OVERLAPPING CLUSTERS FOR DISTRIBUTED COMPUTATION

REID ANDERSEN · MICROSOFT
DAVID F. GLEICH · SANDIA
VAHAB S. MIRROKNI · GOOGLE

1. THE IDEA

Scalable, distributed algorithms must address communication problems. We investigate *overlapping clusters*, or vertex partitions that intersect, for graph computations. This setup stores more of the graph than required but then affords the ease of implementation of vertex partitioned algorithms. Our hope is that this technique allows us to reduce communication in a computation on a distributed graph.

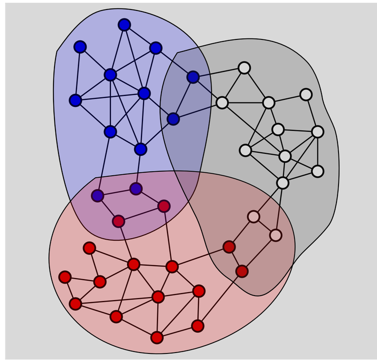

2. RELATED WORK

The motivation above draws on recent work in communication avoiding algorithms. Mohiyuddin et al. (SC09) design a matrix-powers kernel that gives rise to an overlapping partition. Fritzsche et al. (CSC2009) develop an overlapping clustering for a Schwarz method. Both techniques extend an initial partitioning with overlap. Our procedure generates overlap directly. Indeed, Schwarz methods are commonly used to capitalize on overlap. Elsewhere, overlapping communities (Ahn et al, Nature 2009; Mishra et al. WAW2007) are now a popular model of structure in social networks. These have long been studied in statistics (Cole and Wishart, CompJ 1970).

3. PROBLEM SETUP

Let Vol(C) = sum of degrees for $v \in C$; Cut(C) = total edges between C and the rest of the graph. Note that Vol(C) is a proxy for the adjacency data size of vertices in C.

Given a graph G, an overlapping clustering $(\mathcal{C}, \tau)$ is a set of clusters $\mathcal{C}$ and a mapping from each vertex to a home cluster τ. The total number of edges in a cluster (Vol(C)) is constrained by MaxVol. In a random walk on an overlapping clustering, the walk moves from cluster to cluster. On leaving a cluster, it goes to the home cluster of the new vertex: e.g. In the illustrations here, the color indicates the home vertices for a cluster. (See the example above too.) A transition between clusters is a swap, and requires a communication if the underlying graph is distributed. We thus wish to minimize swaps in a random walk. Let $\rho_T(v)$ = the expected fraction of steps that swap in a T-step walk starting from v. We study: $\rho_\infty = \lim_{T\to\infty} \frac{1}{n} \sum_v \rho_T(v)$, the fraction of steps with swaps for a long walk. For a cycle graph, we can prove that overlap reduces the communication.

THEOREM Consider a large cycle C_n of $n = M\ell$ nodes for a large number $M > 0$, and let the maximum volume of a cluster MaxVol be ℓ. Let P be the optimal partitioning of G to non-overlapping clusters of size at most MaxVol and ρ^*_∞ be the swapping probability of P. There exists an overlapping cover with TotalVol of 2Vol(G) whose swapping probability ρ'_∞ is less than $\rho^*_\infty/\Omega(\text{MaxVol})$.

(Proof Sketch) The overlapping clustering that achieves this bound is:

