Parallel Complex Network Partitioning

George M. Slota and Kamesh Madduri (advisor)

gms5016@psu.edu, madduri@cse.psu.edu

Partitioning for Graph Analytics

Partitioning large-scale real-world graphs for parallel analysis is challenging. We ideally want a partition that satisfies the following:

- Balanced number of vertices and edges per partition to evenly distribute computation among tasks.
- Small maximal per-part edge cut to balance communication requirements among tasks.
- Small total edge cut to minimize total communication requirements.
- Fast partitioning with minimal memory and computational overhead.

Visualization of PULP: Partitioning Using Label Propagation

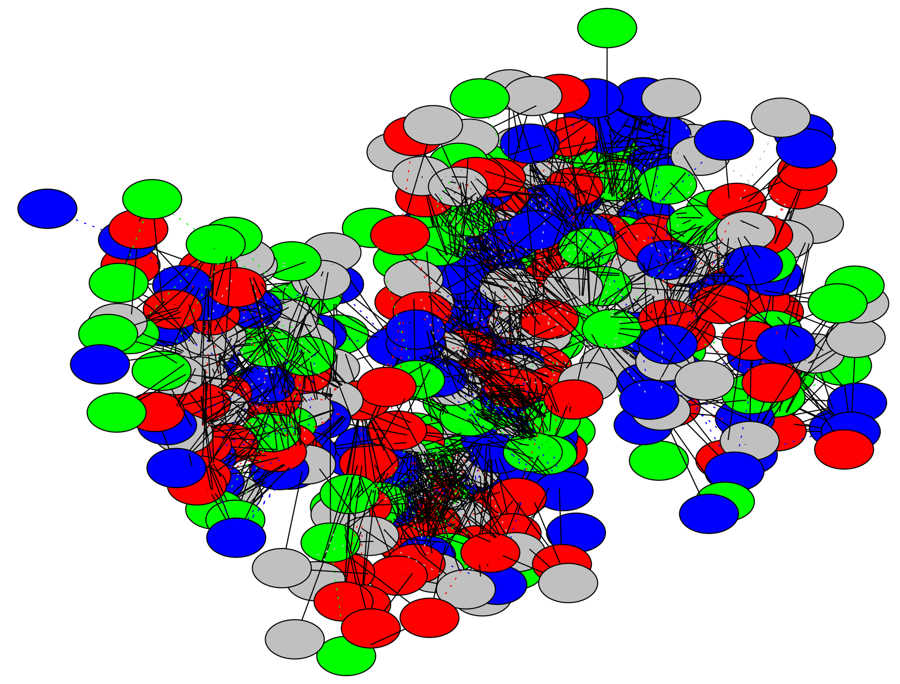

1.) Randomly initialize the partition. Solid lines indicate cut edges and dotted lines indicate intra-part edges.

