EDGAR

- **E**nergy-efficient **D**ata and **G**raph **A**lgorithms **R**esearch
- Funded by Applied Math, ASCR
- Program Manager: Steven Lee
- Early Career Research Program (start: 2013)
- PI: Aydın Buluç (Berkeley Lab)
- Research Scientist: Ariful Azad (Berkeley Lab)
- Students:
  - Veronika Strnadova-Neeley (Oct 2015-2016, UCSB)
  - Carl Yang (March 2016-present, UC Davis)

The Reverse Cuthill-McKee Algorithm in Distributed-Memory

- **Goal:** Find a permutation **P** of a sparse matrix **A** so that the bandwidth of $\mathbf{PAP}^T$ is small.
- **Application:** Faster iterative solvers, e.g., preconditioned conjugate gradients (PCG).
- **Innovation in Parallel RCM Algorithm:**
  1. Step1: level-by-level vertex exploration and ordering. Approach: specialized breadth-first search using sparse matrix-sparse vector multiplication (SpMSpV) over a semiring
  2. Step2: Ordering of vertices in each level by (parents' order, degree) pairs. Approach: parallel partial sorting.
- **Performance:**
  - First ever distributed-memory RCM algorithm that scales up to 4096 cores on NERSC/Edison.
  - Attains up to 38x speedup on 1028 cores.

A.Azad, M. Jacquelin, **A. Buluç,** E.Ng. The Reverse Cuthill-McKee Algorithm in Distributed-Memory. IPDPS'17 (accepted)

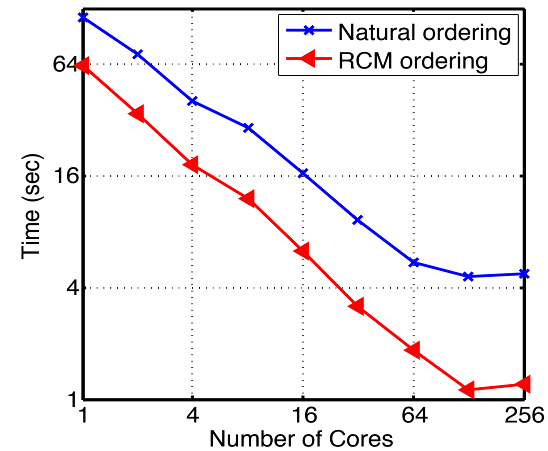

Fig2: Solving PCG in PETSc with/without RCM ordering (on thermal2 matrix)

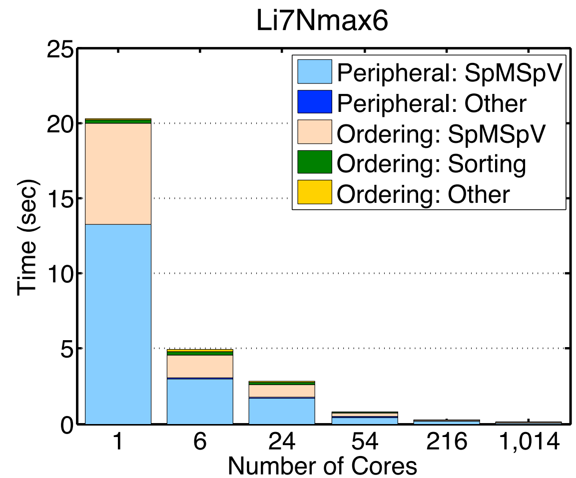

Fig2: Strong scaling of the RCM algorithm on NERSC/Edison

An work-efficient parallel algorithm for sparse matrix-sparse vector multiplication (SpMSpV)

- **Goal:** A scalable SpMSpV algorithm without doing more work on higher concurrency
- **Application:** Breadth-first search, graph matching, support vector machines, etc.
- **Algorithmic innovation:**
  - Attains work-efficiency by arranging necessary columns of the matrix into buckets where each bucket is processed by a single thread
  - Avoids synchronization by row-wise partitioning of the matrix on the fly
- **Performance:**
  - First ever work-efficient algorithm for SpMSpV that attains up to 15x speedup on a 24-core Intel Ivy Bridge processor and up to 49x speedup on a 64-core KNL processor
  - Up to an order of magnitude faster than its competitors, especially for sparser vector