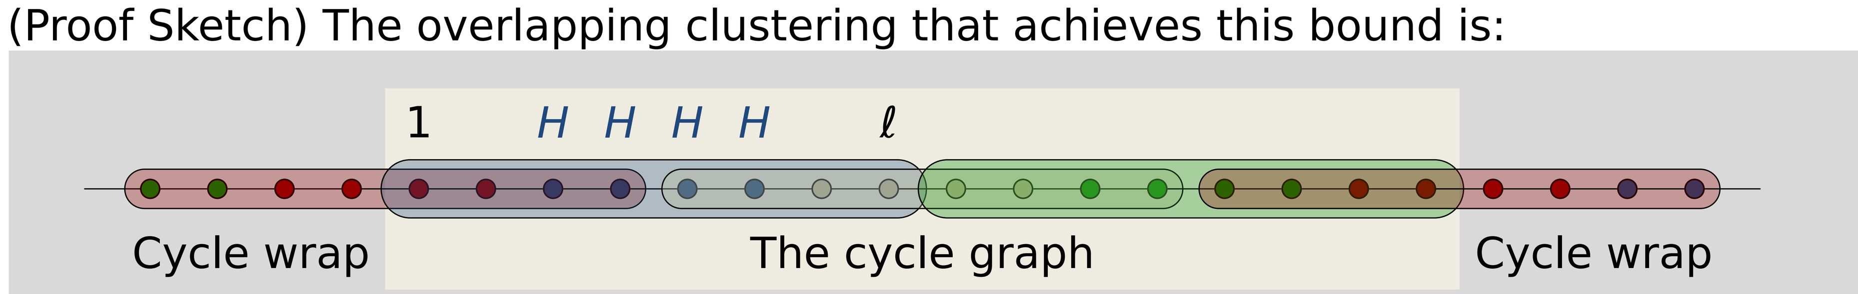

Each cluster has ℓ vertices, and the home vertices are the "middle" ones, as in the four vertices labeled H above for the blue cluster. The best ρ_∞ for a partitioning is $\frac{2}{\ell}$ because $\rho_\infty = \frac{1}{\text{Vol}(G)} \sum_{C \in \mathcal{P}} \text{Cut}(C)$ for a partitioning. A random walk travels $O(\sqrt{t})$ distance in t steps. The edge of an overlapping cluster is always in the center of another cluster, and so it will take $\ell^2/4$ steps to exit after a swap, yielding $\rho_\infty = \frac{4}{\ell^2}$.

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.

4. HEURISTICS FOR OVERLAPPING CLUSTERS

Optimizing ρ_∞ with a MaxVol constraint is NP-hard by a relaxation from minimum bisection. To produce clusters with a small ρ_∞ we use a multi-stage heuristic:

1. Identify candidate clusters. Use a PageRank clustering heuristic or METIS to find small conductance clusters up to size MaxVol.

2. Compute well-contained sets. For each vertex, compute the time for a random walk to leave a cluster starting there and use this to pick home vertices.

3. Cover with cluster cores. Approximately solve a set-cover problem to pick a subset of clusters.

4. Combine clusters. Finally, we combine any small clusters until the final size of each is about MaxVol.

THE OUTPUT

A PageRank cluster
(color is containment, white=best)

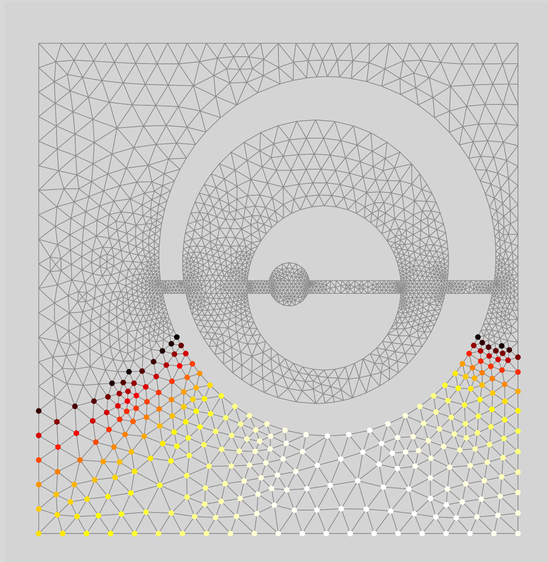

... and another

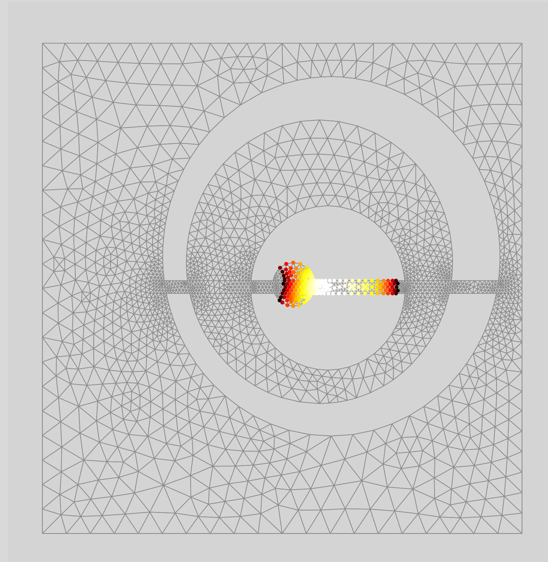

A combined cluster
(red = home vertices)

