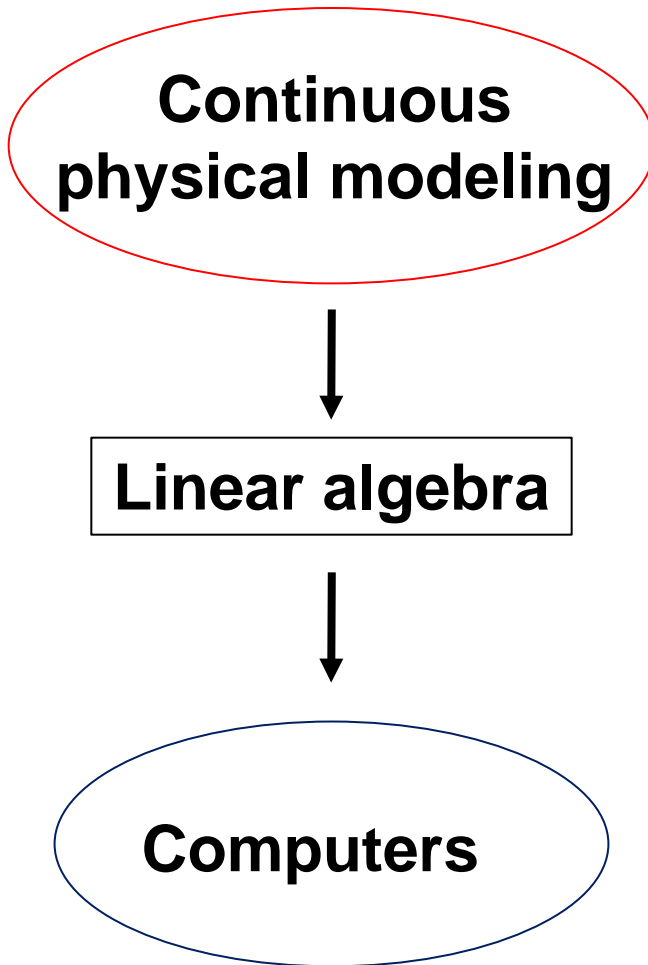
University of California, Santa Barbara

Aydin Buluç (LBNL)

Adam Lugowski (UCSB)

SIAM Minisymposium on
Analyzing Massive Real-World Graphs
July 12, 2010


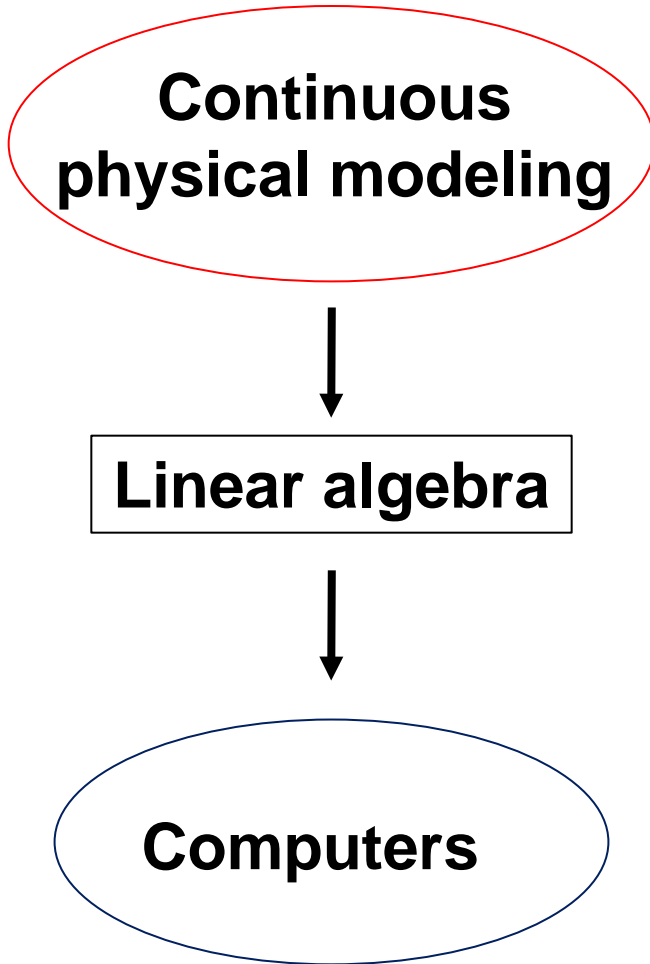
An analogy?



