

UCSB



Parallel Combinatorial BLAS and Applications in Graph Computations

Aydın Buluç

John R. Gilbert


University of California, Santa Barbara

SIAM ANNUAL MEETING 2009

July 8, 2009

Primitives for Graph Computations

- By analogy to numerical linear algebra,
- What would the combinatorial BLAS look like?

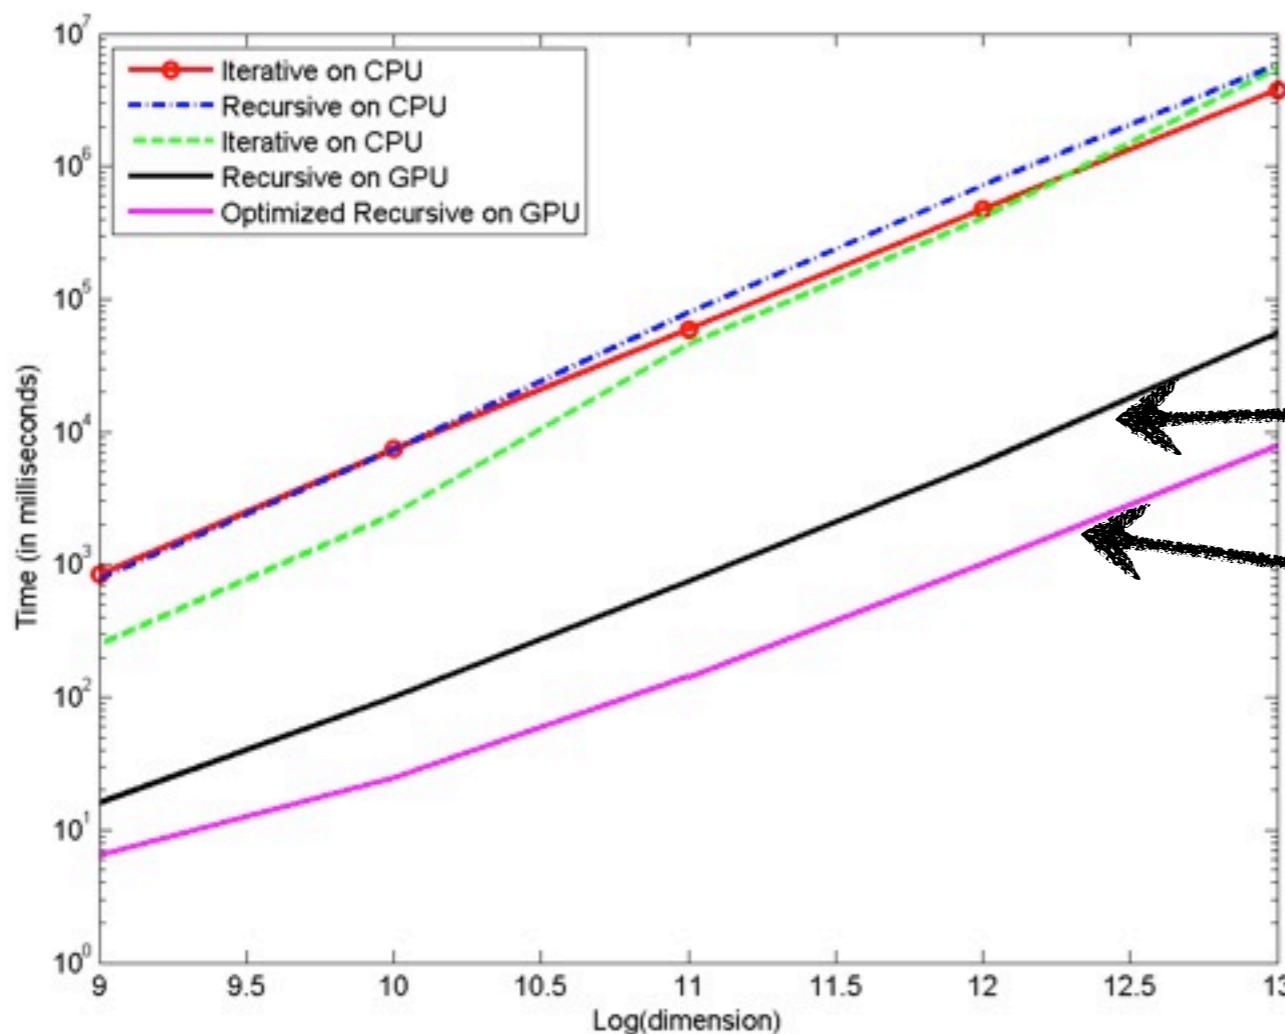


BLAS 3 (n-by-n matrix-matrix multiply)
BLAS 2 (n-by-n matrix-vector multiply)
BLAS 1 (sum of scaled n-vectors)

The Case for Primitives

It takes a “certain” level of expertise to get any kind of performance in this jungle of parallel computing

- I think you’ll agree with me by the end of the talk :)



What’s bandwidth anyway?

480x

I can just implement it (w/ enough coffee)

The right primitive !

