

# Graph Analytics in Arkouda



**David A. Bader**

 [@Prof\\_DavidBader](https://twitter.com/Prof_DavidBader)

<http://www.cs.njit.edu/~bader>



# David A. Bader

Distinguished Professor and  
Director, Institute for Data Science

- IEEE Fellow, SIAM Fellow, AAAS Fellow
- IEEE Computer Society Sidney Fernbach Award
- Recent Service:
  - White House's National Strategic Computing Initiative (NSCI) panel
  - Computing Research Association Board
  - NSF Advisory Committee on Cyberinfrastructure
  - Council on Competitiveness HPC Advisory Committee
  - IEEE Computer Society Board of Governors
  - IEEE IPDPS Steering Committee
  - Editor-in-Chief, ACM Transactions on Parallel Computing
  - Editor-in-Chief, IEEE Transactions on Parallel and Distributed Systems
- Over \$185M of research awards
- 300+ publications,  $\geq 11,000$  citations, h-index  $\geq 61$
- National Science Foundation CAREER Award recipient
- Directed: Facebook AI Systems
- Directed: NVIDIA GPU Center of Excellence, NVIDIA AI Lab (NVAIL)
- Directed: Sony-Toshiba-IBM Center for the Cell/B.E. Processor
- Founder: Graph500 List benchmarking "Big Data" platforms
- Recognized as a "RockStar" of High Performance Computing by InsideHPC in 2012 and as HPCwire's People to Watch in 2012 and 2014.