As the “middleware” of scientific computing, linear algebra has supplied or enabled:


- Mathematical tools
- “Impedance match” to computer operations
- High-level primitives
- High-quality software libraries
- Ways to extract performance from computer architecture
- Interactive environments

An analogy?



An analogy? Well, we're not there yet


- ✓ Mathematical tools
- ? “Impedance match” to computer operations
- ? High-level primitives
- ? High-quality software libs
- ? Ways to extract performance from computer architecture
- ? Interactive environments



The Primitives Challenge

- By analogy to numerical scientific computing. . .
- What should the combinatorial BLAS look like?

Basic Linear Algebra Subroutines (BLAS): Speed (MFlops) vs. Matrix Size (n)



Primitives should ...

- Supply a common notation to express computations
- Have broad scope but fit into a concise framework
- Allow programming at the appropriate level of abstraction and granularity
- Scale seamlessly from desktop to supercomputer
- Hide architecture-specific details from users


The case for sparse matrices

Many irregular applications contain coarse-grained parallelism that can be exploited by abstractions at the proper level.


Traditional graph computations	Graphs in the language of linear algebra
Data driven, unpredictable communication.	Fixed communication patterns
Irregular and unstructured, poor locality of reference	Operations on matrix blocks exploit memory hierarchy
Fine grained data accesses, dominated by latency	Coarse grained parallelism, bandwidth limited

Sparse array-based primitives


Sparse matrix-matrix multiplication (SpGEMM)




Sparse matrix-dense vector multiplication



Element-wise operations





Sparse matrix indexing




Matrices on various semirings: $(x, +)$, (and, or) , $(+, \min)$, ...


Multiple-source breadth-first search




Multiple-source breadth-first search




A^T




X



$A^T X$



Multiple-source breadth-first search



- Sparse array representation => space efficient
- Sparse matrix-matrix multiplication => work efficient
- Three possible levels of parallelism: searches, vertices, edges


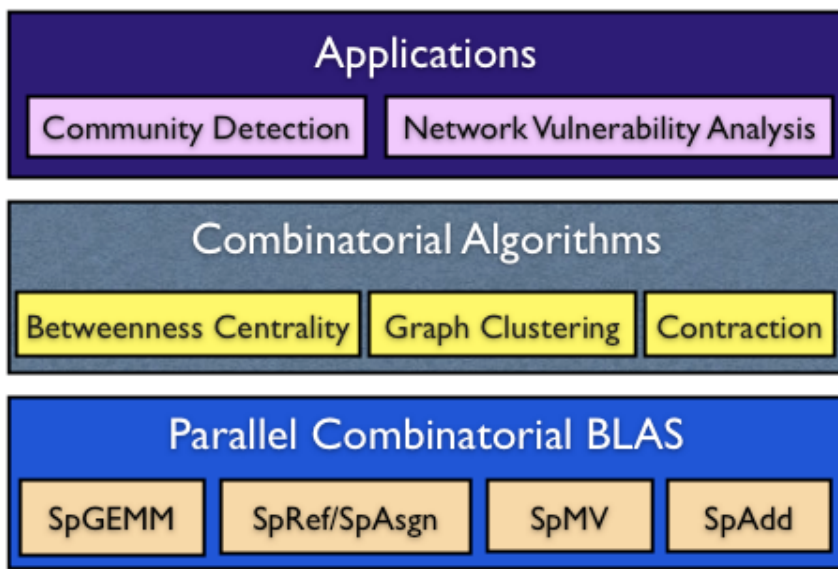
A Few Examples

Combinatorial BLAS

[Buluc, G]