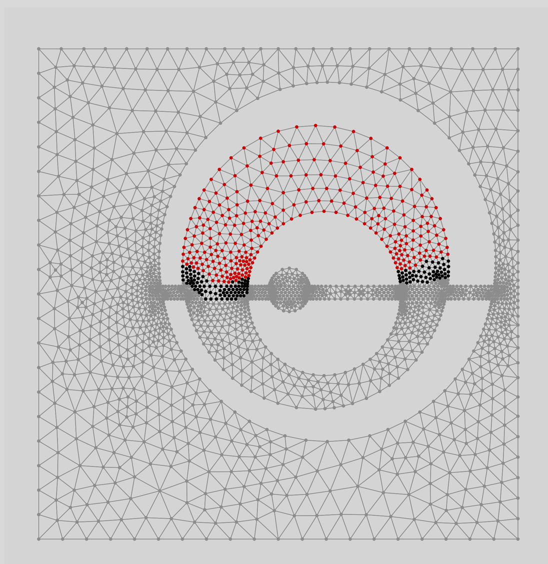

... and another

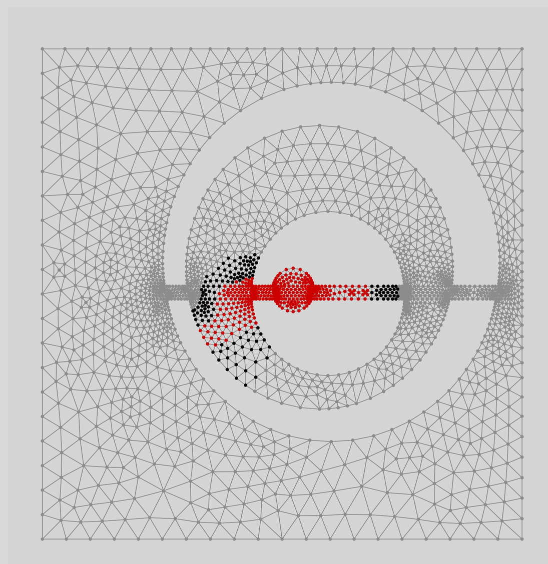

Overall best containment
(white=best)

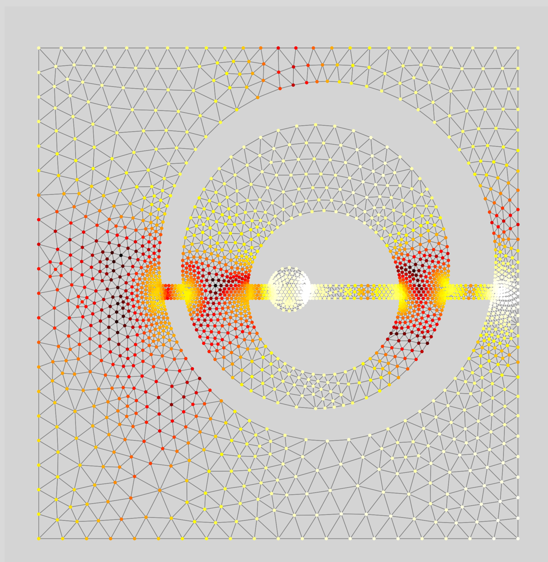

Vertex overlap
(gray=1, black=2, red=2)

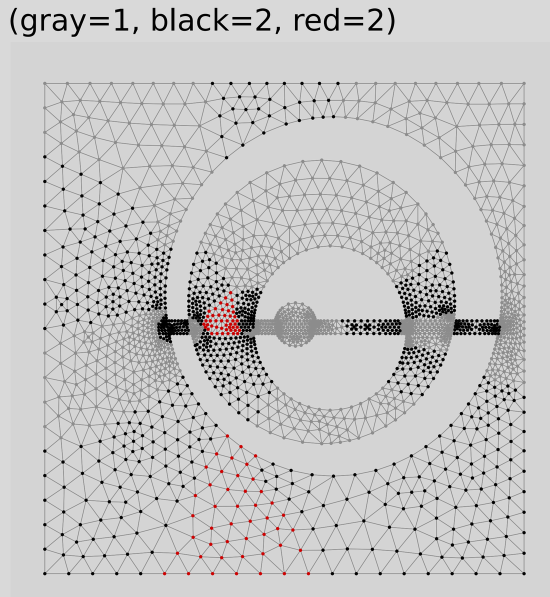

5. DATA

We empirically study this idea on 10 public graphs.

