


# Engineering a Scalable High Quality Graph Partitioner

Manuel Holtgrewe, Peter Sanders, Christian Schulz


Institute for Theoretical Computer Science, Algorithmics II

Novemeber 24, 2009

# Graph Partitioning




# Well-Known Multi-Level Approach



# Match and Contract

- ▶ Sequential Matching: **GPA** [Maue, Sanders, 2007].
  1. 1st: greedy algorithm finds **heavy paths and cycles**
  2. 2nd: **optimal matching** on those using dynamic programming
- ▶ GPA, Greedy, HEM achieve half-approx. in worst case
- ▶ GPA gives considerably **better results** (empirically)
- ▶ Parallel Matching running on the gap graph:  
Implementation of Manne and Bisseling's sketch [Manne, Bisseling, 2007].
- ▶ Use **edge ratings** for prioritization



# Edge Ratings - Few Basic Principles

- ▶ good to contract heavy edges because this decreases the cut size
- ▶ avoid clusters with many outgoing edges
- ▶ contract light nodes to keep the node weight at any level reasonably uniform.

weight “number of edges in original graph”

expansion edge weight divided by the number of vertices

**expansion\*** edge weight divided by the product of number of vertices



**inner-outer** edge weight divided by number of edges out of cluster

# Initial Partitioning

- ▶ Every PE can perform initial partitioning using different seeds
- ▶ Compared PARTY, PMETIS, and SCOTCH
- ▶ SCOTCH yielded best results
- ▶ we will try other partitioners as well


# Uncontract and Local Improvement

- ▶ we adopt basic approach from [Fiduccia, Mettheyses, 1982] (linear time)
- ▶ **basic idea**: at any time, each PE may work on **one pair** of neighboring blocks
- ▶ local search constrained to moving nodes between these two blocks graph
- ▶ assign pairs of blocks to PEs, through **edge colorings** of *quotient graph*  $Q$





# Uncontract and Local Improvement

- ▶ BFS to reduce communication volume
- ▶ **Queue Selection Strategies** of local search improved quality
  1. use queue promising larger gain (**TopGain**)
  2. use queue of heavier block (**MaxLoad**) if one block is overloaded





# Scalability



- ▶ Streetnetwork Europe ( $|V| = 18M, |E| = 44M$ )
- ▶ DelaunayGraph ( $|V| = 33M, |E| = 100M$ )

# Quality - Walshaw Benchmark

## Five percent case highlights

Graph	2		4		8		16		32		64	
wing	+ 826	<b>770</b>	** 1734	<b>1636</b>	* 2632	<b>2551</b>	* 4106	<b>4015</b>	* 6063	<b>5806</b>	+ 8300	<b>7991</b>
brack2	** <b>660</b>	660	** <b>2739</b>	2755	* <b>6776</b>	6883	+ <b>11557</b>	11558	* 17617	<b>17529</b>	+ 26555	<b>26281</b>
finan512	** <b>162</b>	162	** <b>324</b>	324	** <b>648</b>	648	* <b>1296</b>	1296	** <b>2592</b>	2592	** 10909	<b>10560</b>
fetooth	** 3785	<b>3773</b>	* <b>6863</b>	7027	+ <b>11498</b>	11957	* <b>17509</b>	18090	* 25641	<b>25624</b>	+ 35795	<b>35476</b>
ferotor	** <b>1955</b>	1957	+ <b>7031</b>	7520	* <b>12643</b>	12678	** <b>20098</b>	20263	+ 31611	<b>31576</b>	+ 47186	<b>46608</b>
598a	** 2344	<b>2336</b>	+ <b>7837</b>	7978	** <b>15794</b>	16031	* <b>25782</b>	26257	** <b>39478</b>	40179	** <b>58180</b>	58307
feocean	** <b>311</b>	311	** <b>1688</b>	1704	* <b>3952</b>	4019	+ <b>7671</b>	7838	+ 12953	<b>12746</b>	+ <b>20660</b>	21784
144	* 6502	<b>6362</b>	+ 15313	<b>15250</b>	** <b>25529</b>	25611	+ <b>38182</b>	38478	+ <b>57202</b>	57354	+ 80653	<b>80257</b>
wave	** 8613	<b>8563</b>	+ <b>16780</b>	17306	+ <b>28753</b>	29776	+ <b>42810</b>	43791	** <b>62382</b>	63675	** <b>86867</b>	87654
m14b	** 3844	<b>3802</b>	* <b>13124</b>	13285	** <b>25701</b>	26153	+ <b>42644</b>	43747	* <b>66845</b>	67551	+ <b>99460</b>	100183
auto	* 9587	<b>9450</b>	+ <b>25805</b>	26097	** <b>44915</b>	48174	+ <b>76500</b>	80495	** <b>121988</b>	124251	** <b>174173</b>	174904