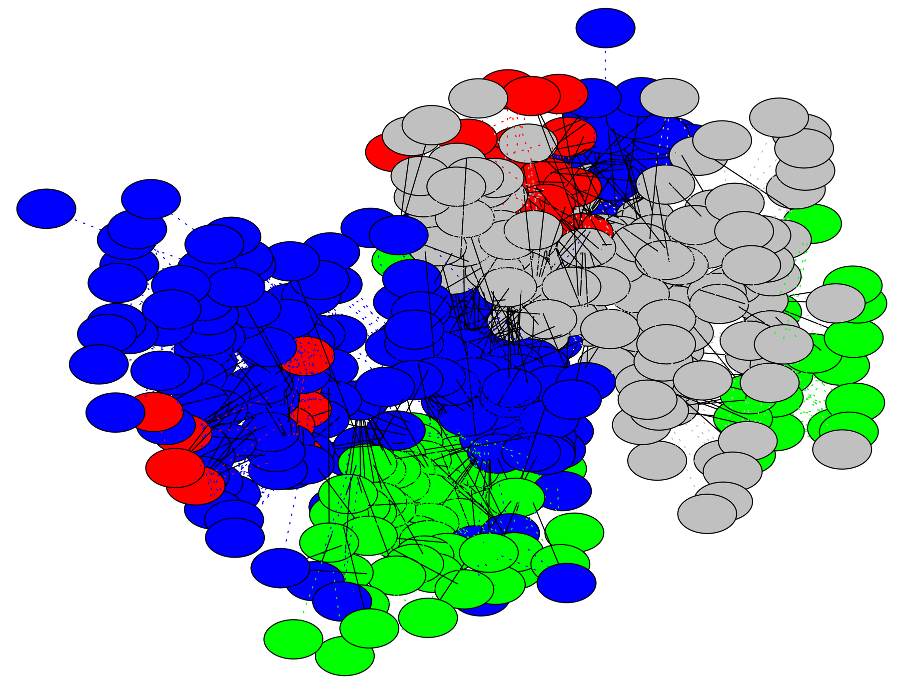

2.) Perform label propagation to create clusters and minimize edge cut.

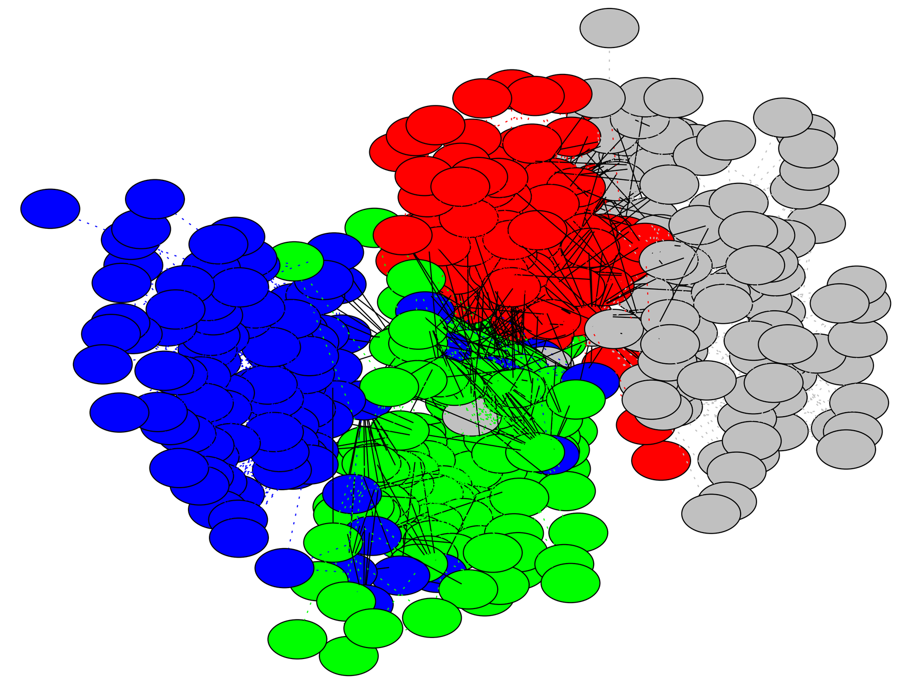

3.) Balance the partition for vertices while additionally refining to minimize edge cut.

