



Berkeley
UNIVERSITY OF CALIFORNIA

Sparse Matrices for Scaling Data-Intensive Problems on Distributed-Memory Systems

Aydın Buluç

Computational Research Division, LBNL
EECS Department, UC Berkeley

SIAM CSE'21
March 1, 2021

Schedule

GraphBLAS: Tools, Algorithms, and Applications (Monday)

9:45 am: Aydin Buluc (Berkeley)
10:05 am: Tim Davis (Texas A&M)
10:25 am: Scott McMillan (CMU SEI)
10:45 am: Roger Pearce (LLNL)
11:05 am: Michel Pelletier (Graphegon)

2:15 pm: Ariful Azad (Indiana Univ.)
2:35 pm: Federico Busato (NVIDIA)
2:55 pm: Andrey Prokopenko (ORNL)
3:15 pm: Ben Brock (Berkeley)
3:35 pm: Vijay Thakkar (GATech)

Building Graph Algorithms with the GraphBLAS (Minitutorial)

Tuesday 2:15pm- 3:55pm (Part 1)

Wednesday 2:15pm- 3:55pm (Part 2)

The GraphBLAS forum

Standards for Graph Algorithm Primitives

Tim Mattson (Intel Corporation), David Bader (Georgia Institute of Technology), Jon Berry (Sandia National Laboratory), Aydin Buluc (Lawrence Berkeley National Laboratory), Jack Dongarra (University of Tennessee), Christos Faloutsos (Carnegie Mellon University), John Feo (Pacific Northwest National Laboratory), John Gilbert (University of California at Santa Barbara), Joseph Gonzalez (University of California at Berkeley), Bruce Hendrickson (Sandia National Laboratory), Jeremy Kepner (Massachusetts Institute of Technology), Charles Leiserson (Massachusetts Institute of Technology), Andrew Lumsdaine (Indiana University), David Padua (University of Illinois at Urbana-Champaign), Stephen Poole (Oak Ridge National Laboratory), Steve Reinhardt (Cray Corporation), Mike Stonebraker (Massachusetts Institute of Technology), Steve Wallach (Convey Corporation), Andrew Yoo (Lawrence Livermore National Laboratory)

Abstract-- It is our view that the state of the art in constructing a large collection of graph algorithms in terms of linear algebraic operations is mature enough to support the emergence of a standard set of primitive building blocks. This paper is a position paper defining the problem and announcing our intention to launch an open effort to define this standard.

“If you want to go fast, go alone. If you want to go far, go together.” -- unknown

<https://graphblas.github.io/>



GraphBLAS C API Specification

- **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]);
```

$$C(-M) \oplus = A^T \oplus \cdot \otimes B^T$$

Pattern 1: Sparse matrix times sparse vector (SpMSpV)



Single-source traversal:


BFS, connected components, matching, ordering, etc.

`GrB_mxv(y, p, <semiring>, A, x, <desc>)`

A: sparse adjacency matrix

x: sparse input vector (previous frontier)

p: mask (already discovered vertices)



Pattern 2: Sparse matrix times sparse matrix (SpGEMM)

Multi-source traversal:


Ex: multi-source BFS, betweenness centrality, triangle counting*, Markov clustering*

`GrB_mxm(Y, P, <semiring>, A, X, <desc>)`

A: sparse adjacency matrix

X: sparse input matrix (previous frontier), n-by-b where b is the #sources


P: mask (already discovered vertices), multi-vector version of p from previous slide






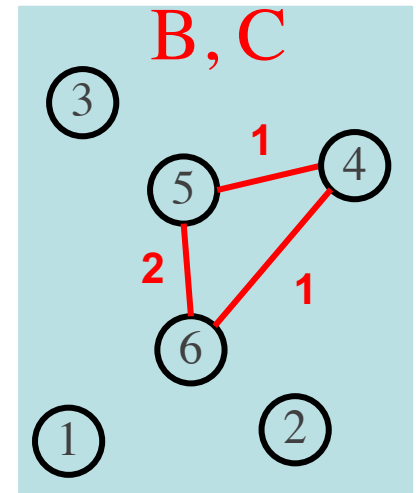
Pattern 2: Sparse matrix times sparse matrix (SpGEMM)

Triangle counting is also multi-source (in fact, all sources) traversal: It just stops after one traversal iteration only, discovering all wedges

GrB_mxm(C, A, <semiring>, L, U, <desc>)



$$\begin{aligned}
 A &= L + U && (\text{hi} \rightarrow \text{lo} + \text{lo} \rightarrow \text{hi}) \\
 L \times U &= B && (\text{wedge, low hinge}) \\
 A \wedge B &= C && (\text{closed wedge}) \\
 \text{sum}(C)/2 &= && \mathbf{4 \text{ triangles}}
 \end{aligned}$$



Pattern 2: Sparse matrix times sparse matrix (SpGEMM)


Markov clustering is also multi-source (in fact, all sources) traversal:

It alternates between SpGEMM and element-wise or column-wise pruning

