

Reproducibility in Parallel Graph Algorithms

Samuel D. Pollard, Sudharshan Srinivasan, and Boyana Norris



SIAM CSE 2019, MS168: Reproducibility in Network Algorithms II
February 26, 2019

Steps to (Compute a Graph Property Quickly)



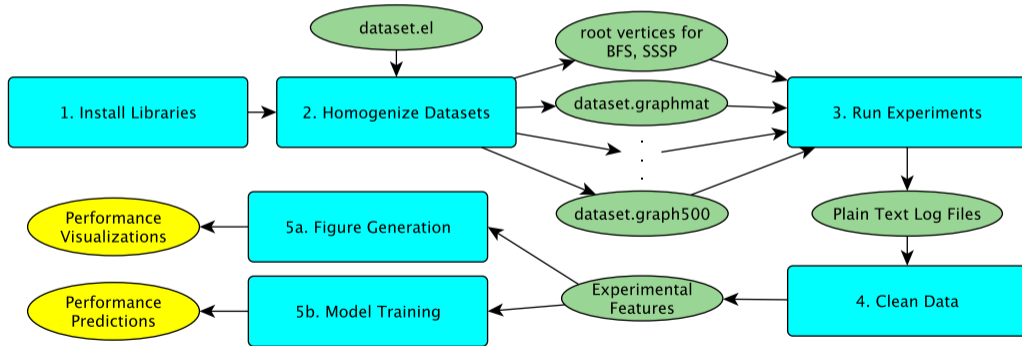
- 1 Find a package
- 2 Nail down dependencies
- 3 Reformat graph input file
- 4 Learn how to use CLI/API
- 5 Run gamut of experiments
- 6 Measure performance of experiments
- 7 Combine, analyze, and decide
- 8 Link with existing workflow

Steps to Quickly (Compute a Graph Property Quickly)



- ~~1 Find a package~~
- ~~2 Nail down dependencies~~
- ~~3 Reformat graph input file~~
- ~~4 Learn how to use CLI/API~~
- 5 Run gamut of experiments on my computer
- 6 Measure performance of such experiments
- ~~7 Combine, analyze, and decide~~
- 8 Link with existing workflow

- ▶ Social network analysis [2]
 - Twitter has its own graph processing package (GraphJet), Google has its own language (Pregel)
- ▶ Bioinformatics [4]
 - Tend to have more complex analysis of smaller datasets
 - non-e.g. Needleman-Wunsch algorithm for global sequence alignment



- ▶ Ensure we're actually measuring the right times!
- ▶ File read time was mistakenly used as performance measurement
- ▶ Abridged log file from GraphMat
 - finished file read. time: **2.65211**
 - load graph: 5.91229 sec
 - initialize engine: 8.32081e-05 sec
 - run algorithm 1 (count degree): 0.0555639 sec
 - run algorithm 2 (compute PageRank): 0.149445 sec
 - print output: 0.0641179 sec
 - deinitialize engine: 0.00022006 sec

- ▶ Binary can be much faster (factor of at least 3)
 - No parsing, can just store into array
 - Less portable
 - May be serialization of internal data structures
- ▶ 0-indexed or 1-indexed?
 - May not be interchangeable

this?

```
55555 44444
44444 22222
11111 66666
```

or this?

```
5 4
4 2
1 6
```

or this?


```
4 3
3 1
0 5
```

► PageRank stopping criterion

- $\|p_t - p_{t-1}\|_1$
- $\|p_t - p_{t-1}\|_\infty$
- Stop when no weights change (machine ϵ)


► Triangle counting

- Count both directions of triangle?
- One, two, or three triangles?



Starting (Root) Vertices

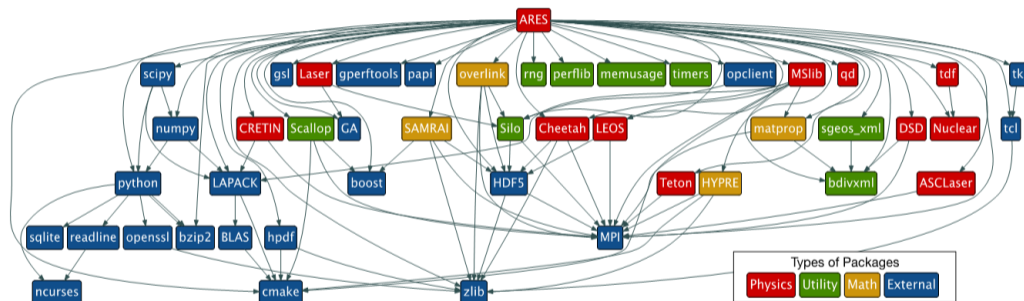
- ▶ Performance of BFS and SSSP depends on where you start
- ▶ More reachable vertices



- ▶ Running experiments is expensive
 - Read in a 10+GB file just to do BFS once
 - Batched is attractive but requires per-package modification

NAME	ST	TIME	NODES
epg1t-22	R	12:24:12	1
epg28t-22	R	4:45:20	1
epg32t-22	R	3:15:10	1
epg48t-22	R	2:45:54	1
epg54t-22	R	2:45:54	1
epg56t-22	R	2:45:54	1
epg40t-22	R	2:47:34	1

How Reproducible is Reproducible Enough?