Launched in **July 2019**, with inaugural director

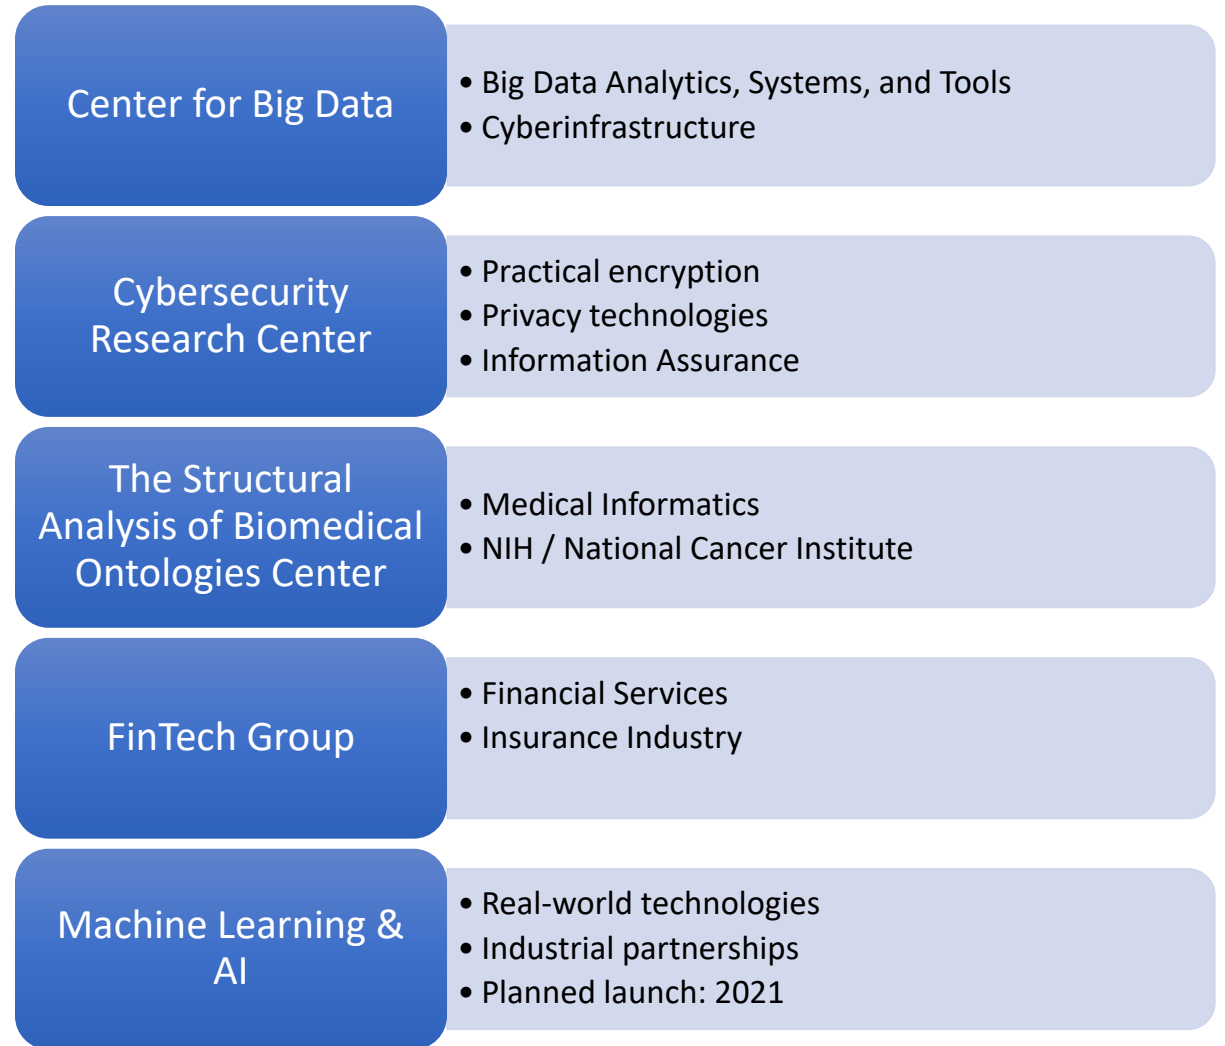
**David A. Bader**

(~35 faculty in current centers)

**NJIT Data Science Seminar Series**

Wednesday's 4pm ET

<https://njit-institute-for-data-science.eventbrite.com/>



# Upcoming NJIT Data Science Seminars

<https://datascience.njit.edu/events>

- October 6: Gordon Bell
- October 20: Jack Dongarra
- October 27: Larry Smarr
- November 3: Laura Haas



<https://www.youtube.com/c/NJITInstituteforDataScience/>



## Joint work with

- Dr. Zihui Du (NJIT)
- Oliver Alvarado Rodriguez (PhD student)
- Joseph Patchett (MS student)

# High Performance Algorithms for Interactive Data Science at Scale

(PI: Bader) 3/2021 – 2/2022, NSF CCF-2109988



A real-world challenge in data science is to develop interactive methods for quickly analyzing new and novel data sets that are potentially of massive scale. This award will design and implement fundamental algorithms for high performance computing solutions that enable the interactive large-scale data analysis of massive data sets.

This project focuses on these three important data structures for data analytics:

- 1) suffix array construction,
  - 2) 'treap' construction, and
  - 3) distributed memory join algorithms,
- useful for analyzing large scale strings, implementing random search in large string data sets, and generating new relations, respectively.

To evaluate and show the effectiveness of the proposed algorithms, these algorithms will be implemented in and contribute to an open source NumPy-like software framework that aims to provide productive data discovery tools on massive, dozens-of-terabytes data sets by bringing together the productivity of Python with world-class high performance computing.

# Institute for Data Science Aims to Democratize Supercomputing With NSF Grant

Written by: Evan Koblenz

Published: Wednesday, March 17, 2021



New algorithms from at NJIT can make supercomputer power available to almost anyone

Ordinary people could soon have greater ability to analyze massive amounts of information, based on new algorithms and software tools being designed at NJIT, intended to simplify

<https://news.njit.edu/institute-data-science-aims-democratize-supercomputing-nsf-grant>