A parallel graph library based on distributed-memory sparse arrays and algebraic graph primitives

Typical software stack



Betweenness Centrality (BC)


What fraction of shortest paths pass through this node?

$$C_B(v) = \sum_{\substack{s \neq v \neq t \in V \\ s \neq t}} \frac{\sigma_{st}(v)}{\sigma_{st}}$$

Brandes' algorithm

BC performance in distributed memory

BC performance



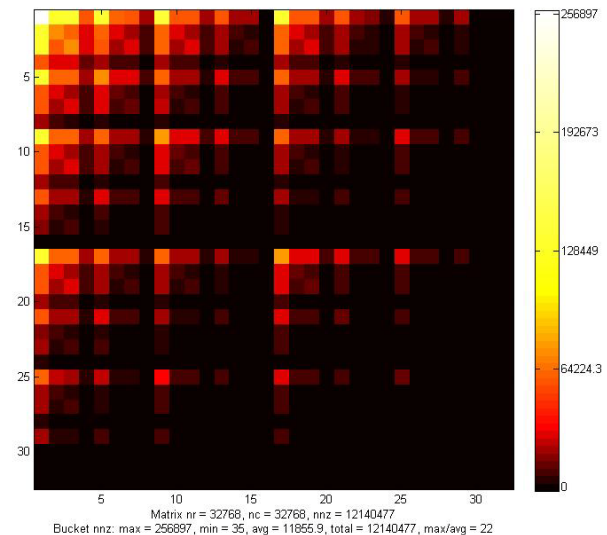


- TEPS = Traversed Edges Per Second
- One page of code using C-BLAS

KDT: A toolbox for graph analysis and pattern discovery

[G, Reinhardt, Shah]

Layer 1: Graph Theoretic Tools

- Graph operations
- Global structure of graphs
- Graph partitioning and clustering
- Graph generators
- Visualization and graphics
- Scan and combining operations
- Utilities


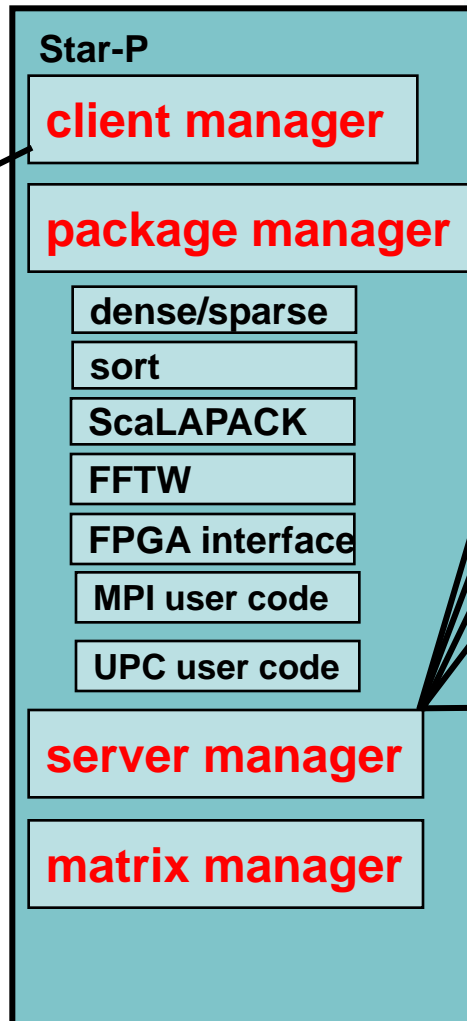


Star-P architecture

MATLAB®



Ordinary Matlab variables




Distributed matrices

Landscape connectivity modeling



- Habitat quality, gene flow, corridor identification, conservation planning
- Pumas in southern California: 12 million nodes, < 1 hour
- Targeting larger problems: Yellowstone-to-Yukon corridor




