MAINTAINING CONNECTIVITY IN PARALLEL GRAPH PARTITIONING

Christopher I. Jones jonesc10@rpi.edu

Ian Bogle boglei@rpi.edu George M. Slota slotag@rpi.edu



MOTIVATION

Partitioning graphs is an important preprocessing task in parallel scientific computing and graph analysis applications to balance work per-rank and minimize global communication. In many application, it is desirable for convergence and solution stability that each part within a partition remain fully **connected**. This work focuses on developing partitioning methods that ensure such connectivity.

APPLICATIONS

General applications we consider:

- Scientific computing on fine meshes, which desires connectivity and multiweight vertices
- Problem solving in biology, power grids, image processing, and many more [3]
- **Redistricting**, which requires connectivity, and can reduce *gerrymandering* issues in U.S. politics [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: connectivity. It is common for partitioning algorithms to not guarantee connected partitions, but it can be a highly desired feature. The work done in this project is built on an existing algorithm: **Partitioning using La**bel Propagation (PuLP) [1].

METHODS EXPLORED 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.

Merge Small Components



neighboring component.



Restrict Vertex Movement

Calculate BFS trees on each com-Ideally, each part has one compo- ponent, and only allow the leaves nent. Every some number of itera- to move. Or, every so many vertex • Refine works on minimizing tions, merge all components that are movements, calculate bi-connectivity, the edge cut relative to each not the largest in their part into a and prevent **articulation points** from moving.

Additional Metrics

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

 $gain(P) = (n_P + \boldsymbol{a_P}) \cdot w_P$ (1) $gain(P) = n_P \cdot w_P / \boldsymbol{a_P}$ (2)

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

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.

N/A

RESULTS

vertex.

Merge Small Components	Restrict Vertex Movement			Additional Metrics					
Social Network Metrics During Merge Iterations	C	Vertex Imbalance	Edge Cut	Small Comp. Count		Vertex Imb.	Edge Cut	$\frac{\text{Small } \mathbf{C}}{\text{Count}}$	Comps Size
250000	0	1.5 - 2.5	>1,000K	0	Normal alg.	1.039	$125.4\mathrm{K}$	$6,\!100$	22.17
	1	1.1 - 1.2	800K	$4,\!000$	Method (1)	1.040	$125.5\mathrm{K}$	$5,\!800$	22.53
200000	2	≤ 1.1	$650 \mathrm{K}$	$5,\!000$	Method (2)	1.057	$434.4\mathrm{K}$	$14,\!400$	13.94
150000 1.5 Verve	∞	<1.1	150K	5,500	Add Merging as a Final Step:				
Edge Cut	Leaves (Only Meth	nod: We d	consider a BFS on	Normal alg.	1.182	$94.5\mathrm{K}$	0	N/A
Vert Overweight	each part	from a ce	ntral verte	ex. We then only	Method (1)	1.174	$95.3\mathrm{K}$	0	N/A



0	1.5 - 2.5	>1,000K	0
1	1.1 - 1.2	800K	$4,\!000$
2	≤ 1.1	$650 \mathrm{K}$	$5,\!000$
∞	<1.1	$150\mathrm{K}$	$5,\!500$

each part from a central vertex. We then only allow vertices with C or less children to change parts.

Merging Method: Here, we iteratively parti-	
tion \rightarrow merge disconnected components \rightarrow par-	
tition. This merging procedure works well for	
regular graphs and meshes, while sacrificing some	\mathbf{A}
part balance and edge cut. But for more complex	ro
graphs, like a social network, part imbalance can	ev
be considerable.	m

R	10K	$2.5\mathrm{K}$	$1,\!250$	320	80
Comp. Count	$2,\!000$	1,500	$1,\!400$	$1,\!275$	1,200

ent of articulation/cut vertices for each part.

 $\overline{\mathbf{U}}$ Using an Additional Metric: These numbers ____ are from a web-crawl graph when adding the Av-- erage Number of Children metric, a_P when using **rticulation Point Method:** Running on a the above gain equations of (1) and (2). Method bad network, part biconnectivity is recalculated (2) produces many components that are relatively very R vertex updates. We restrict the move-small, allowing those components to be easily merged.

356K

1.081

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 borders. Previous work includes 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.



Redistricting introduces new metrics 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 Competitiveness squared. 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.

REFERENCES

Method (2)

- [1] 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

Gerrymandering: North Carolina districts from 2013-2016.



PuLP: A redistricting using PuLP with small component merging.

parallel evolutionary computation approach for political redistricting optimization and analysis. In Swarm and Evolutionary Computation, 30, 78-92, 2016.

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

ACKNOWLEDGEMENTS

This work was supported in part by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Scientific Discovery through Advanced Computing (Sci-DAC) program through the FASTMath Institute under Contract No. DÉ-AC02-05CH11231 at Rensselaer Polytechnic Institute and Sandia National 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.