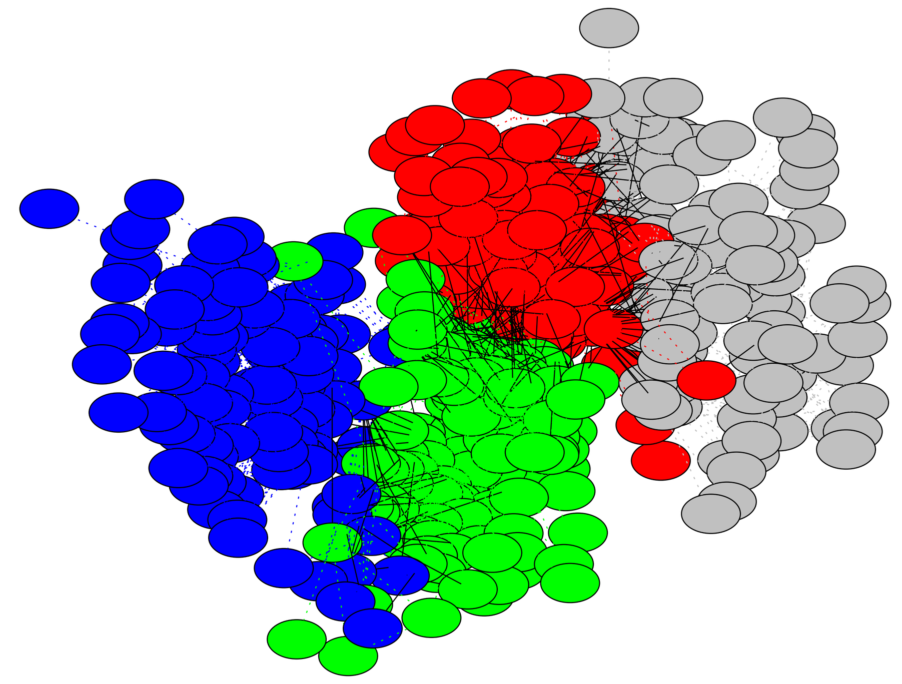

4.) Further balance the partition for edges and minimize maximal per-part edge cut.

Network shown is the Infectious network dataset from KONECT (http://konect.uni-koblenz.de/)

PULP, PULP-M, PULP-MM

To satisfy the above objectives, we employ an efficient iterative label propagation-based [1] partitioning scheme to exploit the community structures inherent in most small-world graphs. More algorithmic detail is available in [4].

```
Initialize p random partitions.
Execute degree-weighted label propagation.
for k1 iterations do
    for k2 iterations do
        Balance partitions to satisfy constraint 1.
        Refine partitions to minimize objective 1.
    for k3 iterations do
        Balance partitions to satisfy constraint 2
          and minimize objective 2.
        Refine partitions to minimize objective 1.
```

We call the above algorithm **PULP-MM**. By running parts of the above we also have **PULP** (minimize edge cut and balance vertices) and **PULP-M** (minimize edge cut and balance vertices and edges).

Results: Superior Scalability and Performance Compared to Multilevel Partitioners