Figures courtesy of Brad McRae

- Predicting gene flow with resistive networks
- Matlab, Python, and Star-P (parallel) implementations
- Combinatorics:
 - Initial discrete grid: ideally 100m resolution (for pumas)
 - Partition landscape into connected components
 - Graph contraction: habitats become nodes in resistive network
- Numerics:
 - Resistance computations for pairs of habitats in the landscape
 - Iterative linear solvers invoked via Star-P: Hypre (PCG+AMG)

A Few Nuts & Bolts


SpGEMM: sparse matrix \times sparse matrix

- Graph clustering (Markov, peer pressure)
- Subgraph / submatrix indexing
- Shortest path calculations
- Betweenness centrality
- Graph contraction
- Cycle detection
- Multigrid interpolation & restriction
- Colored intersection searching
- Applying constraints in finite element computations
- Context-free parsing ...




$$\begin{bmatrix} & 1 & & \\ & & & \\ & & & \\ & & & 1 \end{bmatrix} \times \begin{bmatrix} & & & \\ \color{green}1 & \color{green}1 & & \\ & & & \\ \color{green}1 & \color{green}1 & & \end{bmatrix} \times \begin{bmatrix} 1 & & \\ & 1 & \\ & & 1 \end{bmatrix}$$

Two Versions of Sparse GEMM




1D block-column distribution



2D block distribution

Parallelism in multiple-source BFS




Three levels of parallelism from 2-D data decomposition:


- columns of X : over multiple simultaneous searches
- rows of X & columns of A^T : over multiple frontier nodes
- rows of A^T : over edges incident on high-degree frontier nodes