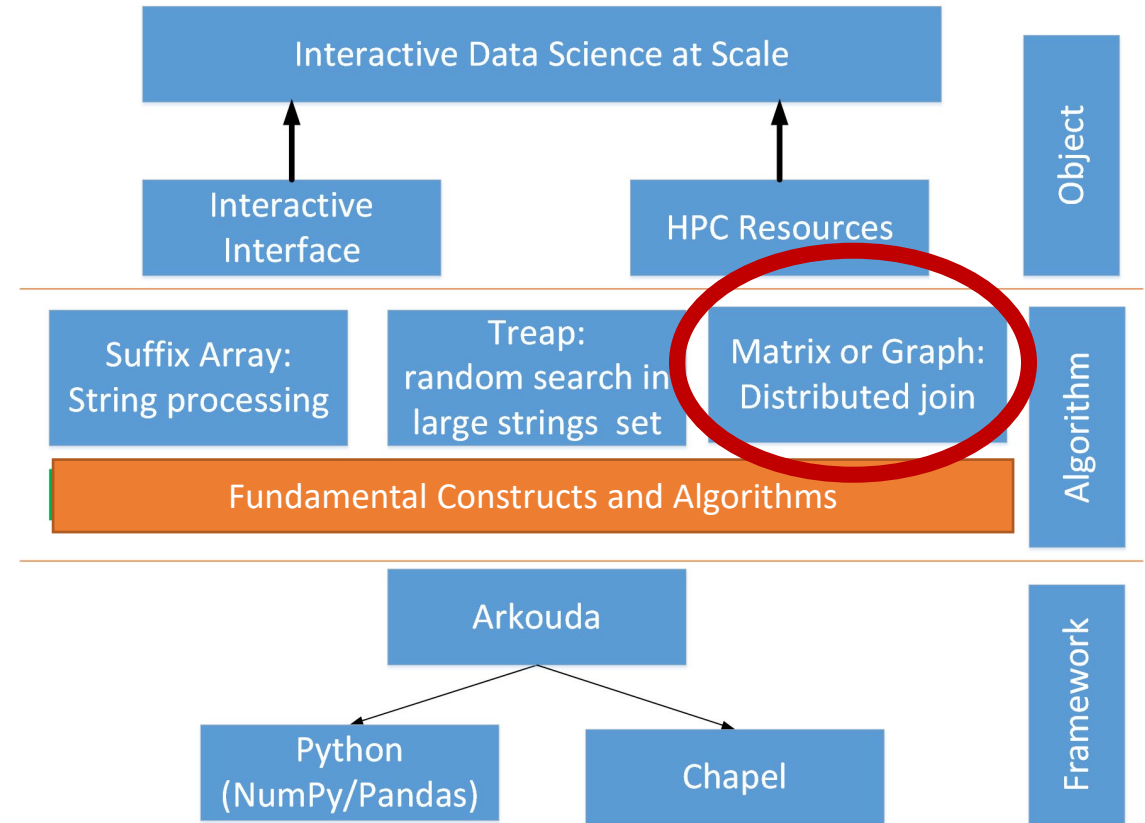
6 October 2021

David A. Bader

7

# Interactive Data Science at Scale

- Objective:
  - One-stop solution for non-HPC experts with massive data sets.
- Developmental focus:
  - Data structures and algorithms for graph and other problems.
- Framework:
  - Arkouda

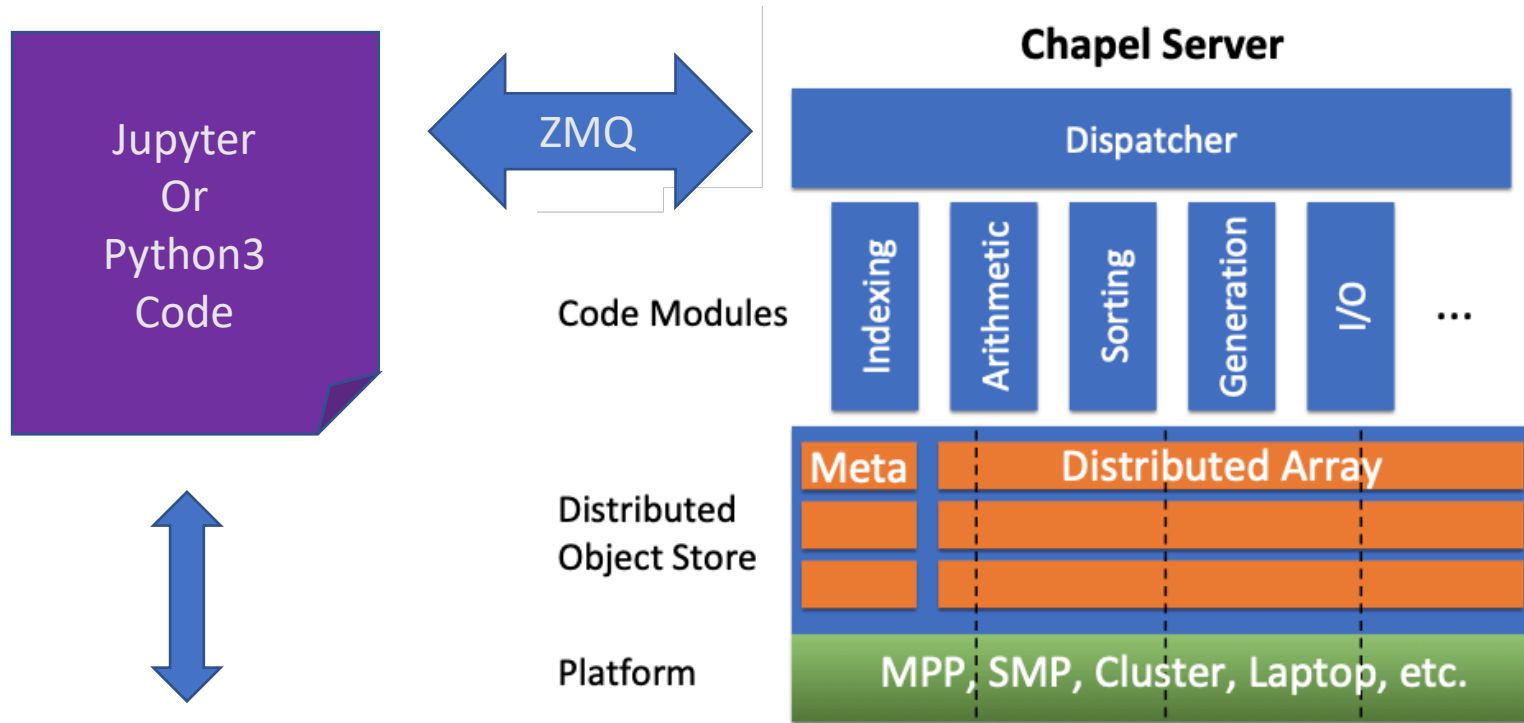




# Productivity + Performance



# Typical Environment Set-Up




## Python3 Implementation:

- Parray class
- Rely on Python to reduce complexity
- Integrate with and use NumPy