All pairs shortest paths on the GPU

The Case for Sparse Matrices

- Many irregular applications contain sufficient coarse-grained parallelism that can ONLY be exploited using abstractions at proper level.

| Traditional graph computations | Graphs in the language of linear algebra |
|--|---|
| Data driven. Unpredictable communication patterns | Fixed communication patterns. Overlapping opportunities |
| Irregular and unstructured. Poor locality of reference | Operations on matrix blocks. Exploits memory hierarchy |
| Fine grained data accesses. Dominated by latency | Coarse grained parallelism. Bandwidth limited |

Identification of Primitives


- ▶ Sparse matrix-matrix multiplication (SpGEMM)
Most general and challenging parallel primitive.
- ▶ Sparse matrix-vector multiplication (SpMV)
- ▶ Sparse matrix-transpose-vector multiplication (SpMVT)
Equivalently, multiplication from the left
- ▶ Addition and other point-wise operations (SpAdd)
Included in SpGEMM, “proudly” parallel
- ▶ Indexing and assignment (SpRef, SpAsgn)
 $A(I,J)$ where I and J are arrays of indices
Reduces to SpGEMM

Matrices on semirings, e.g. $(\times, +)$, (and, or), $(+, \min)$

Why focus on SpGEMM?


$$\begin{array}{c} \text{length}(I) \\ \left\{ \begin{array}{|c|c|c|c|} \hline 0 & 1 & 0 & 0 \\ \hline 0 & 0 & 0 & 1 \\ \hline \end{array} \right\} \times \begin{array}{|c|c|c|c|} \hline \square & \square & \square & \square \\ \hline \color{green}\square & \color{green}\square & \color{green}\square & \square \\ \hline \square & \square & \square & \square \\ \hline \color{green}\square & \color{green}\square & \color{green}\square & \square \\ \hline \end{array} \times \begin{array}{c} \text{length}(J) \\ \left\{ \begin{array}{|c|c|c|} \hline 1 & 0 & 0 \\ \hline 0 & 1 & 0 \\ \hline 0 & 0 & 1 \\ \hline 0 & 0 & 0 \\ \hline \end{array} \right\} \end{array} = \begin{array}{|c|c|c|} \hline \color{green}\square & \color{green}\square & \color{green}\square \\ \hline \color{green}\square & \color{green}\square & \color{green}\square \\ \hline \end{array}$$

m *n*





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

Why focus on SpGEMM?