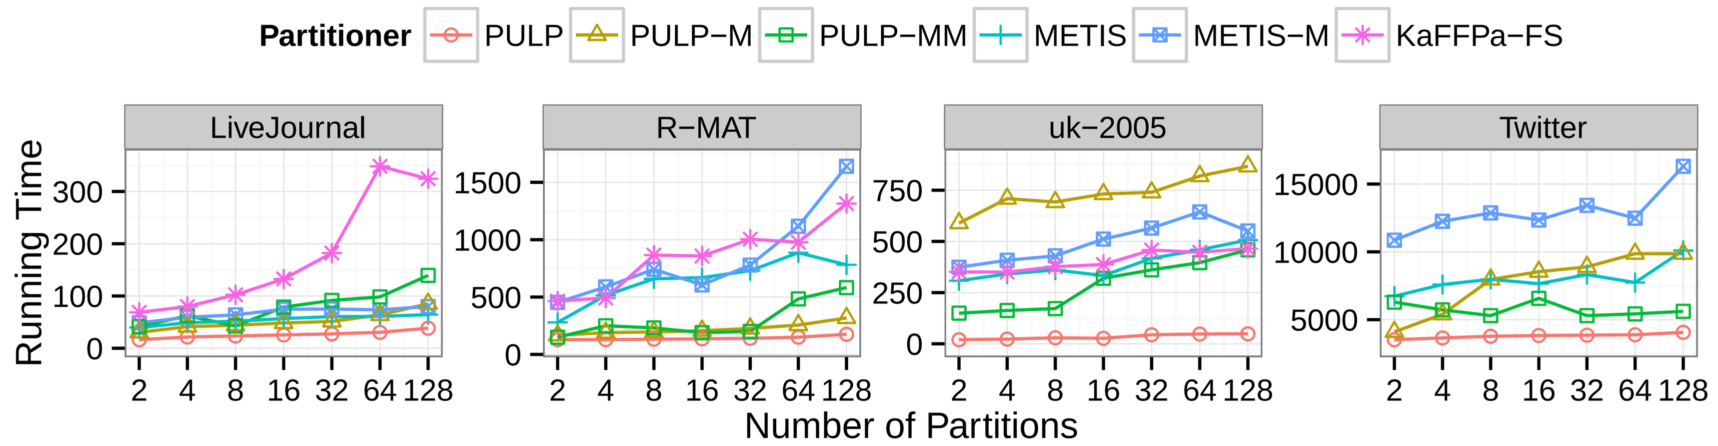

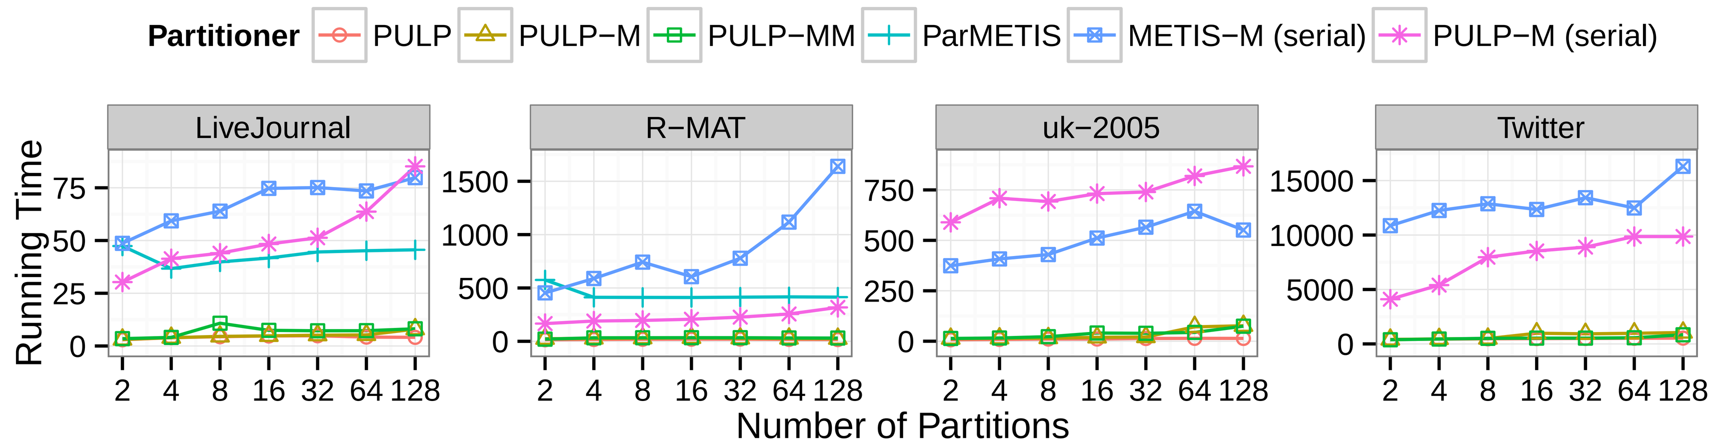

Execution Time: PULP gives up to an order of magnitude speedup in both serial and parallel partitioning relative to KaFFPa [3] and METIS/ParMETIS [2].

Network	Memory Utilization for 128 Parts				Decrease
	METIS-M	KaFFPa	PULP-MM	Graph Size	vs. Best
LiveJournal	7.2 GB	5.0 GB	0.44 GB	0.33 GB	11×
R-MAT	42 GB	64 GB	1.2 GB	1.02 GB	35×
uk-2005	41 GB	27 GB	7.9 GB	7.12 GB	3.4×
Twitter	487 GB	-	14 GB	12.2 GB	39×

Memory Savings: PULP's direct single-level approach avoids the considerable memory overheads of multi-level schemes.