## Server Implementation:

- High-level language with C-comparable performance
- Great parallelism handling
- Great distributed array support
- Portable code: laptop --> HPC



Where can I get it?:

Image: <https://chapel-lang.org/CHI UW/2020/Reus.pdf>

Software: <https://github.com/mhmerrill/arkouda>

Our Contribution: <https://github.com/Bader-Research/arkouda/tree/streaming>

# Large-Scale Graph Analytics in Arkouda

- Graph Data Structure (Double-Index: DI)
  - Edge-oriented partitioning scheme
  - [Du, Alvarado Rodriguez, Bader 2021]
- Breadth-First Search
  - A high-level and low-level distributed and parallel implementation of breadth-first search (BFS) with reverse Cuthill-McKee (RCM) relabeling heuristic if needed.
  - [Alvarado Rodriguez, Du, Bader 2021]
- Suffix Array
  - Highly-productive: Python users can process massive collections of strings in almost the same way as small strings
  - [Du, Alvarado Rodriguez, Bader 2021]
- K-Truss
  - Design an optimized multi-locale (distributed) parallel k-Truss algorithm that can significantly improve the performance
  - [Patchett, Du, Bader, 2021]

# Double-Index Graph Data Structure

- Advantage

- $O(1)$  time complexity
  - Locate specific vertex from given edge ID
  - Locate adjacency list from given vertex ID