Modeled limits on speedup, sparse 1-D & 2-D

1-D algorithm

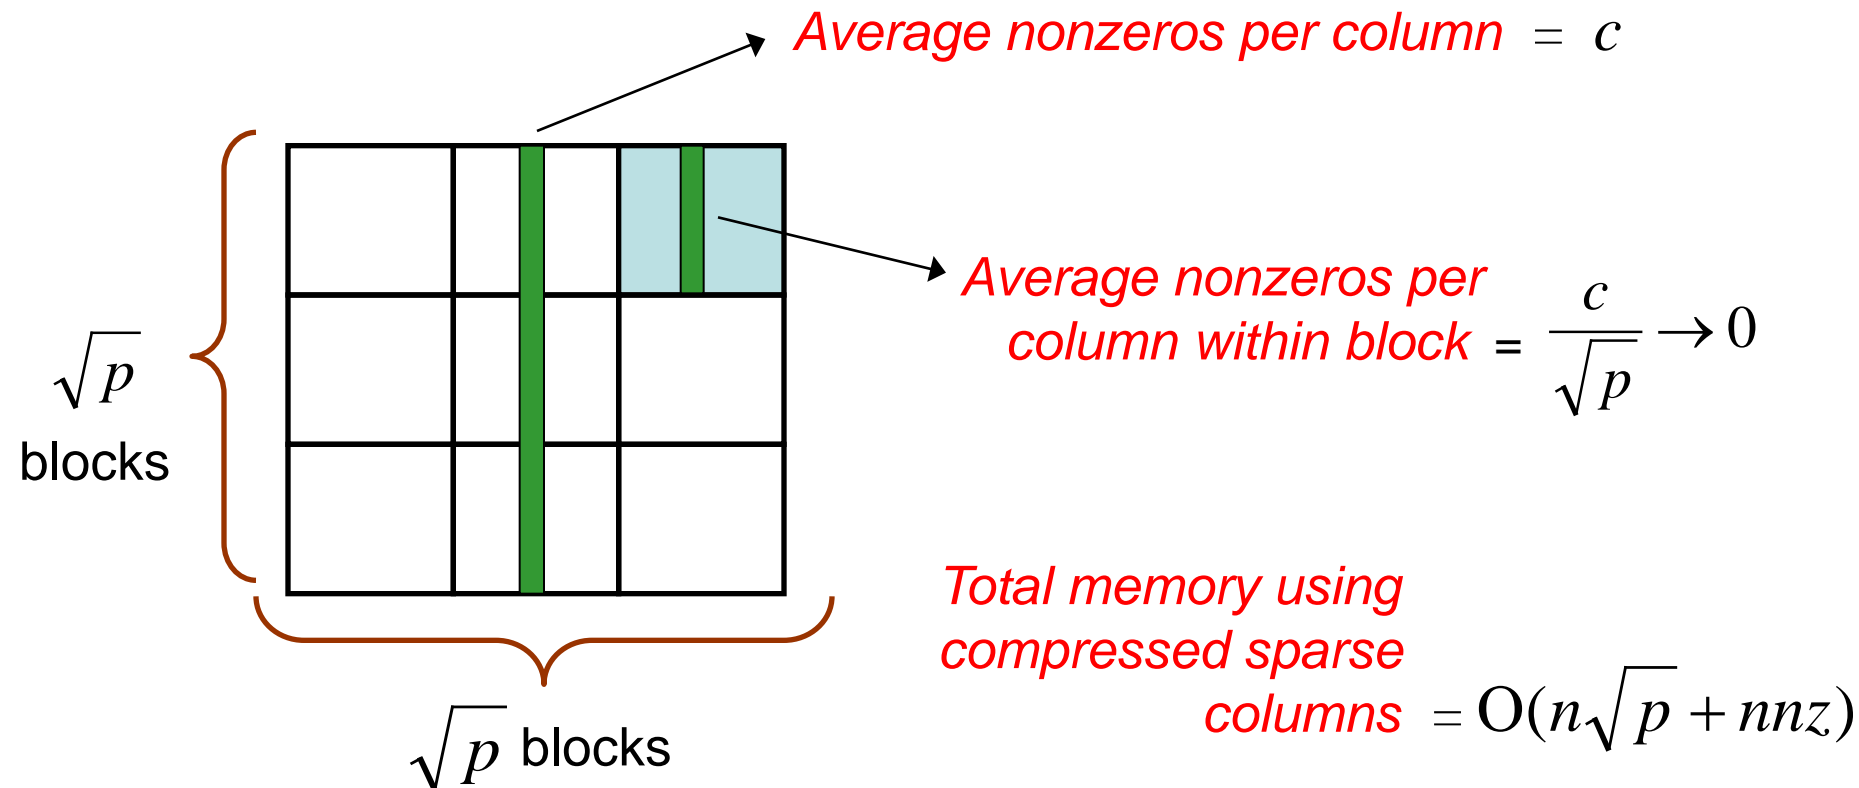


2-D algorithm



- 1-D algorithms do not scale beyond 40x
- Break-even point is around 50 processors


Submatrices are hypersparse ($nnz \ll n$)




Any algorithm whose complexity depends on matrix dimension n is asymptotically too wasteful.

Distributed-memory sparse matrix-matrix multiplication

- 2D block layout
- Outer product formulation
- Sequential “hypersparse” kernel



Parallel PSpGEMM Scalability, Rmat-Scale20

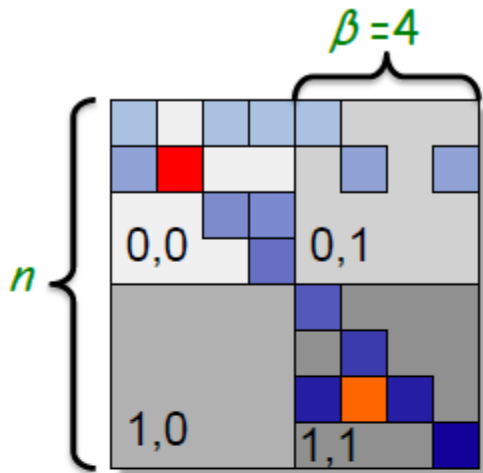


Time vs Number of cores -- 1M-vertex RMat

- Scales well to hundreds of processors
- Betweenness centrality benchmark: over 200 MTEPS
- Experiments: TACC Lonestar cluster


