

MOTIVATION

Partitioning graphs is an important task in parallel scientific computing and graph analysis applications to balance work per-task and minimize communication. For large graphs and meshes, parallelizing the partitioning process is heavily desired. If parts within a partition remain **connected**, they can allow simpler computations and reduce the number of very small components that are disconnected from the rest of their part.

Applications for partitioning:

- Parallel graph analysis using **clusters**
- Scientific computing of fine meshes, which desires connectivity and multi-weight vertices
- Problem solving in biology, power grids, image processing, and many more [3]
- **Redistricting**, which requires connectivity, and can reduce *gerrymandering* issues [2]

BACKGROUND

Graph partitioning typically is the process of creating vertex-disjoint sets on a graph that minimize edge cut and constrain weight per-part. This project focuses on adding an additional constraint: **connec**tivity. It is common for partitioning algorithms to not guarantee connected partitions, but it can be a highly desired feature.

Several categories of graph partitioning algorithms exist. Global algorithms work with the entire graph and compute a direct solution, and they're mainly used for smaller graphs. Iterative **improvement** heuristics gradually improve partitions from an initial starting position. These take a small set of vertices at a time and move them to the best suited part. Multilevel algorithms condense the graph small enough to do an initial partitioning, then work back and uncondense the graph [3].

The work done in this project is built on an existing algorithm: **Partitioning using Label Propa**gation (PuLP) [1]. PuLP has achieved good vertex balancing and edge cut, while obtaining greatly improved runtime and memory usage over other partitioning algorithms. The results of this will be integrated into PuLP to address connectivity issues.

Merging works decently for some graphs, such as a road network, while sacrificing some of the vertex balancing and edge cut. But for other complex graphs, like a social network, vertices are thrown way off balance.

MAINTAINING CONNECTIVITY IN PARALLEL GRAPH PARTITIONING Christopher I. Jones Ian Bogle George M. Slota Renselaer slotag@rpi.edu jonesc10@rpi.edu boglei@rpi.edu

BFS Trees

METHODS FOR MAINTAINING CONNECTIVITY

We are building upon an existing PuLP [1] algorithm, adding features and functionality to aid connectivity. The original algorithm has 2 inner iterations:

• **Balance** focuses on balancing the vertices between parts. Each vertex considers the *gain* to being in part P:

$$gain(P) = n_P \cdot w_P$$

 n_P is the number of neighbors in P. w_P is how underweight P is.

• **Refine** works on minimizing the edge cut relative to each vertex.

Merge Small Components



Merge

Ideally, each part has one component. Calculate BFS trees on each component, Every some number of iterations, merge and only allow the **leaves** to move. Or, all components that are not the largest every so many vertex movements, calcuin their part into a neighboring compo- late bi-connectivity, and prevent **artic**ulation points from moving. nent.

RESULTS



Leaves where or

Com

APPLICATION: POLITICAL REDISTRICTING

Gerrymandering is the issue of politicians manipulating districts based on where voters live. Translating the redistricting problem to graph partitioning, census blocks are vertices, **districts** are parts, and edges represent shared bor-Previous work includes ders. PEAR [2], which is an evolutionary algorithm. We've applied PuLP to counties in North Carolina, using population demographics as vertex weights and border lengths as edge weights.



PuLP: A redistricting using PuLP with small component merging.

Restrict Vertex Movement





Restrict Vertex Movement						Additional Metrics				
C	Vertex Overweight	Edge Cut	Small	Comp Count	• t		Vertex Overweight	Edge Cut	$\frac{\text{Small Cor}}{\text{Count}}$	nponents Size
$egin{array}{c} 0 \ 1 \ 2 \end{array}$	1.5-2.5 1.1-1.2 ≤ 1.1	>1,000K 0 800K 4,000 650K 5,000)))	Normal alg. Method (1) Method (2)	$1.039 \\ 1.040 \\ 1.057$	125.4K 125.5K 434.4K	$6,100 \\ 5,800 \\ 14,400$	$22.17 \\ 22.53 \\ 13.94$
∞	∞ <1.1 150K 5,500 Duly Mothod: Results are from a web gravel graph					Add Merging as a Final Step:				
nly vertices with C or less children are allowed to move.						Normal alg. Method (1)	$1.182 \\ 1.174$	$94.5\mathrm{K}$ $95.3\mathrm{K}$	0 0	N/A N/A
	10K	2.5K 1	,250	320	80	Method (2)	1.081	356K	0	N/A
p. Co	ount 2,000	1,500 1	,400 1	,275	$1,\!200$	Using an Ad	ditional Met	ric: Thes	e numbers	are from a

The original algorithm outputs only around 130 components.

Articulation Point Method: Running on a road network web-crawl graph when adding the Average Number of Children graph, biconnectivity is recalculated every R vertex updates. metric, a_P . Method (2) produces many components that are relatively small, allowing those components to be easily merged.

ACKNOWLEDGEMENTS

This work was supported in part by the U.S. Department of En-ergy, Office of Science, Office of Advanced Scientific Computing Re-search, Scientific Discovery through Advanced Computing (SciDAC) program through the FASTMath Institute under Contract No. DE-AC02-05CH11231 at Rensselaer Polytechnic Institute and Sandia Na-tional Laboratories. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.

REFERENCES

- 30, 78-92, 2016.

rics that can be taken into consideration, such as compactness and competitiveness, as defined in [2]. **Compactness** is related to how well connected a district is and is defined as the area divided by perimeter squared. Competitiveness takes the numbers of voters in opposing parties to measure how evenly they're distributed in the districts. Other demographic data can also be considered.

Redistricting introduces new met-

Additional Metrics

Add a new metric to the *gain* equation by 2 methods:

 $gain(P) = (n_P + \boldsymbol{a_P}) \cdot w_P$ $gain(P) = n_P \cdot w_P / \boldsymbol{a_P}$

$$a_P = \sum_{\substack{u \in N(v)\\u \in P}} \frac{children(u)}{d(u)}$$

Where N(v) is the neighborhood of v, children(u) is the number of children u has in the component's BFS tree, and d(u) is the degree of u.

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

[2] Y. Y. Liu, W. K. T. Cho, S. Wang PEAR: a massively parallel evolutionary computation approach for political redistricting optimization and analysis. In Swarm and Evolutionary Computation,

[3] A. Buluc, H. Meyerhenke, I. Safro, P. Sanders, C. Schulz Recent Advances in Graph Partitioning In Algorithm Engineering, 2016.