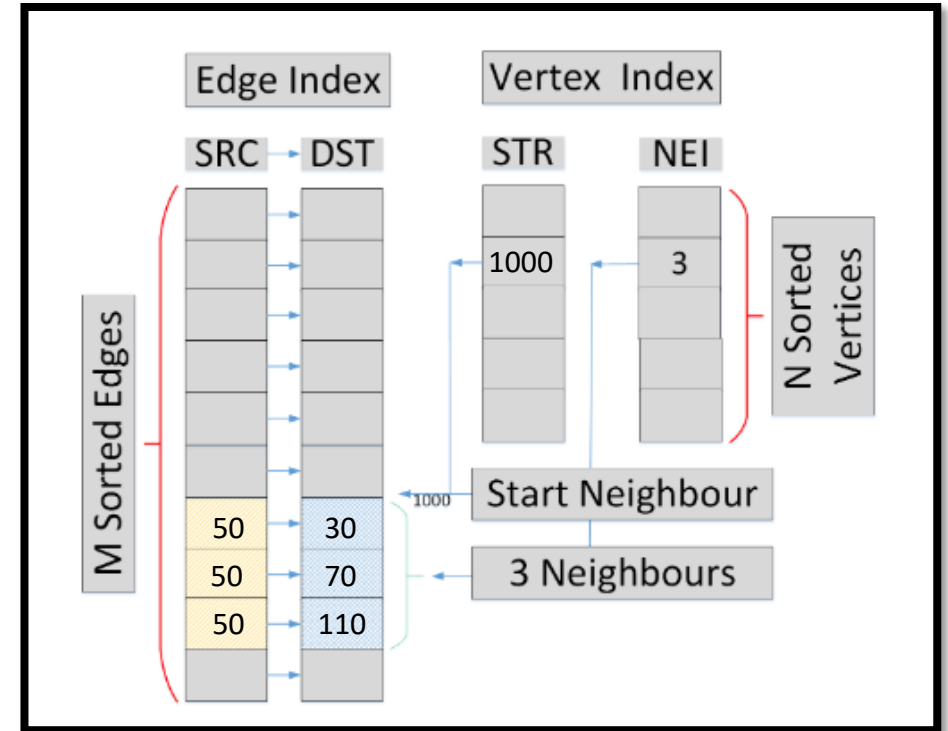
- Comparison with CSR

- Similarity

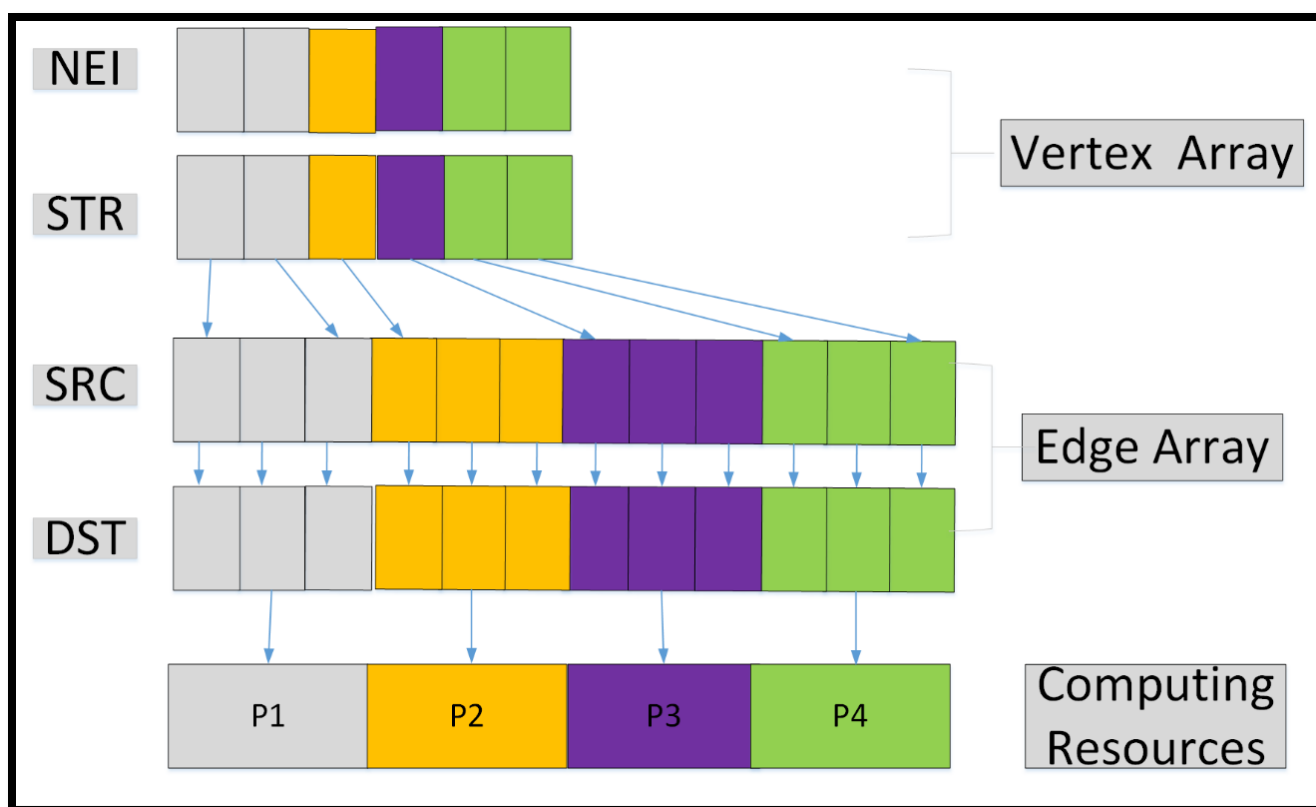
- Value array  $\leftrightarrow$  SRC/DST, array size is NNZ
- Column index  $\leftrightarrow$  STR, array size is  $|V|$
- Row index  $\leftrightarrow$  NEI, array size is  $|V|+1$  and  $|V|$

- Difference

- We can search from edge ID to vertex ID, CSR cannot.
- We need a bit more memory (another NNZ array) than CSR



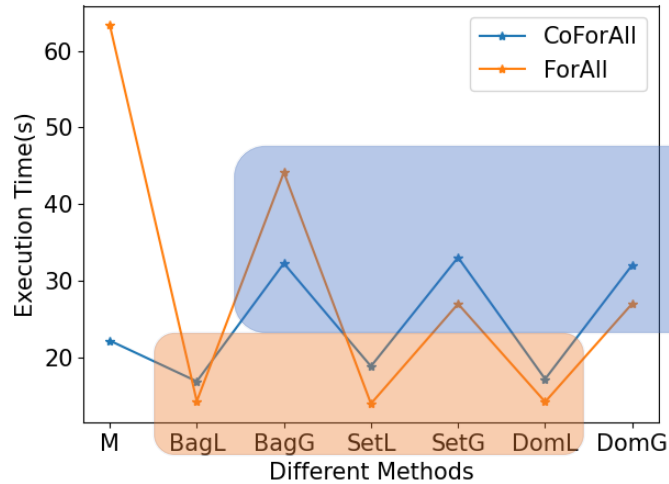
# Support Edge-Oriented Graph Partition



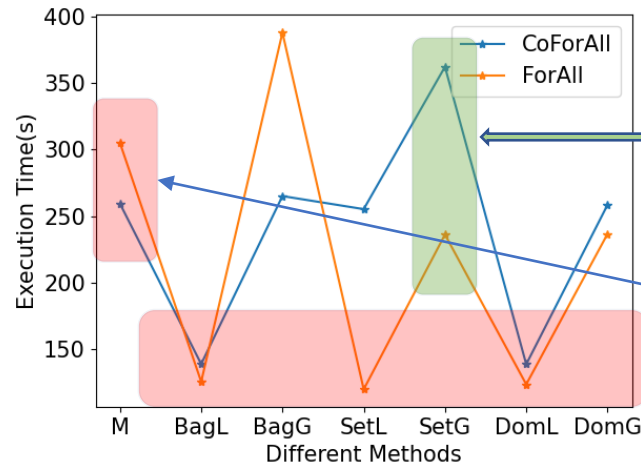
- Imbalance in vertex's degree
  - Power law degree distribution of real-world graphs
    - One vertex can have very different number edges
- Edge oriented partition
  - Edge array is much larger than vertex array
    - Align vertices to corresponding edge
    - load balancing