CSB: Compressed sparse block storage

[Buluc, Fineman, Frigo, G, Leiserson]



$n \times n$ matrix with nnz nonzeros, in $\beta \times \beta$ blocks

Block pointer *Dense collection of sparse blocks*




- Dense 2-D array of sparse blocks
- Blocks stored in row-major order
- Entries within blocks stored in **Z-morton order**
- Storage = 1 word / edge + 1 word / vertex



CSB for parallel Ax and $A^T x$

[Buluc, Fineman, Frigo, G, Leiserson]

- Efficient multiplication of a sparse matrix *and its transpose* by a vector
- Compressed sparse block storage
- Critical path never more than $\sim \sqrt{n} \cdot \log(n)$
- Multicore / multsocket architectures



From Semirings to Computational Patterns

Matrices over semirings

- Matrix multiplication $\mathbf{C} = \mathbf{AB}$ (or matrix/vector):

$$\mathbf{C}_{i,j} = \mathbf{A}_{i,1} \times \mathbf{B}_{1,j} + \mathbf{A}_{i,2} \times \mathbf{B}_{2,j} + \dots + \mathbf{A}_{i,n} \times \mathbf{B}_{n,j}$$

- Replace scalar operations \times and $+$ by

\otimes : associative, distributes over \oplus , identity 1

\oplus : associative, commutative, identity 0 annihilates under \otimes

- Then $\mathbf{C}_{i,j} = \mathbf{A}_{i,1} \otimes \mathbf{B}_{1,j} \oplus \mathbf{A}_{i,2} \otimes \mathbf{B}_{2,j} \oplus \dots \oplus \mathbf{A}_{i,n} \otimes \mathbf{B}_{n,j}$

- Examples: $(\times, +)$; (and, or) ; $(+, \min)$; . . .

- No change to data reference pattern or control flow

From semirings to computational patterns

Sparse matrix times vector as a semiring operation:

- Given vertex data \mathbf{x}_i and edge data $\mathbf{a}_{i,j}$
- For each vertex j of interest, compute

$$\mathbf{y}_j = \mathbf{a}_{i_1,j} \otimes \mathbf{x}_{i_1} \oplus \mathbf{a}_{i_2,j} \otimes \mathbf{x}_{i_2} \oplus \dots \oplus \mathbf{a}_{i_k,j} \otimes \mathbf{x}_{i_k}$$

- User specifies: definition of operations \otimes and \oplus


From semirings to computational patterns

Sparse matrix times vector as a computational pattern:

- Given vertex data and edge data
- For each vertex of interest, combine data from neighboring vertices and edges
- User specifies: desired computation on data from neighbors

SpGEMM as a computational pattern

- Explore length-two paths that use specified vertices
- Possibly do some filtering, accumulation, or other computation with vertex and edge attributes
- E.g. “friends of friends” (per Lars Backstrom)
- May or may not want to form the product graph explicitly
- Formulation as semiring matrix multiplication is often possible but sometimes clumsy
- Same data flow and communication patterns as in SpGEMM




Graph BLAS: A pattern-based library

- User-specified operations and attributes give the performance benefits of algebraic primitives with a more intuitive and flexible interface.
- Common framework integrates algebraic (edge-based), visitor (traversal-based), and map-reduce patterns.
- 2D compressed sparse block structure supports user-defined edge/vertex/attribute types and operations.
- “Hypersparse” kernels tuned to reduce data movement.
- Initial target: manycore and multsocket shared memory.

Challenge: Complete the analogy . . .

- ✓ Mathematical tools
- ? “Impedance match” to computer operations
- ? High-level primitives
- ? High-quality software libs
- ? Ways to extract performance from computer architecture
- ? Interactive environments



Some challenges

- Fault tolerance
- Uncertainty and probabilistic attributes
- Fine-grained dynamic updates
- Interactive environments for exploration & productivity
- Hardware architecture