```
GrB_mxm(C, GrB_NULL, <semiring>, A, A, <desc>)
```

A: sparse normalized adjacency matrix

C: denser (but still sparse) pre-pruned matrix for next iteration



At each iteration:

Step 1 (Expansion): Squaring the matrix while pruning (a) small entries, (b) denser columns

Naïve implementation: sparse matrix-matrix product (SpGEMM), followed by column-wise top-K selection and column-wise pruning

Step 2 (Inflation): taking powers entry-wise

Pattern 3: Sparse matrix times tall-skinny dense matrix (SpMM)

Feature aggregation from neighbors:

Used in Graph neural networks, graph embedding, etc.

```
GrB_mxm(W, GrB_NULL, <semiring>, A, H, <desc>)
```


A: sparse adjacency matrix, n-by-n

H: input dense matrix, n-by-f where $f \ll n$ is the feature dimension


W: output dense matrix, new features

O(f) feature vector


3.2 5.4 ... 1.3



$W =$




A^T




H

Graph Neural Networks (GNNs)



GNNs are finding success in many challenging scientific problems that involve interconnected data.

Interdependencies between samples (nodes of the graph) make stochastic gradient non-trivial without graph sampling



- GNNs are computationally intensive to train. Distributed training need to scale to large GPU/node counts despite challenging sparsity.
- CAGNET (Communication-Avoiding Graph Neural nETworks) full gradient descent to avoid inaccurate (and expensive) graph sampling

<https://github.com/PASSIONLab/CAGNET/>

GNN Training

- Each node is initialized with a feature vector
 - H^0 has initial feature vector per node ($n \times f$)
- Each node aggregates vectors of its neighbors, applies a weight
- Each layer computes gradients

for $i = 1 \dots E$

$$A \in n \times n$$

for $l = 1 \dots L$

$$Z^l = A^T * H^{l-1} * W^l$$

$$H^l \in n \times f^l$$

$$H^l = \sigma(Z^l)$$

...

for $l = L-1 \dots 1$

$$G^l = A * G^{l+1} * (W^{l+1})^T \odot \sigma'(Z^l)$$