# Breadth-First Search Algorithms in Arkouda

- High level data structure
  - Distributed bag (Bag), set (Set), and domain (Dom)
  - L is the distributed parallel computing version. G is the shared memory computing version.
  - M is the low-level version
- Parallel construct
  - forall/coforall



Delaunay 17



Delaunay 20

redundant calculations without idle threads

no redundant calculation with idle threads

Different parallel constructs can affect performance

High-level algorithm can compete with low-level algorithm under the same algorithm framework

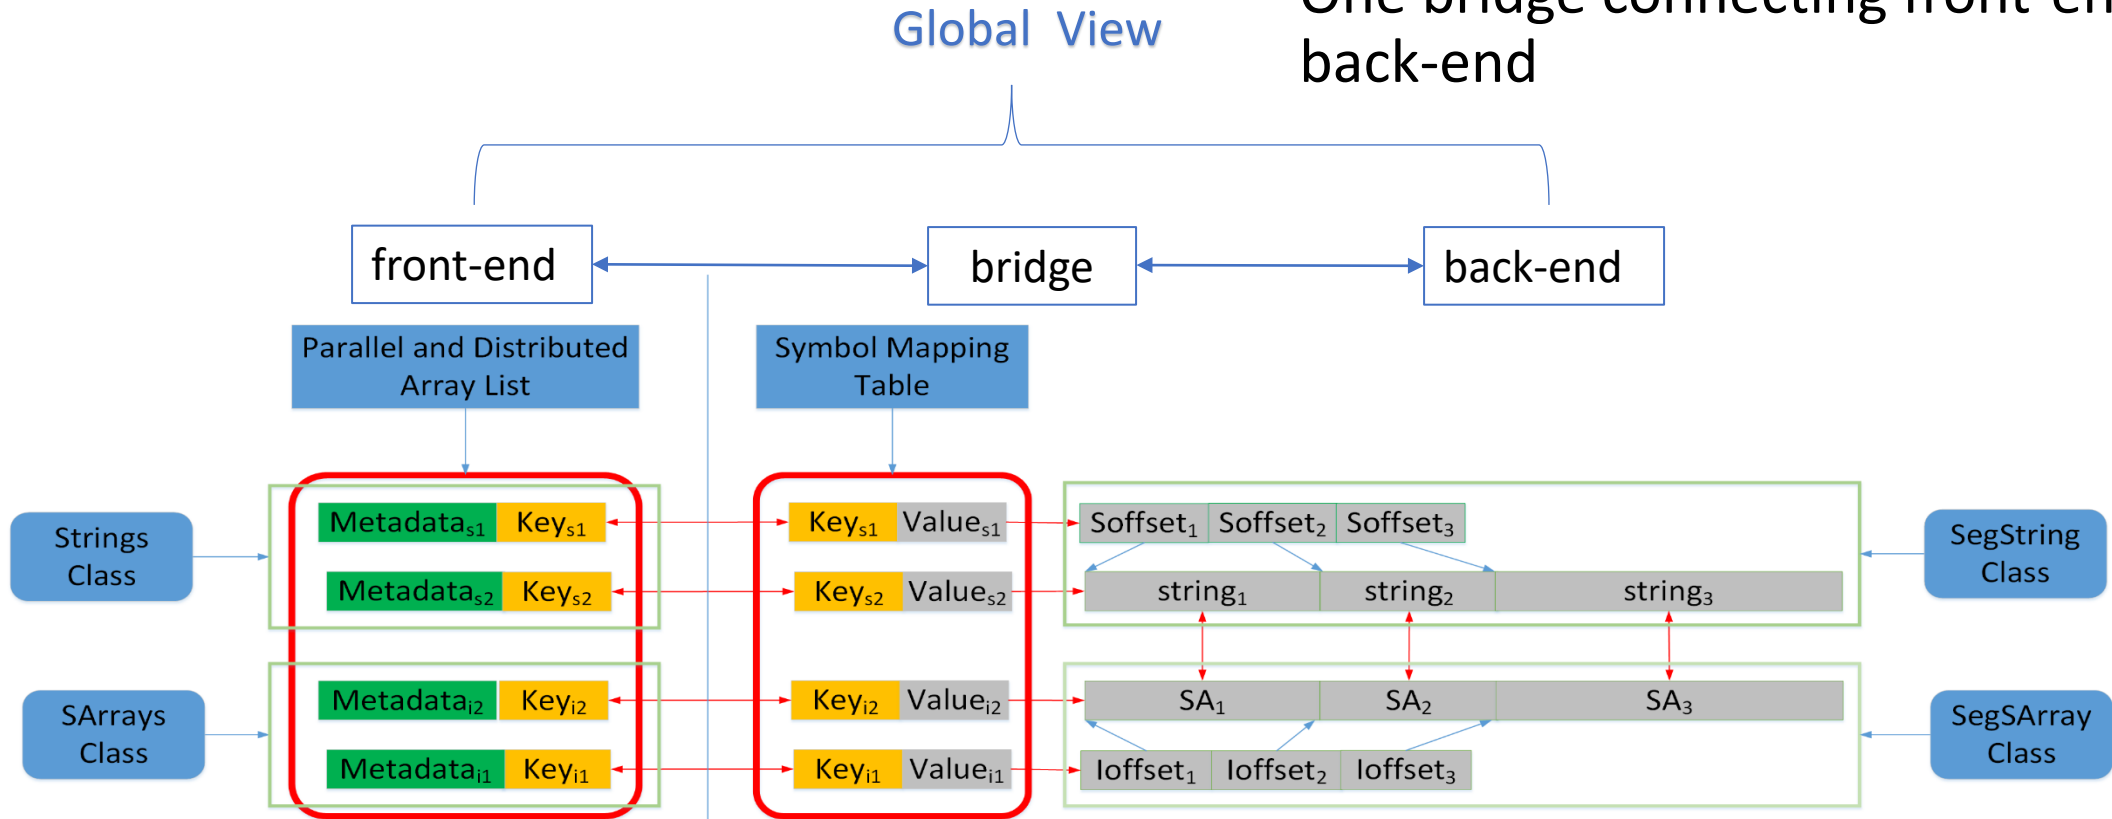
# Suffix arrays for a group of strings

- Example

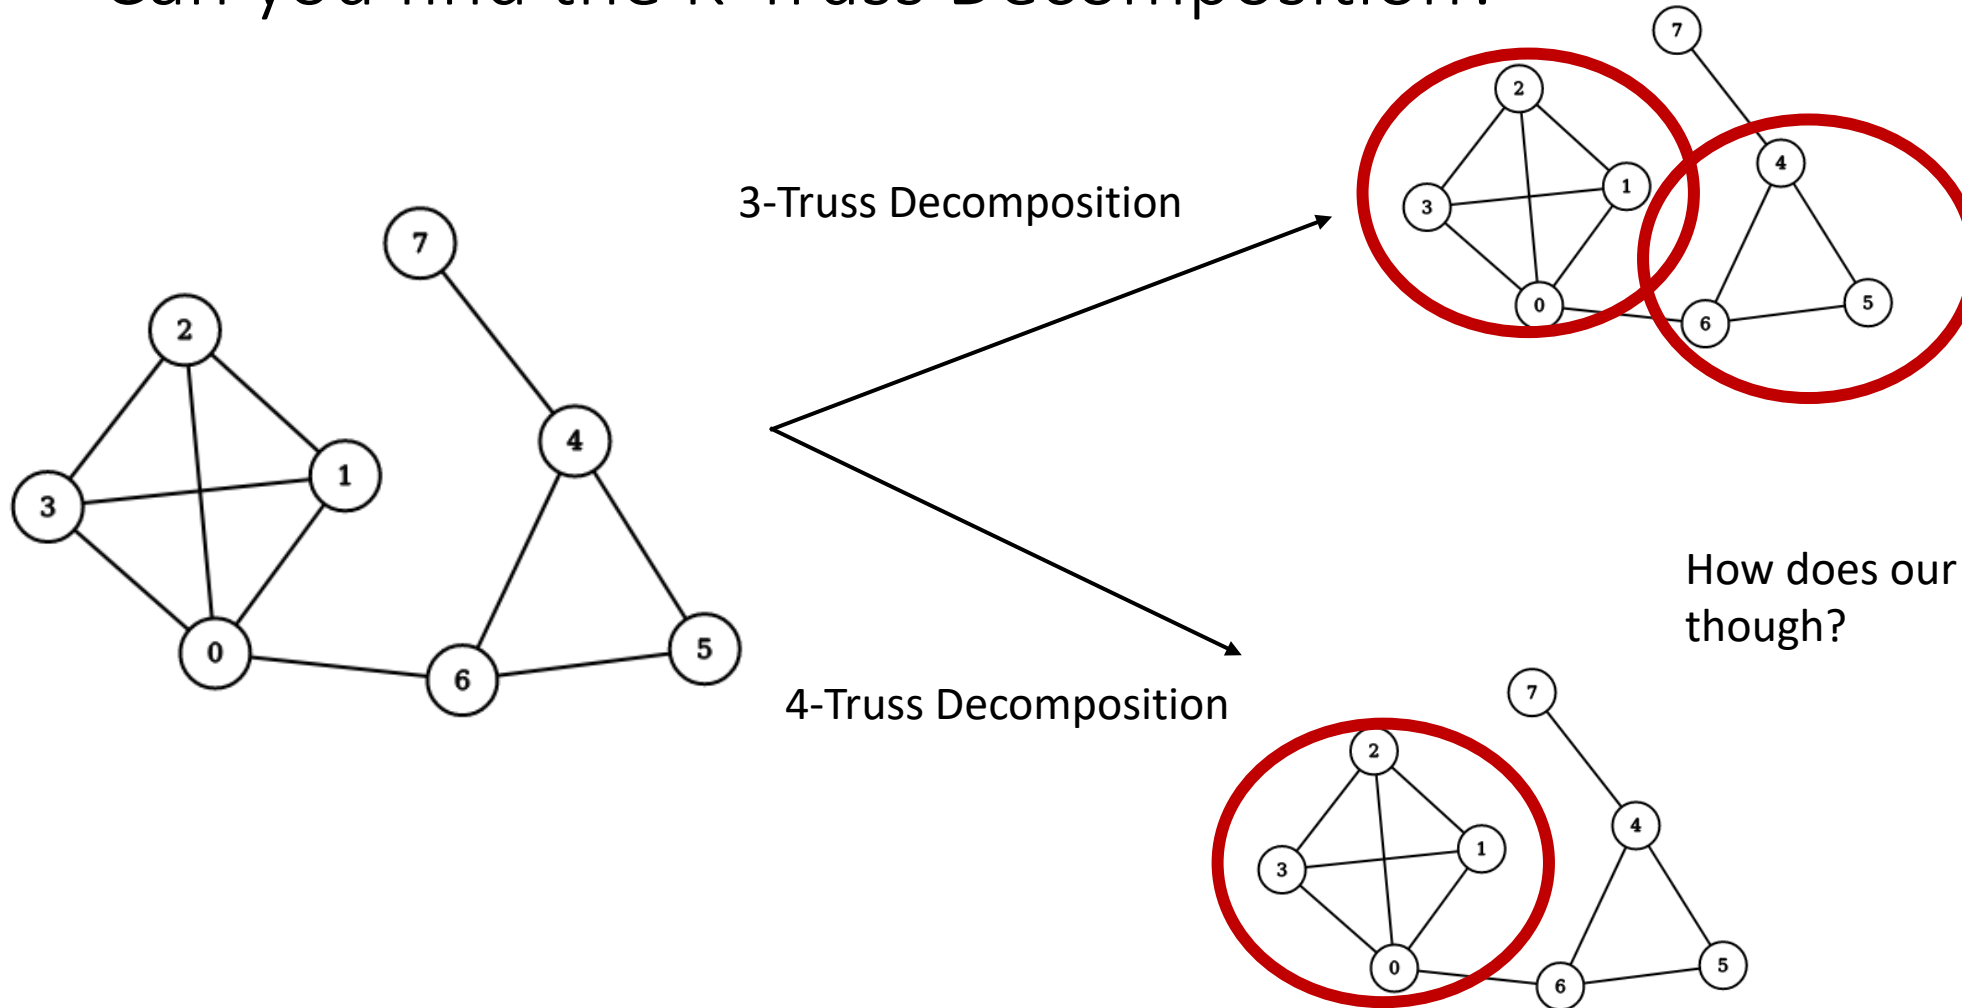
- Input strings  $SG = \{ "cba", "5986" \}$
- Output suffix arrays  $SA = \{ [3, 2, 1], [1, 4, 3, 2] \}$
- Value/offset arrays at back-end