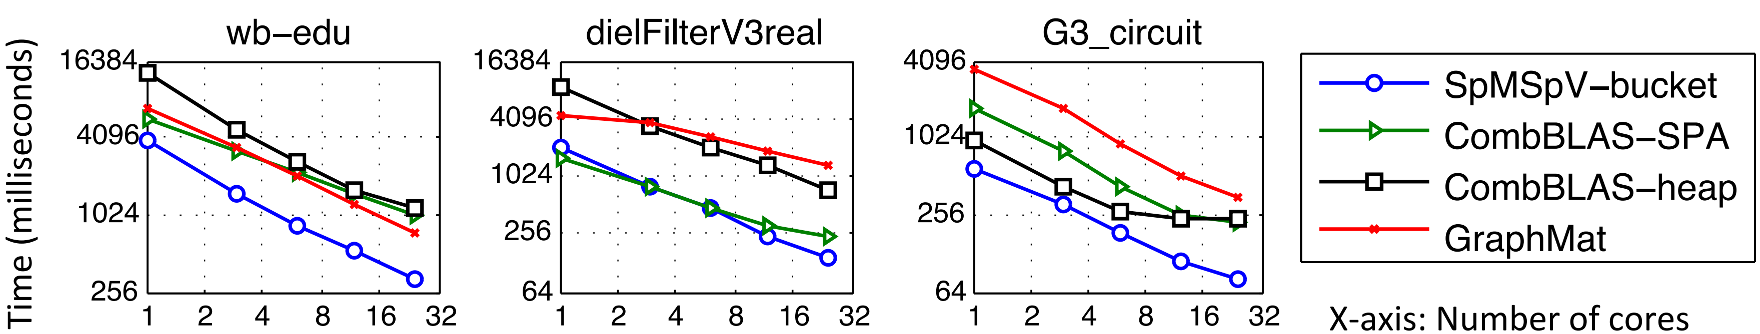

X-axis: Number of cores (Intel Ivy Bridge)

A.Azad, **A. Buluç.** A work-efficient parallel sparse matrix-sparse vector multiplication algorithm. IPDPS'17 (accepted)

GraphBLAS C API Spec (http://graphblas.org)

- **Goal:** A crucial piece of the GraphBLAS effort is to translate the mathematical specification to an actual Application Programming Interface (API) that
  i. is faithful to the mathematics as much as possible, and
  ii. enables efficient implementations on modern hardware.
- **Impact:** All graph and machine learning algorithms that can be expressed in the language of linear algebra
- **Innovation:** Function signatures (e.g. mxm, vxm, assign, extract), parallelism constructs (blocking v. non-blocking), fundamental objects (masks, matrices, vectors, descriptors), a hierarchy of algebras (functions, monoids, and semiring)

```
GrB_info GrB_mxm(GrB_Matrix                *C,      // destination
                 const GrB_Matrix           Mask,
                 const GrB_BinaryOp         accum,
                 const GrB_Semiring         op,
                 const GrB_Matrix           A,
                 const GrB_Matrix           B
               [, const Descriptor          desc]);
```

$$C(\neg M) \oplus= A^T \oplus.\otimes B^T$$

A. Buluç, T. Mattson, S. McMillan, J. Moreira, **C. Yang**. "Proposal for a GraphBLAS C API" (Working document from the GraphBLAS Signatures Subgroup)

Parallel algorithms for sparse-matrix- sparse matrix multiplication (SpGEMM)

- **Goal:** More scalable SpGEMM algorithms in shared and distributed-memory
- **Applications:** Algebraic multigrid (AMG) restriction, graph computations, quantum chemistry, data mining, interior-point optimization
- **Algorithmic innovations:** (1) Novel shared-memory kernel for in-node parallelism, (2) Split-3D-SpGEMM: an efficient implementation of communication-avoiding SpGEMM
- **Performance:** Split-3D-SpGEMM with new shared-memory kernel (red) beats old state-of-the-art (blue) by 8X at large concurrencies