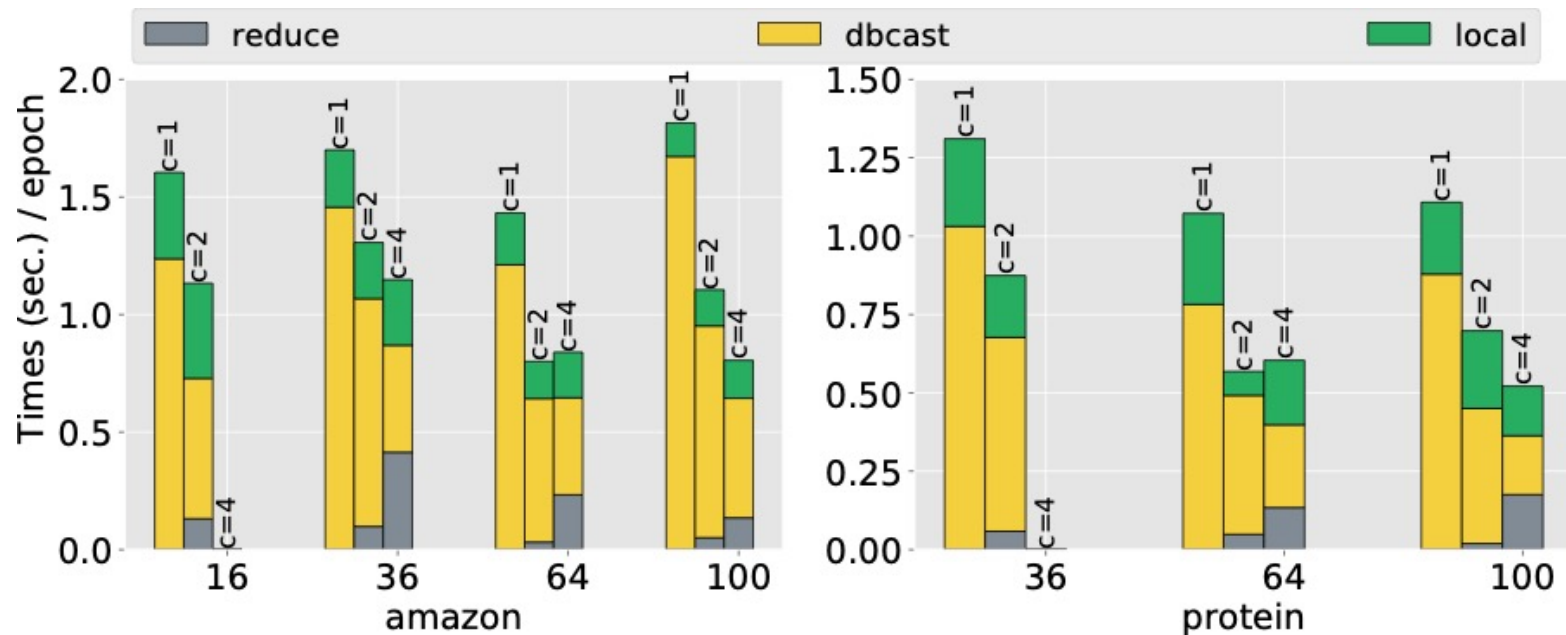
$$G^l \in n \times f^l$$

$$dH/dW = (H^{l-1})^T * A * G^l$$

$$W^l \in f^{l-1} \times f^l$$

- A is sparse and $f \ll n$, so the main workhorse is SpMM (sparse matrix times tall-skinny dense matrix)

Communication avoidance (CA) In GNN Training




- Scales with both P (GPUs – x axis) and c (replication layers in CA algorithms)
- This is 1 GPU/node on Summit (all GPUs per node results in paper)
- Expect to scale with all GPUs / node with future architectures (e.g. Perlmutter)

SpGEMM for DNA read overlapping

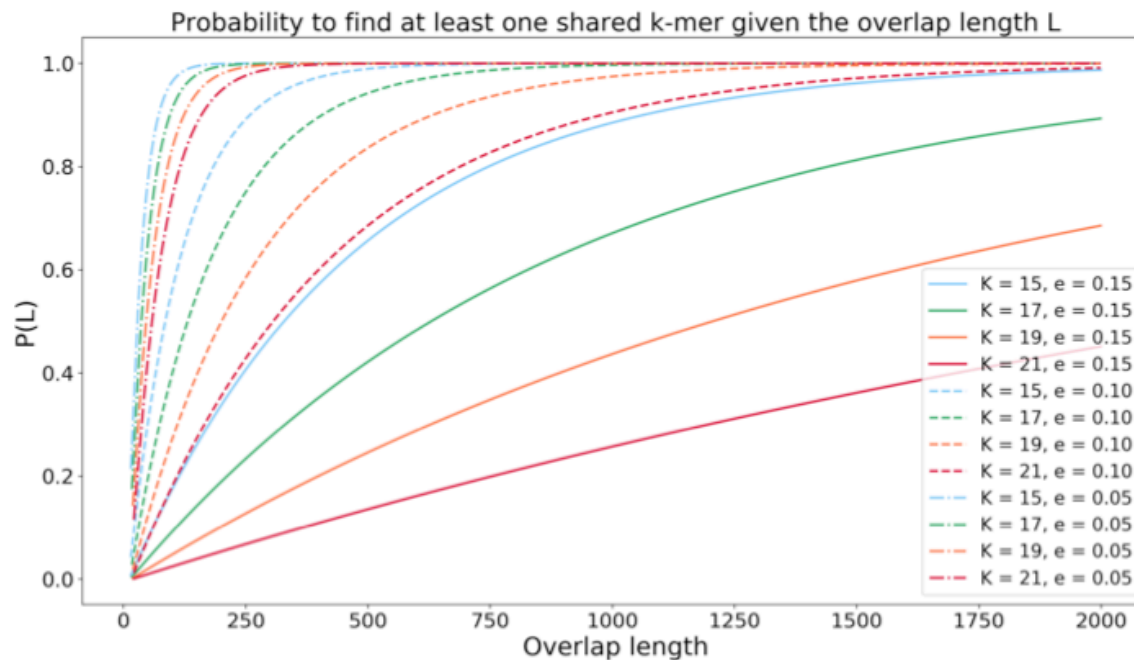
- **Long reads** from PacBio and Oxford Nanopore have the potential to revolutionize de-novo assembly
- **Overlap-Consensus-Layout** paradigm is more suitable than de Bruijn graph paradigm.
- **Overlapping** is the most computationally expensive step.

Overlap-Layout-Consensus



SpGEMM for DNA read overlapping