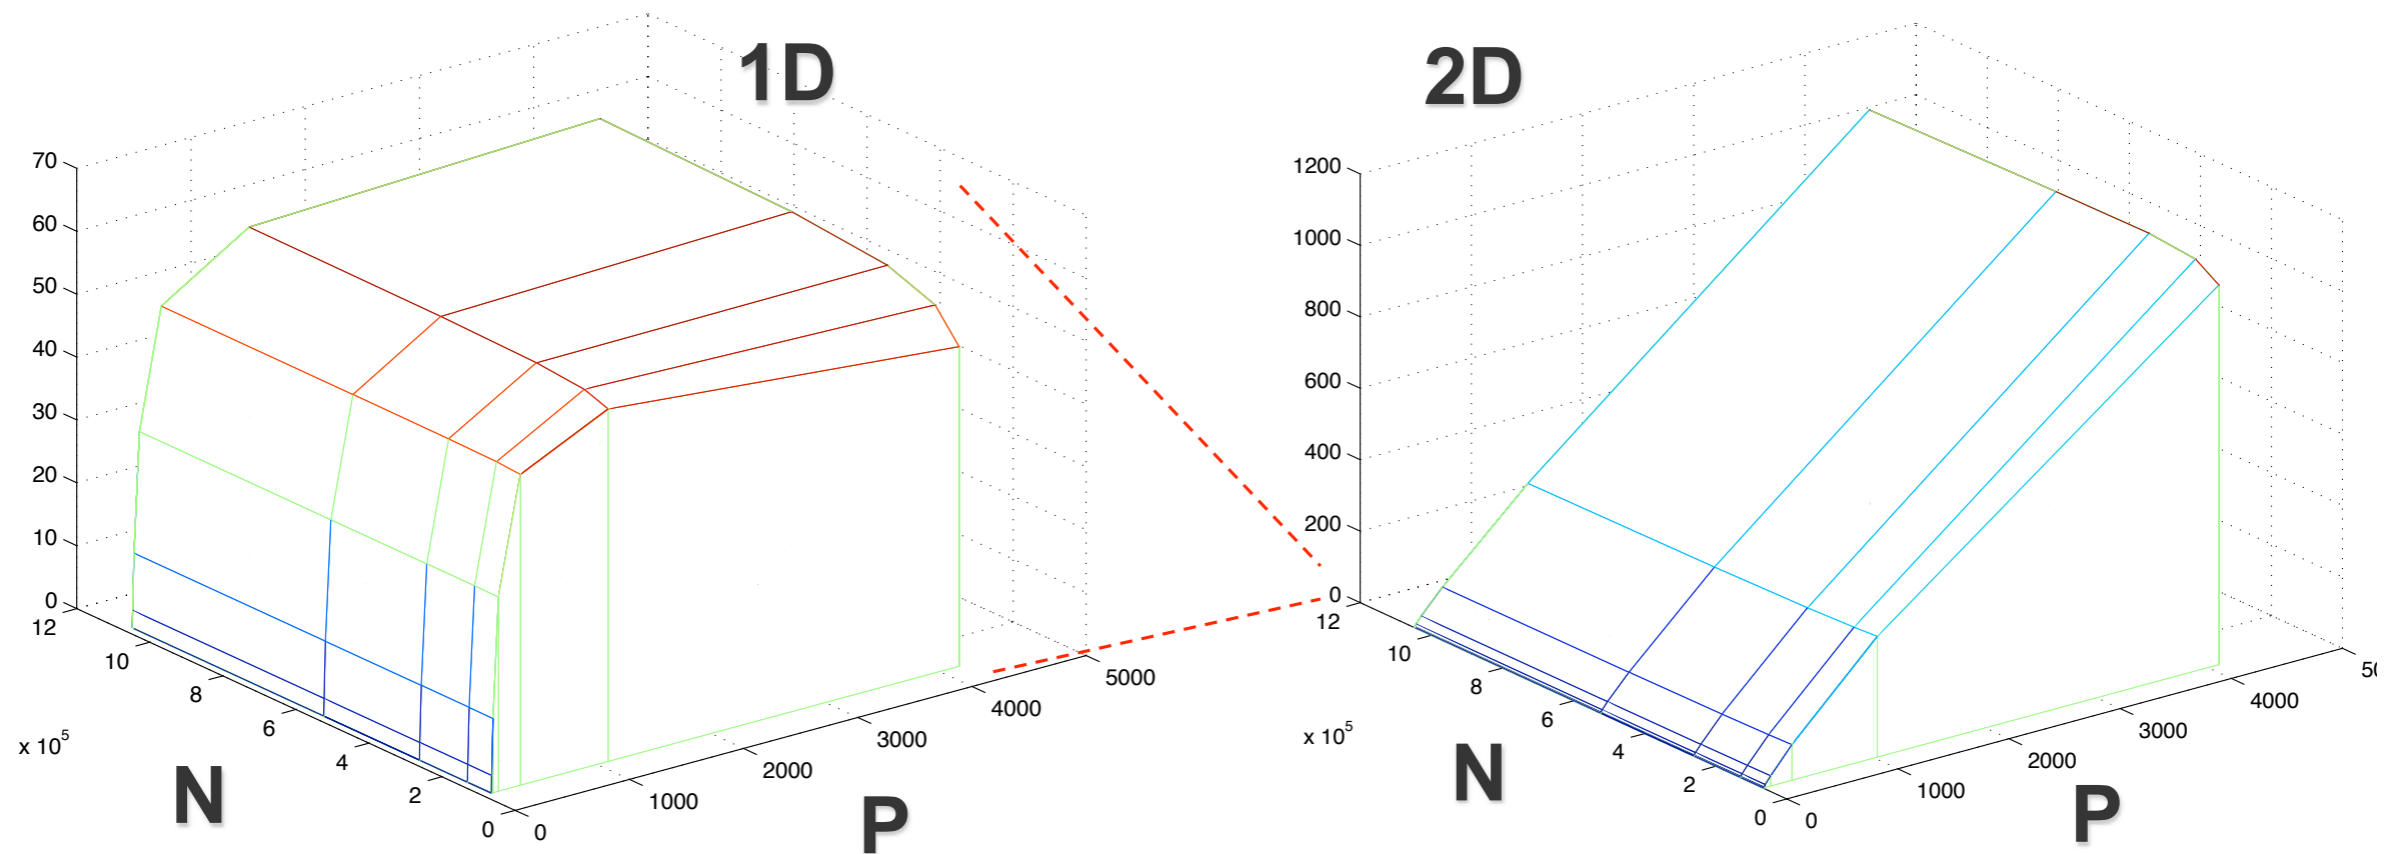


- Graph clustering (Markov, peer pressure)
- Shortest path calculations
- Betweenness centrality
- Subgraph / submatrix indexing
- Graph contraction
- Cycle detection
- Multigrid interpolation & restriction
- Colored intersection searching
- Applying constraints in finite element computations
- Context-free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



Vitals of Combinatorial BLAS

1. **Scalability**, in the presence of increasing *processors*, *problem size*, and *sparsity*.



In practice, 2D algorithms have the potential to scale, if implemented correctly. Overlapping communication, and maintaining load balance are crucial.