- GLB (Global-Local-Bridge) design

- One global view
- Two local implementations
- One bridge connecting front-end and back-end



# Can you find the K-Truss Decomposition?



How does our k-Truss work though?



# High Level k-Truss Algorithm

```
1  $G_{prime} = preprocess(G)$ 
  /*  $G$  is the graph defined by our double
   index data structure */
2 while A change has been made do
3    $TriCount = T(G_{prime})$ 
4   //Compute the Initial triangle counts
5   Initialize current frontier  $SetCurF$ 
6   Initialize next frontier  $SetNextF$ 
7   if Edge  $m$  has  $TriCount[m] < k - 2$  then
8     Assign  $m$  to  $SetCurF$ 
9     Remove  $m$  from  $G_{prime}$ 
10  end
11  coforall (loc in Locales ) do
12    while  $SetCurF$  is not empty do
13      forall Edges  $\langle u, v \rangle$  in the  $SetCurF$  do
14        forall Adjacent edges do
15          if A triangle exists then
16            Decrement  $TriCount$  for the
17            edge
18            if Edge has  $TriCount < k-2$ 
19            then
20              Add the edge to  $SetNextF$ 
21              Remove the edge from
22               $G_{prime}$ 
23            end
24          end
25        end
26      end
27       $SetCurF \Leftarrow SetNextF$ 
28      clear  $SetNextF$ 
29    end
30  end
31  return  $G_{prime}$ 
```

Some edges aren't part of any triangles and can be removed

The frontier maintains edges to be removed

Use DI graph structure to quickly assess adjacency

# Acknowledgement

We appreciate the help from the Chapel and Arkouda teams including: Brad Chamberlain, Elliot Joseph Ronaghan, Engin Kayraklioglu, David Longnecker, Michael Merrill, and Bill Reus.

This research was funded in part by NSF grant number CCF-2109988.