- We need to quickly determine pairs of reads that are *likely to* overlap, without resorting to $O(n^2)$ comparisons
- If two reads do not share any subsequence of length k (aka a k -mer) for a reasonably short k , then they are unlikely to overlap



SpGEMM for DNA read overlapping

A matrix

	k_1	k_2	k_3	k_4	k_5	k_6
r_1	●				●	
r_2						●
r_3	●		●			
r_4			●		●	
r_5					●	
r_6				●		

- Suppose you have counted k -mers and only retained *reliable* k -mers
- Now you can generate this **read-by-kmer** sparse matrix **A**
- These are all linear time computations so far

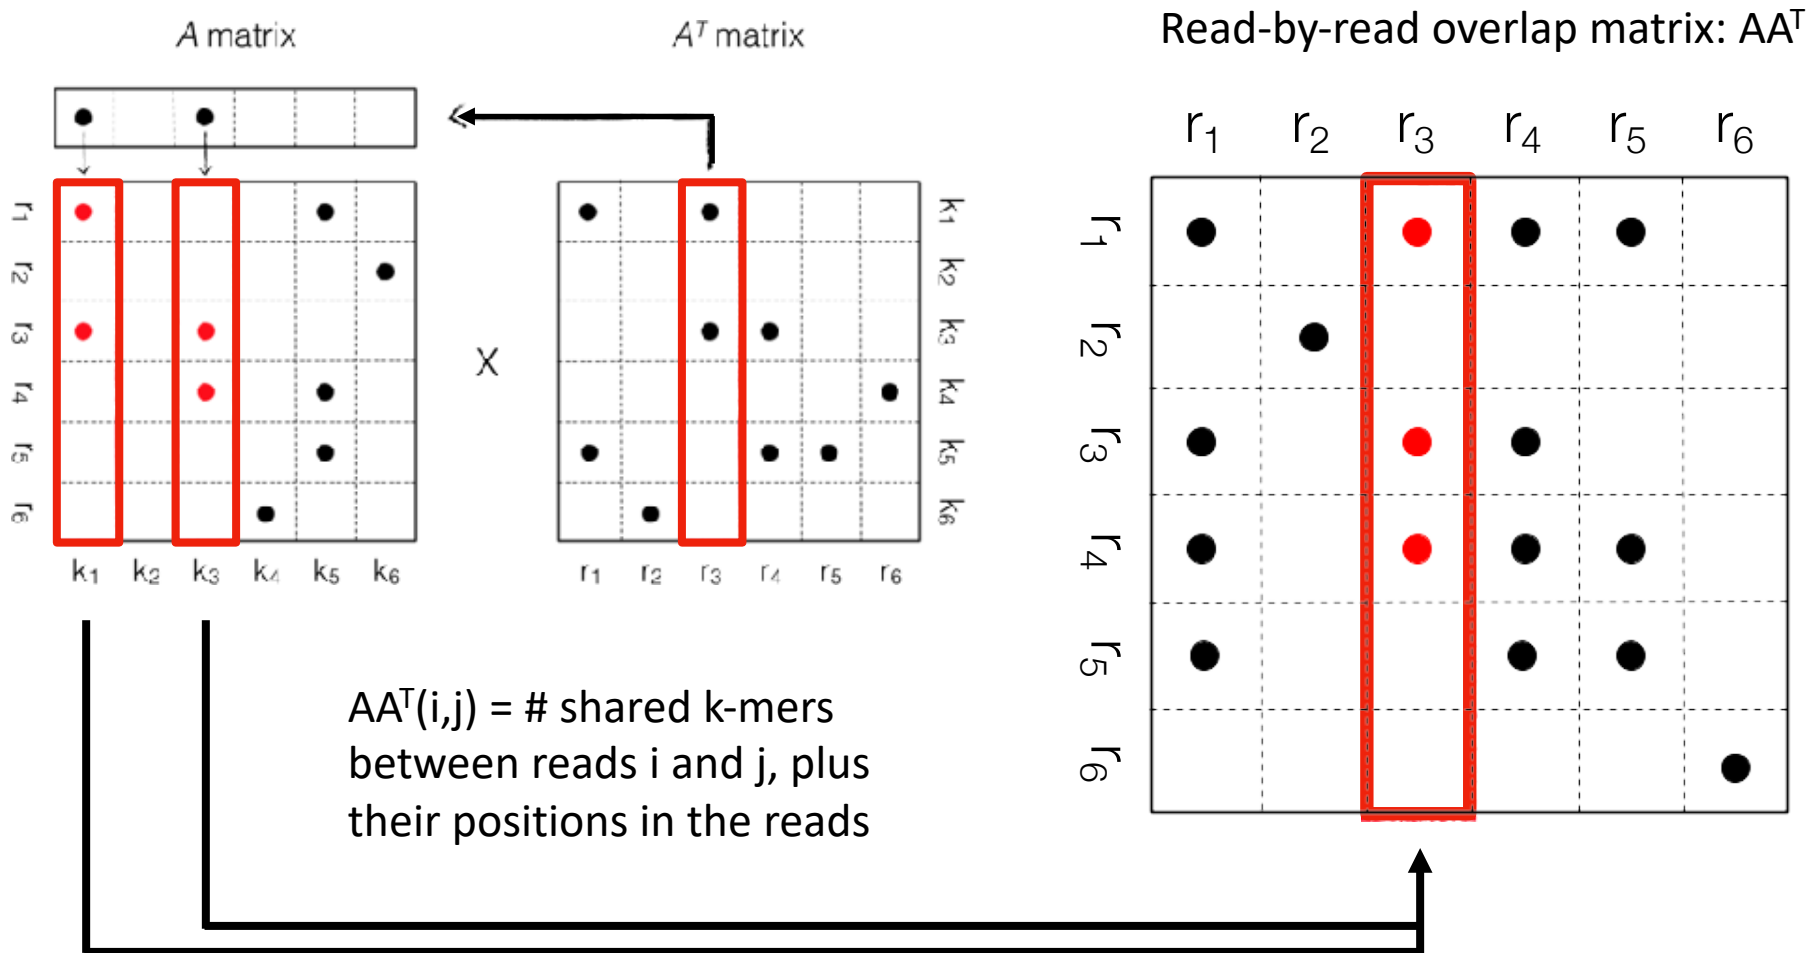
$r_i = i^{\text{th}}$ read

$k_j = j^{\text{th}}$ reliable k -mer

$A(i,j) =$ presence of j^{th} reliable k -mer in i^{th} read, plus its position

Giulia Guidi, Marquita Ellis, Daniel Rokhsar, Kathy Yelick, Aydın Buluç, BELLA: Berkeley Efficient Long-read to Long-Read Overlapper and Aligner, *Biorxiv*.