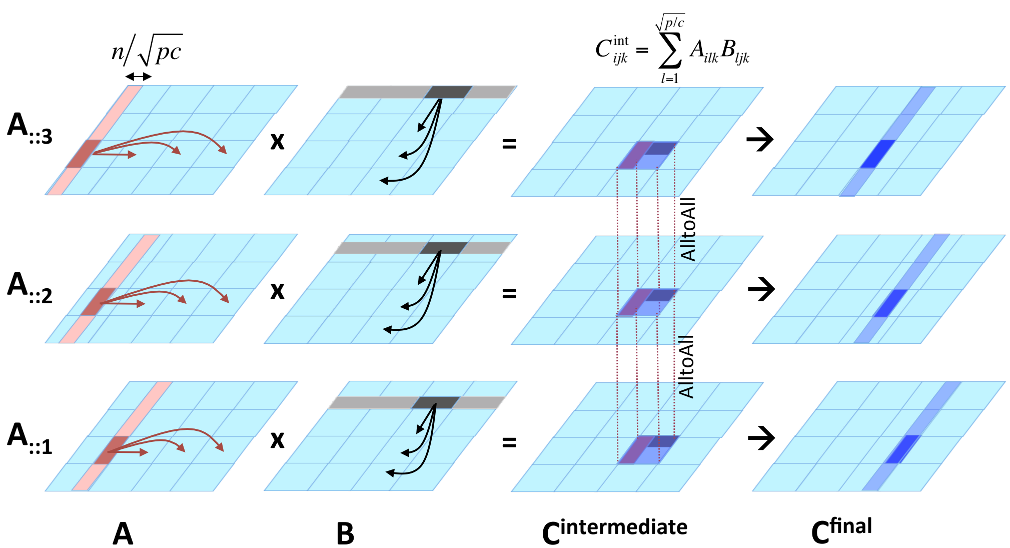

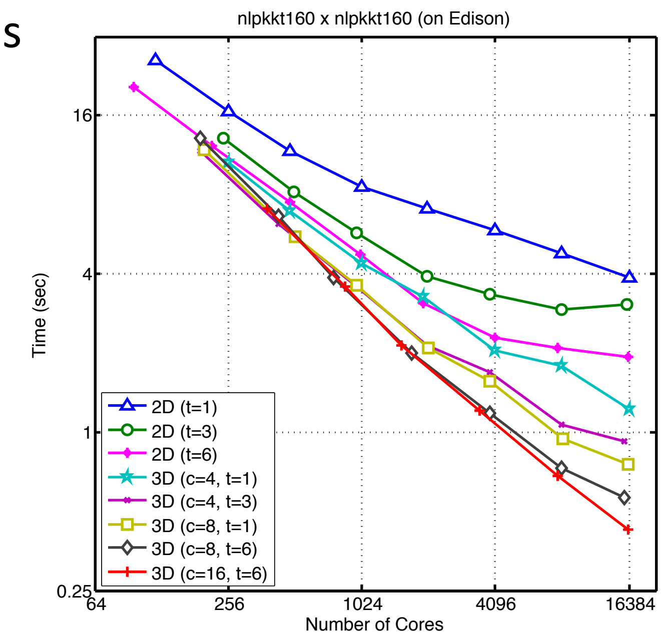

A. Azad, G. Ballard, **A. Buluç**, J. Demmel, L. Grigori, O. Schwartz, S. Toledo, S. Williams. Exploiting multiple levels of parallelism in sparse matrix-matrix multiplication. SIAM Journal of Scientific Computing (SISC), 2016.

Other Important Publications

- P. Koanantakool, **A. Azad**, **A. Buluç**, D. Morozov, S. Oh, L. Oliker, K. Yelick. *Communication-avoiding parallel sparse-dense matrix-matrix multiplica*tion. In Proceedings of the IPDPS, 2016.
- J. Kepner, P. Aaltonen, D. Bader, **A. Buluç**, F. Franchetti, J. Gilbert, D. Hutchison, M. Kumar, A. Lumsdaine, H. Meyerhenke, S. McMillan, J. Moreira, J. Owens, **C. Yang**, M. Zalewski, T. Mattson. *Mathematical foundations of the GraphBLAS*. In IEEE High Performance Extreme Computing (HPEC), 2016
- **A. Azad**, **A. Buluç,** *A matrix-algebraic formulation of distributed-memory maximal cardinality matching algorithms in bipartite graphs*. Parallel Computing, 2016
- **A. Buluç**, H. Meyerhenke, I. Safro, P. Sanders, C. Schulz: *Recent Advances in Graph Partitioning*. Algorithm Engineering, 2016
- **V. Strnadová-Neeley, A. Buluç**, J. Gilbert, L. Oliker, W. Ouyang: *LiRa: A New Likelihood-Based Similarity Score for Collaborative Filtering, CoRR abs/1608.08646*