Graph	$\|V\|$	$\|E\|$	max deg	$\|E\|/\|V\|$
onera	85567	419201	5	4.9
usroads	126146	323900	7	2.6
annulus	500000	2999258	19	6.0
email-Enron	33696	361622	1383	10.7
soc-Slashdot	77360	1015667	2540	13.1
dico	111982	2750576	68191	24.6
lcsh	144791	394186	1025	2.7
web-Google	855802	8582704	6332	10.0
as-skitter	1694616	22188418	35455	13.1
cit-Patents	3764117	33023481	793	8.8

6. RESULTS

We present two types of results: (i) an estimated swapping probability ρ_∞; and (ii) the communication volume of a parallel PageRank solution (link-following $\alpha = 0.85$) using an additive Schwarz method. The volume ratio is the amount of extra storage for the overlap (2 means we store the graph twice). Below, as the ratio increases, the swapping probability and PageRank communication volume decreases.

Two graphs, in detail

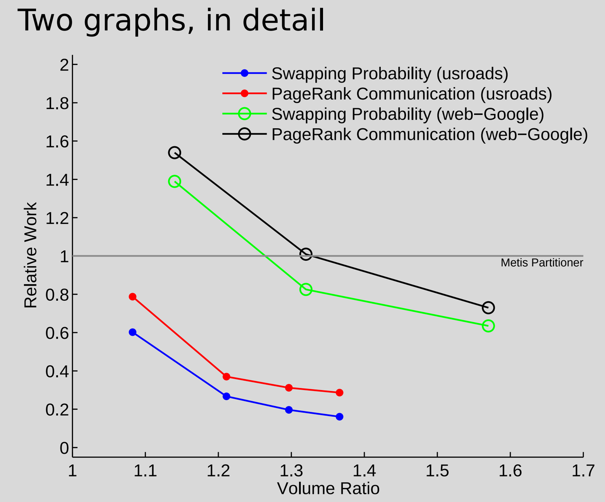

All graphs, multiple experiments.

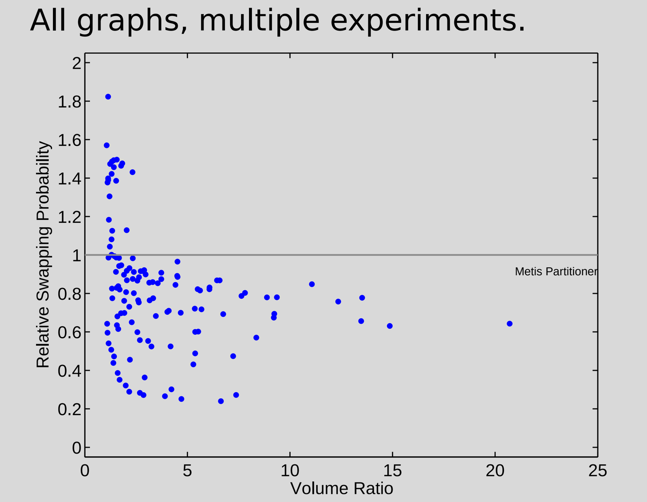

The communication ratio of our best result for the PageRank communication volume compared to METIS or GRACLUS shows that the method works for 6 of them (perf. ratio < 1). The 0 communication result is not a bug.

Graph	Comm. of Partition	Comm. of Overlap	Perf. Ratio	Vol. Ratio
onera	18654	48	0.003	2.82
usroads	3256	0	0.000	1.49
annulus	12074	2	0.000	0.01
email-Enron	194536*	235316	1.210	1.7
soc-Slashdot	875435*	1.3×10^6	1.480	1.78
dico	1.5×10^6*	2.0×10^6	1.320	1.53
lcsh	73000*	48777	0.668	2.17
web-Google	201159*	167609	0.833	1.57
as-skitter	2.4×10^6	3.9×10^6	1.645	1.93
cit-Patents	8.7×10^6	7.3×10^6	0.845	1.34

Finally, we evaluate our heuristic.

At left, the cluster combine procedure reduces 10^6 clusters to around 10^2. Middle, combining clusters can decrease the volume ratio from 10 to around 1. At right, the mean conductance tends to decrease after combining clusters.