SpGEMM for DNA read overlapping

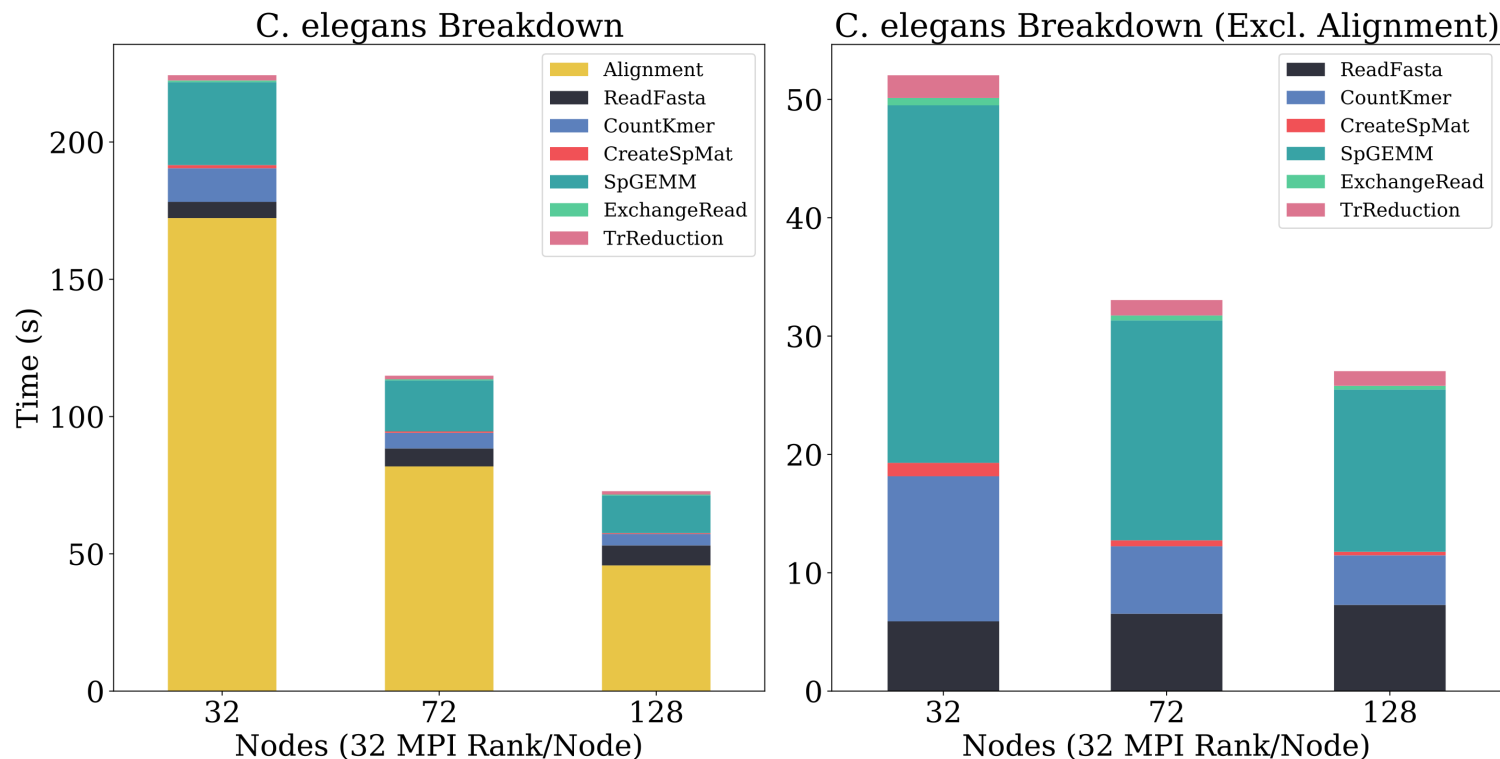


Use any fast SpGEMM algorithm, as long as it runs on *arbitrary semirings*

diBELLA.2D performance results

diBELLA.2D: distributed-memory version of BELLA on 2D process grid
Performs *overlap detection* plus *transitive reduction* (overlap to string graph)

<https://github.com/PASSIONLab/diBELLA.2D>



Giulia Guidi, Oguz Selvitopi, Marquita Ellis, Leonid Oliker, Katherine Yelick, Aydin Buluç. Parallel String Graph Construction and Transitive Reduction for De Novo Genome Assembly. *IPDPS 2021*

SpGEMM for many-to-many protein alignment

PASTIS (<https://github.com/PASSIONLab/PASTIS>) does distributed many-to-many protein sequence similarity search using sparse matrices

Introduce new sparse matrix **S**


Contains substitution information

Each entry has **substitution cost**

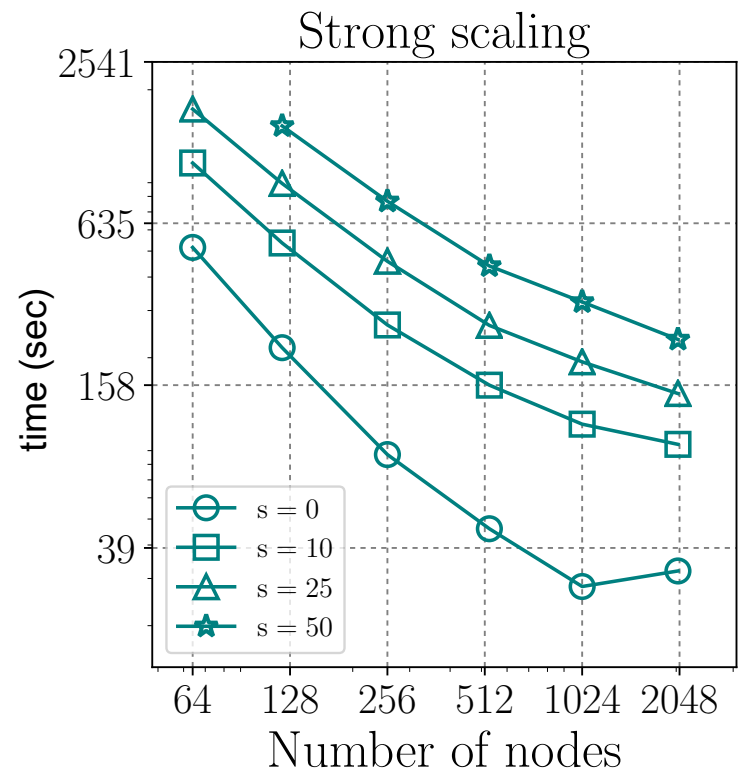