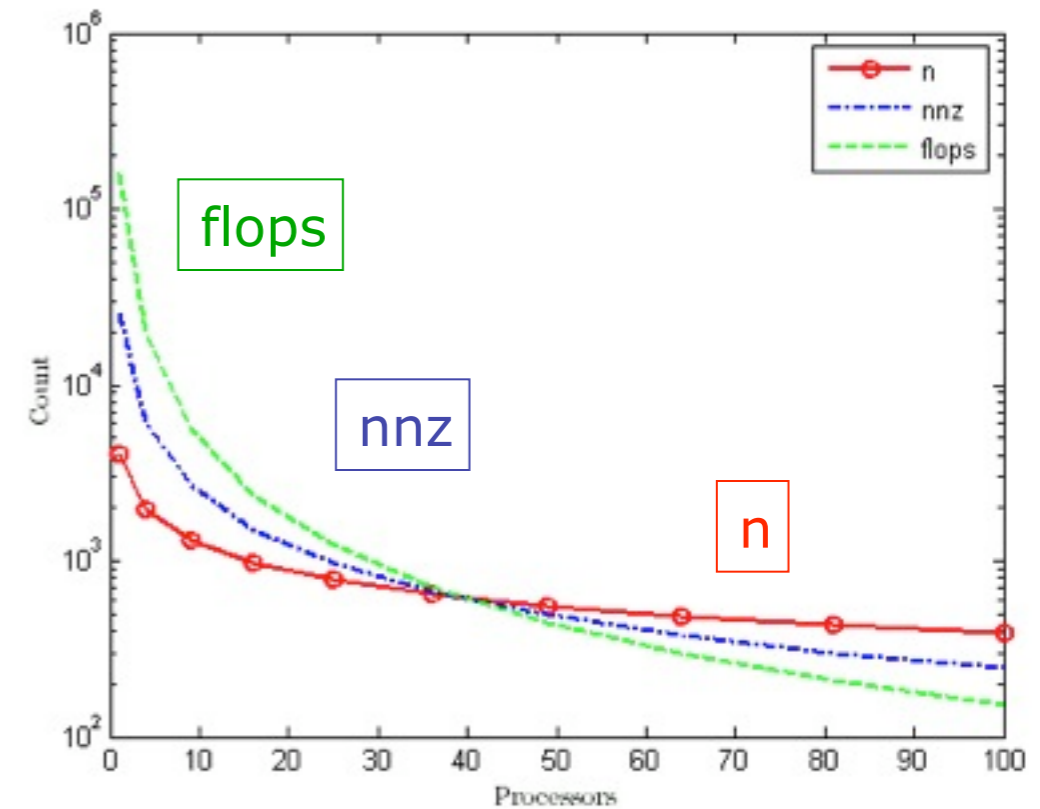
Sequential Kernel

Standard algorithm is $O(nnz + flops + n)$

$$n' \text{ (dimension)} \approx \frac{n}{\sqrt{p}}$$

$$nnz' \text{ (data size)} \approx \frac{nnz}{p}$$


$$flops' \text{ (work)} \approx \frac{flops}{p\sqrt{p}}$$



- Strictly $O(nnz)$ data structure
- Outer-product formulation
- Work-efficient

Node Level Considerations

Submatrices are *hypersparse* (i.e. $nnz \ll n$)


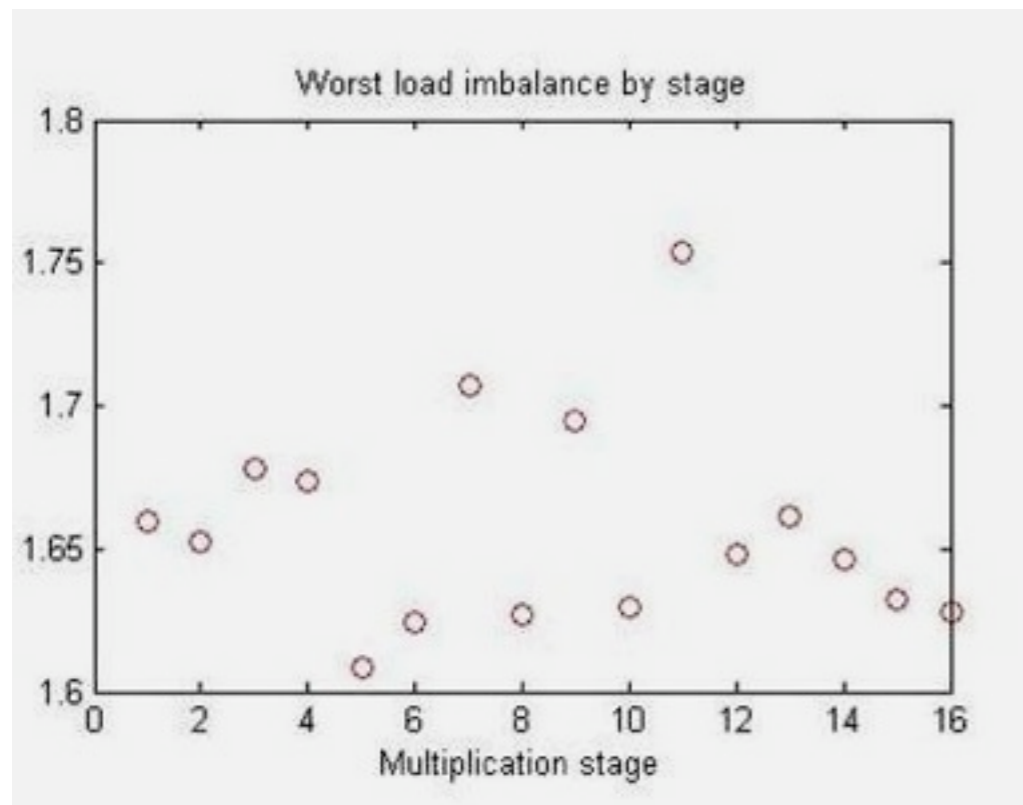


- A data structure or algorithm that depends on the matrix dimension n (e.g. CSR or CSC) is asymptotically too wasteful for submatrices

Addressing the Load Balance

RMat: Model for graphs with high variance on degrees



- Random permutations are useful. But...
- Bulk synchronous algorithms may still suffer:
- Asynchronous algorithms have no notion of stages.
- Overall, no significant imbalance.



Addressing the Load Balance


RMat: Model for graphs with high variance on degrees

- Random permutations are useful. But...
- Bulk synchronous algorithms may still suffer:
- Asynchronous algorithms have no notion of stages.
- Overall, no significant imbalance.