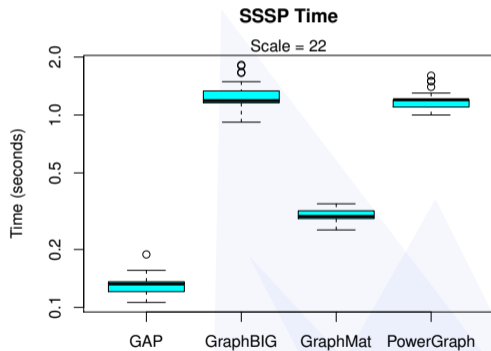



[1]

- ▶ We want better performance!
 - Computer upgrades, new package versions
 - All the lovely breaking changes that come along with them

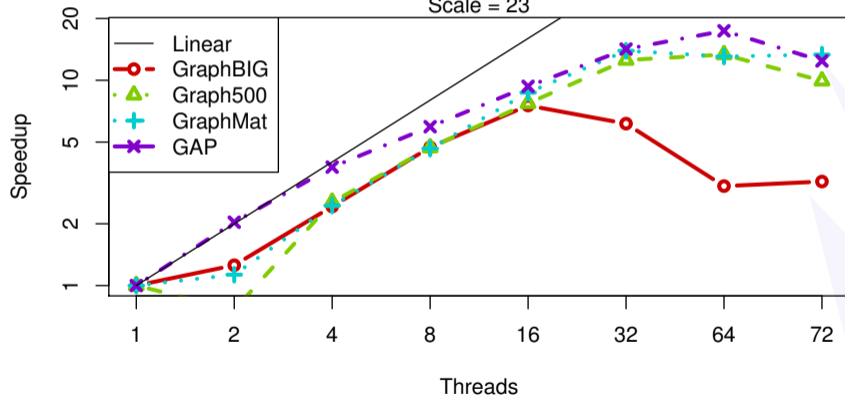
- ▶ BFS - α and β (direction-optimizing). This is a combination of:
 - bottom-up: unvisited nodes searching for visited parents
 - top-down: visited nodes searching for unvisited children
- ▶ SSSP - Δ -stepping; order nodes using ranges of size Δ to distribute work

- ▶ Data structure construction time sometimes not split



BFS Speedup

Scale = 23



$\approx 16 * 2^{23}$
 $\approx 135M$
 edges

- ▶ What if I naïvely wrote serial implementation?
- ▶ Graph500 has Serial, OpenMP, and MPI implementations

- ▶ Pick the best graph package for your hardware and graph
 - 1 Compute features of a graph
 - This may be expensive
 - 2 Train a model based on these features
- ▶ Apply work in linear solver recommendations [5] to graph processing packages

- ▶ We use 12 features (the ones computed on SNAP)
- ▶ e.g. # vertices, # edges, diameter, clustering coefficient


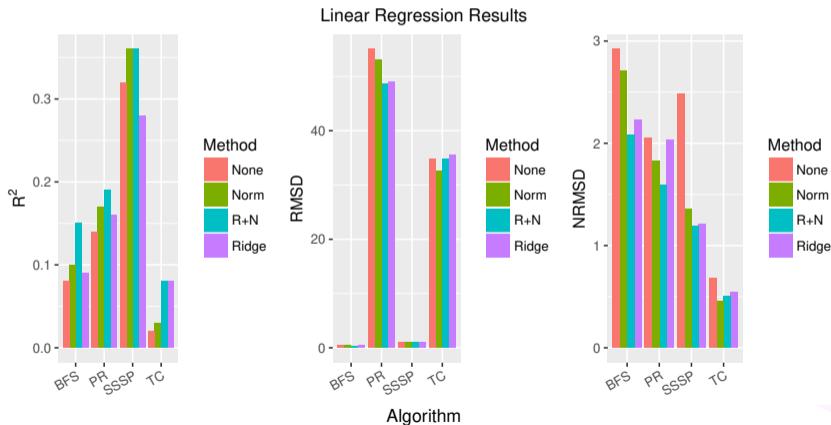


Figure: A scale free network has low diameter [3]



- ▶ Worst case: BFS with factor of 2
- ▶ Best case: TC with factor of 0.5

(b) Random Forest

TC		good	bad
	good	43	1
	bad	0	118
BFS		good	bad
	good	110	0
	bad	0	147
PR		good	bad
	good	75	0
	bad	1	193
SSSP		good	bad
	good	51	11
	bad	9	198

- ▶ Manually selecting between dozens of implementations is infeasible
- ▶ Automated installation, performance experiments can help reproducibility
- ▶ Computing graph features and performance models beforehand can facilitate package selection

- ▶ Leaderboard of performance results (like Graph500)
- ▶ Containerization
 - Singularity
 - Docker
 - Extreme-Scale Scientific Software Stack (E4S) software stack

- [1] Gamblin, T., LeGendre, M., Collette, M. R., Lee, G. L., Moody, A., de Supinski, B. R., and Futral, S.
The spack package manager: bringing order to hpc software chaos.
In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (Nov. 2015), SC '15, pp. 40:1–40:12.
- [2] Kang, U., Meeder, B., and Faloutsos, C.
Spectral analysis for billion-scale graphs: Discoveries and implementation.
In Advances in Knowledge Discovery and Data Mining (Berlin, 2011), vol. 6635 of Lecture Notes in Computer Science, Springer.
- [3] Lamberson, P. J.
Scale-free network.
Available at <http://social-dynamics.org/scale-free-network/>.

- [4] Pavlopoulos, G. A., Secrier, M., Moschopoulos, C. N., Soldatos, T. G., Kossida, S., Aerts, J., Schneider, R., and Bagos, P. G.
Using graph theory to analyze biological networks.
In [BioData Mining](#) (Bethesda, MD, 2011), PubMed Central.
- [5] Sood, K., Norris, B., and Jessup, E.
Comparative performance modeling of parallel preconditioned krylov methods.
In [IEEE 19th International Conference on High Performance Computing and Communications; 15th International Conference on Smart City; 3rd International Conference on Data Science and Systems](#) (Dec. 2017), HPCC/SmartCity/DSS, pp. 26–33.