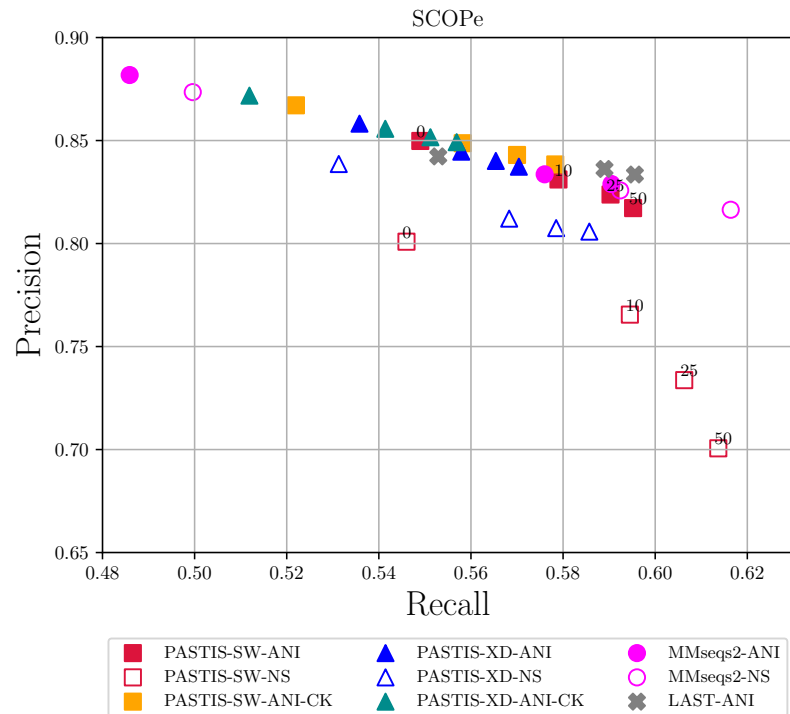
Exact k-mers $\rightarrow C=AA^T$

Substitute k-mers $\rightarrow C=ASA^T$

New semiring



PASTIS performance and accuracy



- *Protein similarity search* is the first and most time-consuming step in discovering protein families (proteins evolved from a common ancestor and who likely have the same function)
- *Protein family identification* is a key step in protein function discovery and taxonomic assignment of newly sequenced organisms

Acknowledgments

Ariful Azad, Ben Brock, Tim Davis, Saliya Ekanayake, Marquita Ellis, John Gilbert, Giulia Guidi, Jeremy Kepner, Nikos Krypides, Tim Mattson, Scott McMillan, Jose Moreira, Lenny Oliker, John Owens, Georgios Pavlopoulos, Dan Rokhsar, Oguz Selvitopi, Yu-Hang Tang, Alok Tripathy, Carl Yang, Kathy Yelick.

- My Research Team: <http://passion.lbl.gov>
- Our Youtube Channel: <http://shorturl.at/lpFRY>
- The GraphBLAS Forum: <http://graphblas.org>