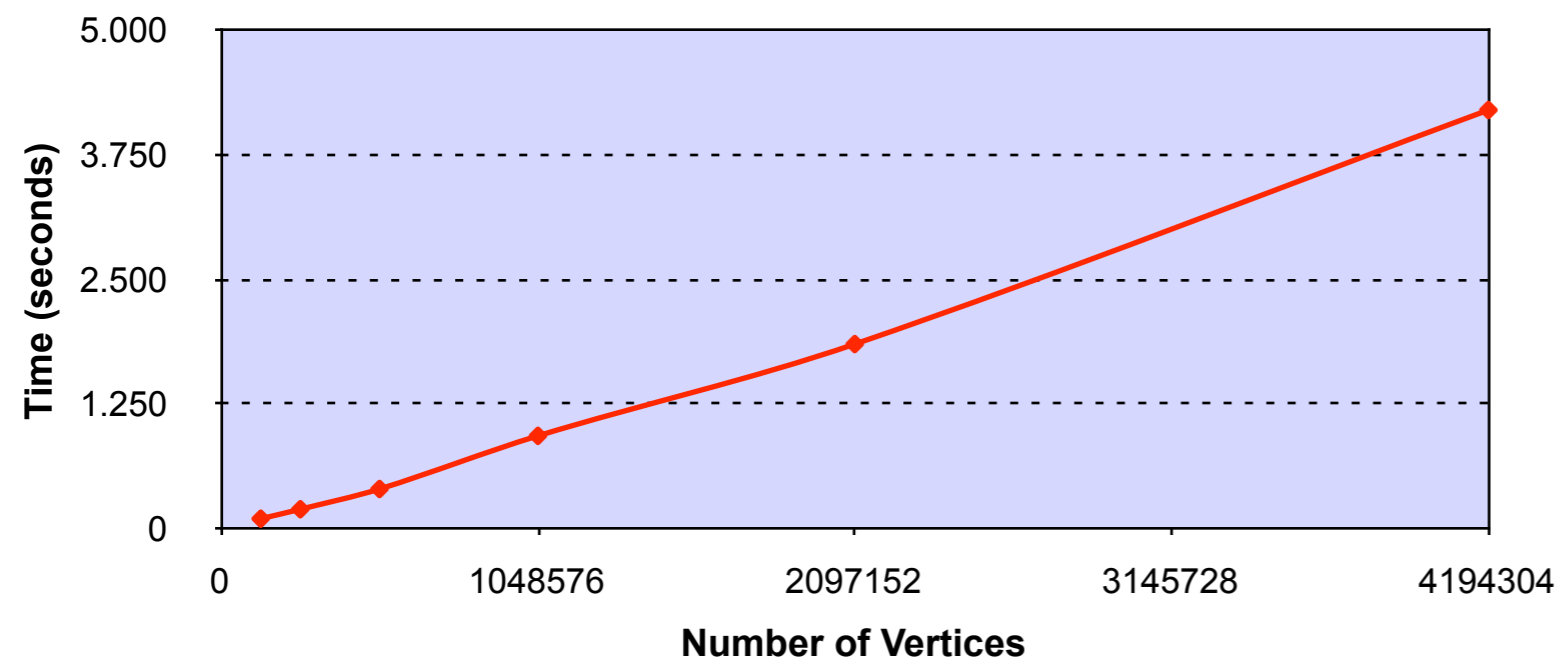
Scaling Results for SpGEMM

Parallel PSpGEMM Scalability, Rmat-Scale20



- Asynchronous implementation
One-sided MPI-2
- Runs on TACC's Lonestar cluster
- Dual-core dual-socket
Intel Xeon 2.66 Ghz
- RMat X RMat product
Average degree (nnz/n) ≈ 8

PSpGEMM Scalability with Increasing Problem Size 64 Processors



Vitals of Combinatorial BLAS

2. **Generality**, of the numeric type of matrix elements, algebraic operation performed, and the library interface.

Without the language abstraction penalty: C++ Templates

```
template <class IT, class NT, class DER>  
class SpMat;
```

- Achieve mixed precision arithmetic: Type traits
 - Enforcing interface and strong type checking: CRTP
 - General semiring operation: Function Objects
- ➡ *Abstraction penalty is not just a programming language issue.*
- ➡ *In particular, view matrices as indexed data structures and stay away from single element access (Interface should discourage)*

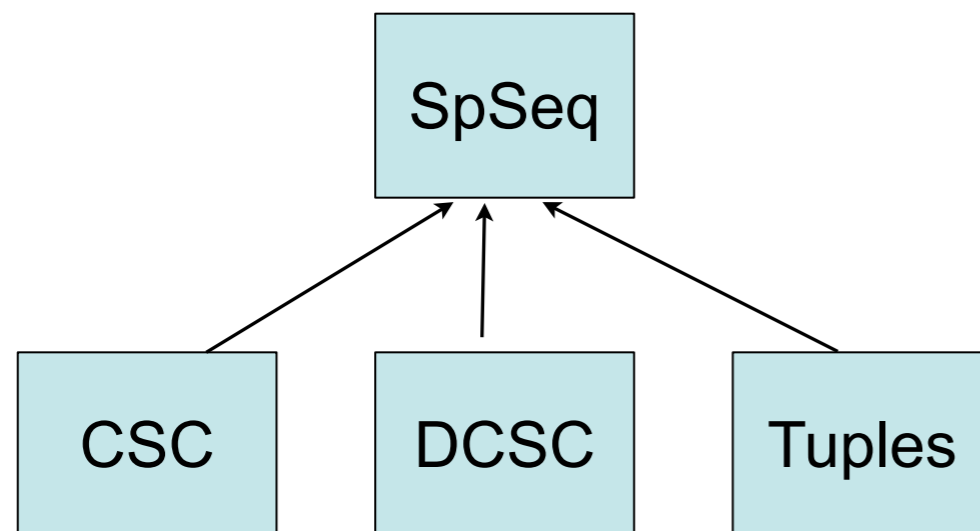
Vitals of Combinatorial BLAS

3. **Extendability**, of the library while maintaining compatibility and seamless upgrades.