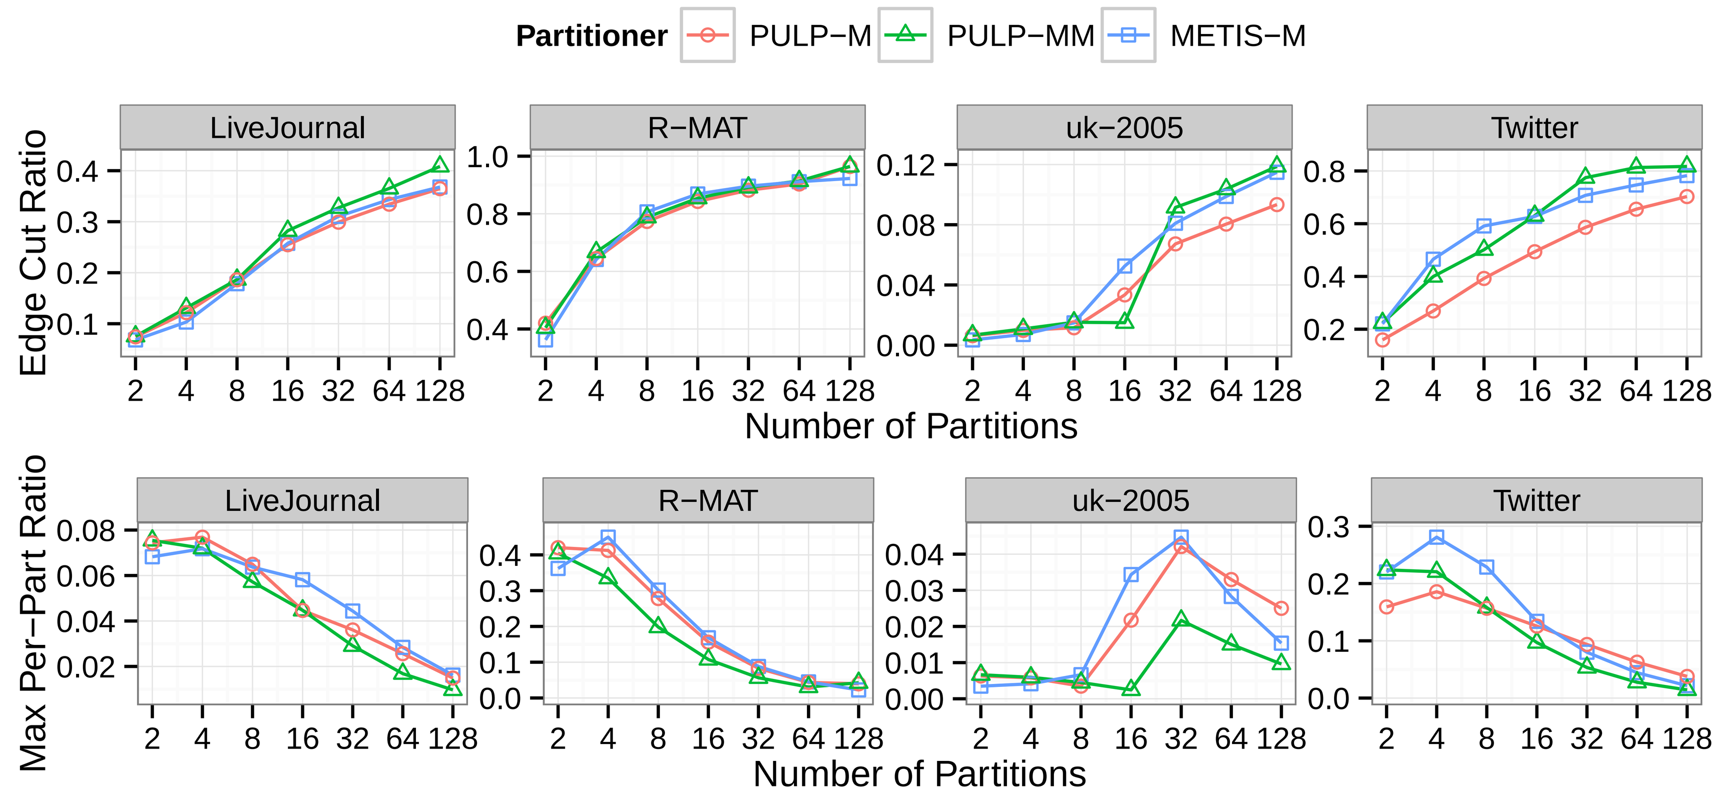

Cut Quality: Our multi-constraint PULP-M demonstrates superior cut quality to multi-constraint METIS-M (top), while PULP-MM demonstrates superior performance in terms of cut balance or maximal per-part cut (bottom).

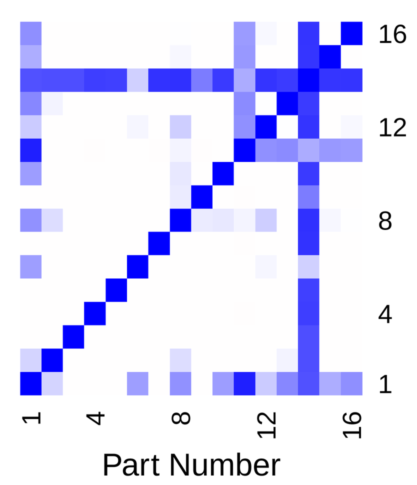

PULP-MM

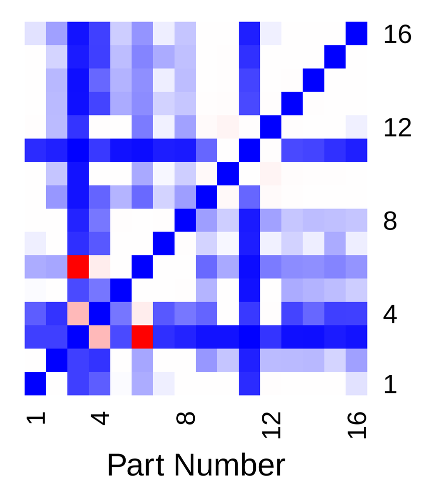

METIS-M

Balanced Communication: PULP-MM demonstrates a much more balanced cut than METIS-M in terms of the cut between all (p_i,p_j) pairs for 16 parts of uk-2005. Dark blue indicates few cuts, white is the average cut, and dark red is the maximal cut.

Conclusions

Utilizing an iterative label propagation-based approach to small-world graph partitioning can offer orders of magnitude improvement in both running time and memory utilization, while producing better cut quality under the complex objectives modern graph analytics requires.

References

[1] U. N. Raghavan, R. Albert, S. Kumara Near linear time algorithm to detect community structures in large-scale networks In *Physical Review E*, vol. 76, no. 3, 2007.

[2] G. Karypis and V. Kumar METIS: A software package for partitioning unstructured graphs, partitioning meshes, and computing fill-reducing orderings of sparse matrices. version 5.1.0, March 2013.

[3] H. Meyerhenke, P. Sanders, C. Schulz Partitioning complex networks via size-constrained clustering In *Proc. SEA*, 2014.

[4] G. Slota, K. Madduri, S. Rajamanickam PULP: Scalable Multi-Objective Multi-Constraint Partitioning for Small-World Networks In *Proc. IEEE BigData Conf.*, 2014.

Acknowledgements

Thanks to Siva Rajamanickam for assistance and guidance. This work is supported by NSF grant ACI-1253881 and the DOE Office of Science through the FASTMath SciDAC Institute. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.