- ➔ Decouple parallel logic from the sequential part.
- ➔ Even Boost' serializable concept might be restrictive (and slow)

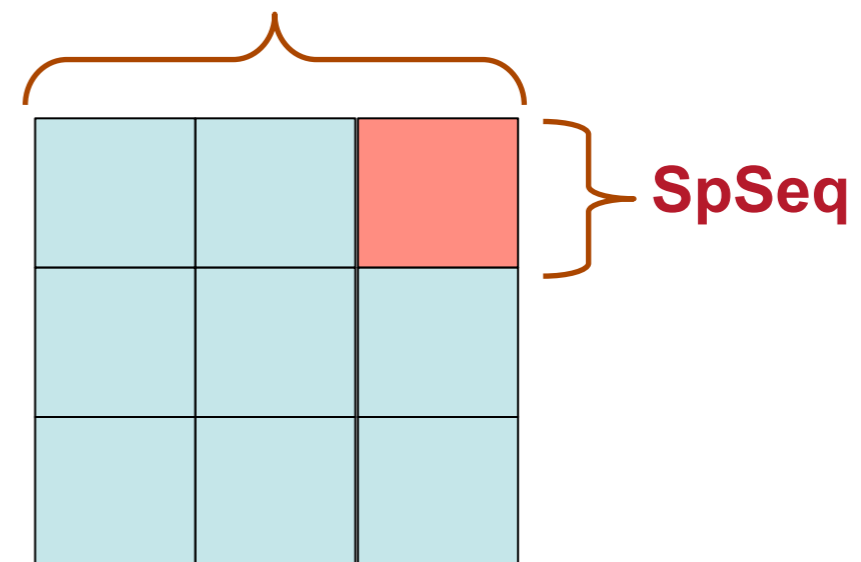
Commonalities:

- Support the sequential API
- Composed of a **number of arrays**

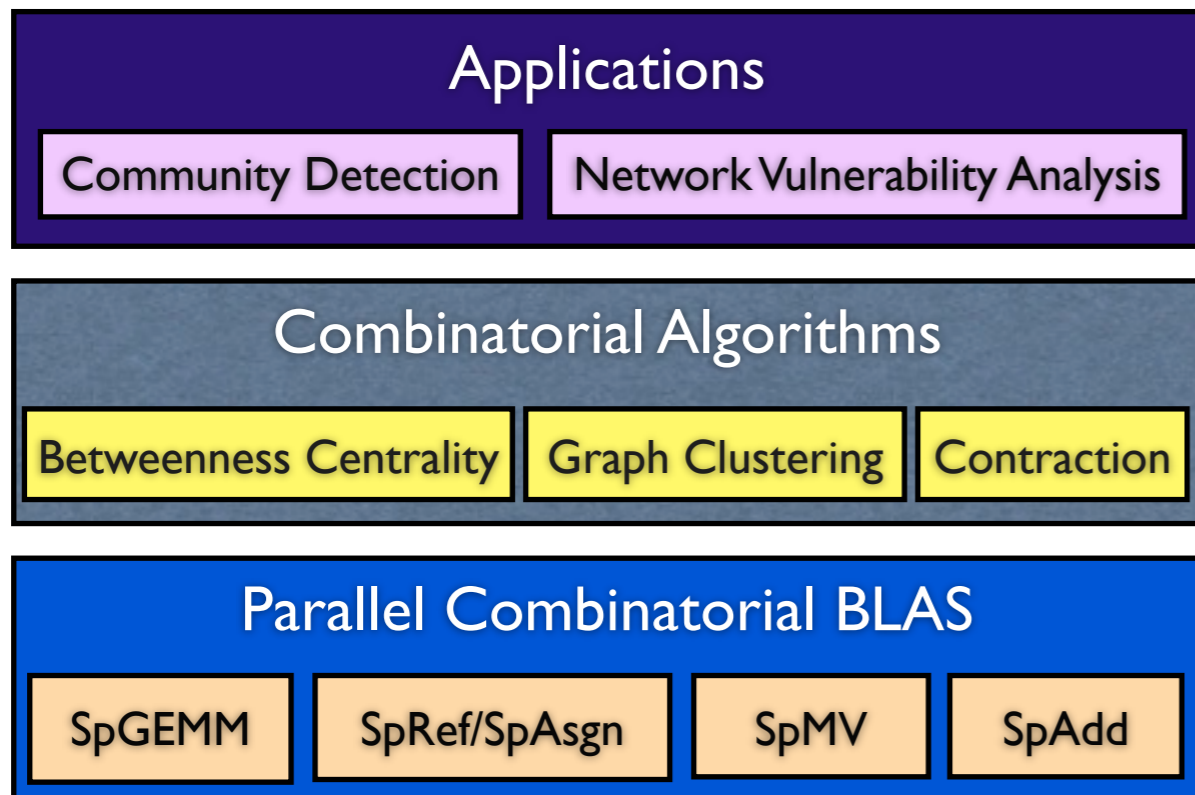


Any parallel logic:
asynchronous, bulk synchronous, etc


SpPar<Comm, SpSeq>



Applications and Algorithms



A typical software stack for an application enabled with the Combinatorial BLAS




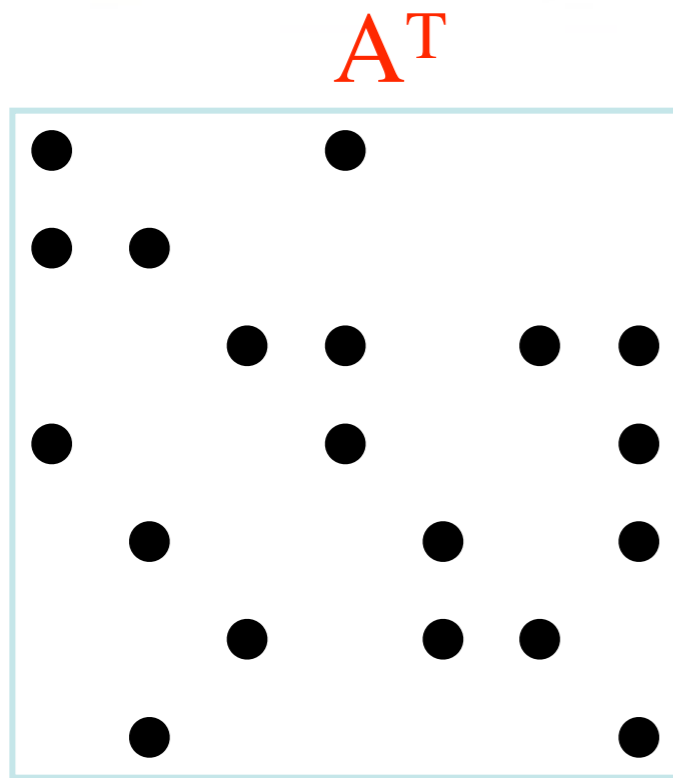
Betweenness Centrality

$C_B(v)$: Among all the shortest paths, what fraction of them pass through the node of interest?

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



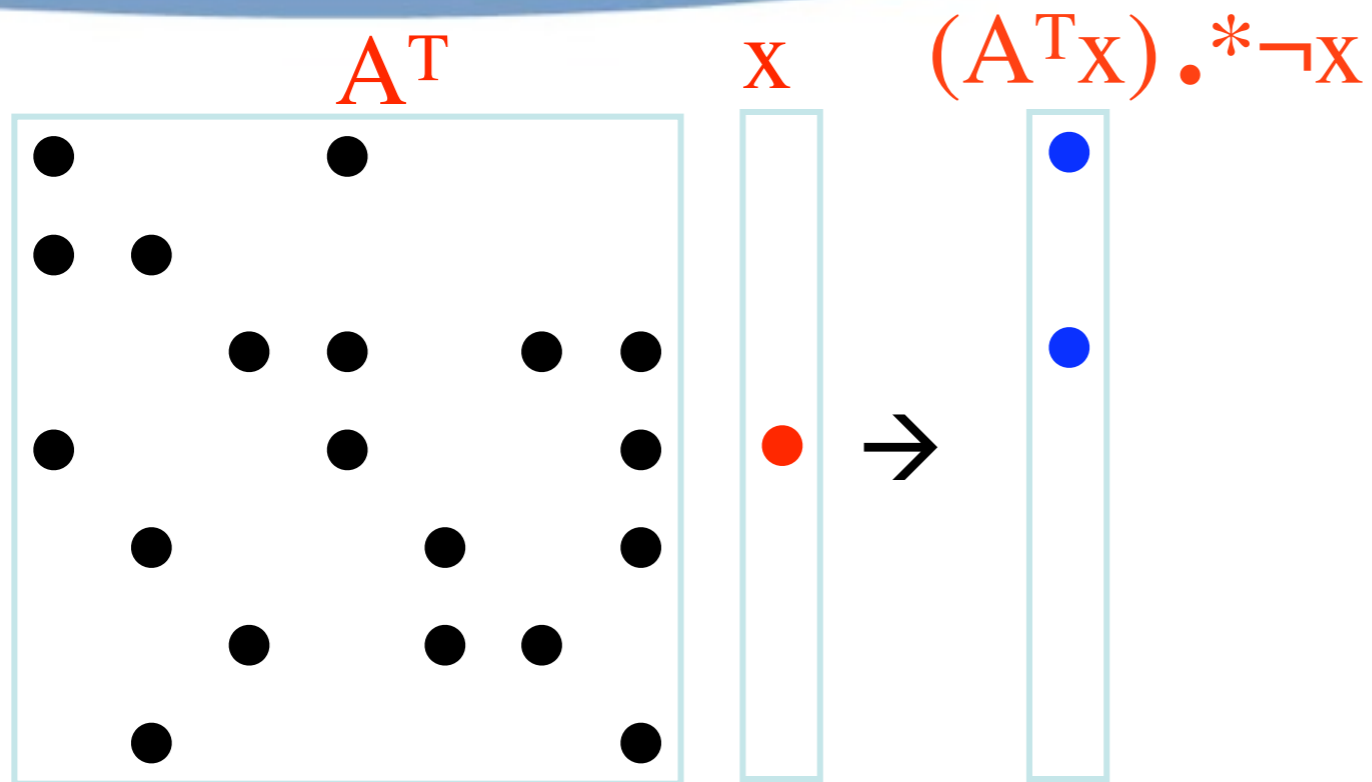
Brandes' algorithm

Betweenness Centrality using Sparse Matrices [Robinson, Kepner]




- Adjacency matrix: sparse array w/ nonzeros for graph edges
- Storage-efficient implementation from sparse data structures
- Betweenness Centrality Algorithm:
 1. Pick a starting vertex, v
 2. Compute shortest paths from v to all other nodes
 3. Starting with most distant nodes, roll back and tally paths

Betweenness Centrality using BFS



- Every iteration, another level of the BFS is discovered.
- Sparsity is preserved, but sparse matrix times sparse vector has very little potential parallelism (has $o(nnz)$ work)

Parallelism: Multiple-source BFS

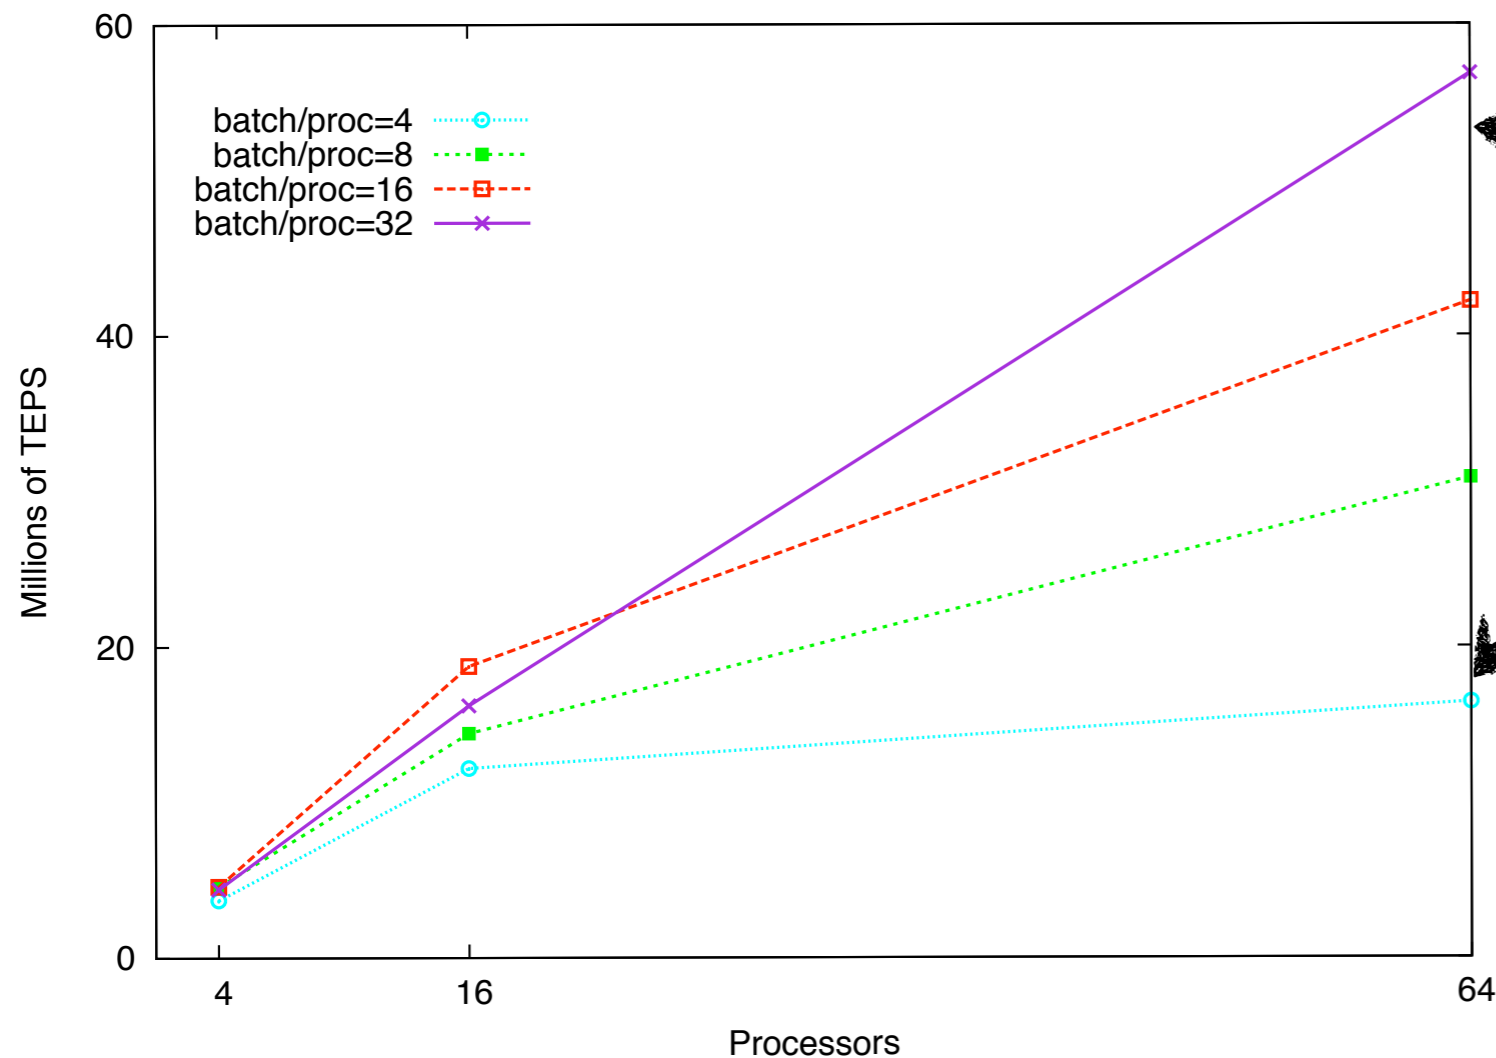


- Batch processing of multiple source vertices
- Sparse matrix-matrix multiplication \Rightarrow work efficient
- Potential parallelism is much higher
- Same applies to the tallying phase

Betweenness Centrality on Combinatorial BLAS

- Semi-basic implementation: *2D matrices*, synchronous matrix multiplication, no overlapping of communication with computation, *some remote DMA*, *mixed type arithmetic*, no template specialization for boolean matrices

Fundamental trade-off:
Parallelism vs memory usage



Batch processing greatly helps for large p

Input: RMAT scale 17

~1M edges only

- Likely to perform better on large inputs

- Code only a few lines longer than Matlab version

Thank You !

Questions?