# Future Work

- ▶ **Base case**  $\Rightarrow$  Meyerhenke's code?
- ▶ Back to **sequential** partitioning
- ▶ Other pairwise improvers? **Flow?** **Diffusion?**
- ▶ Improve overall **search strategies**
- ▶ Put every aspect of multilevel graph partitioning on trial
- ▶ More information: <http://arxiv.org/pdf/0910.2004>


**Thank you for your interest. Any questions?**

# Building Parts





- ▶ Recursive Coordinate Bisection as a prepartitioner
- ▶ Sequential and parallel algorithms for approximating maximal weight matchings
- ▶ Target functions for edges (*edge ratings*)
- ▶ Parallel contraction code
- ▶ Initial partitioning library (SCOTCH [Pellegrini, Roman, 1996])
- ▶ Pairwise refinement algorithm

# Recursive Coordinate Bisection





- ▶ Similar to Bentley's *kd*-trees
- ▶ Parallel QUICK-SELECT
- ▶ Could be improved using Parallel Randomized Selection algorithm by Blume et al.



# References I





-  B. Hendrickson and R. W. Leland  
*A Multi-Level Algorithm For Partitioning Graphs*  
in Supercomputing 1995
-  G. Karypis and V. Kumar  
*Multilevel Graph Partitioning Schemes*  
ICCP (3) 1995, p. 113-122
-  C. Walshaw, M. Cross, M. G. Everett  
*Dynamic mesh partitioning: a unified optimisation and load-balancing algorithm*  
Tech. Rep. 95/IM/06, Univ. Greenwich, London
-  M. J. Berger and S. H. Bokhari  
*A Partitioning Strategy For Nonuniform Problems On Multiprocessors*  
IEEE Trans. Comput., 36(5):570580, 1987

# References II




-  H. D. Simon  
*Partitioning of Unstructured Problems for Parallel Processing*  
Computing Systems in Engineering
-  A. Pothen, H. D. Simon, K. P. Liu  
*Partitioning Sparse Matrices with Eigenvectors of Graphs.*  
SIAM J. on Matrix Analysis and Applications, 11(3):430-452, 1990.
-  J. Maue and P. Sanders  
*Engineering Algorithms for Approximate Weighted Matching*  
Experimental Algorithms, p. 242-255
-  B. Monien, R. Preis  
*Quality Matching and Local Improvement for Multilevel Graph-Partitioning*  
Parallel Computing, 26(12):1609-1634, 2000



# References III

-  B. W. Kernighan, S. Lin  
*An Efficient Heuristic Procedure for Partitioning Graphs*  
Bell Systems Technical Journal 14:291-307
-  C. Schulz  
*Scalable Parallel Refinement of Graph Partitions*  
Diplomarbeit, Universitt Karlsruhe (TH)
-  C. Fiduccia, R. Mattheyses  
*A Linear-Time Heuristic for Improving Network Partitions*  
Design Automation, 1982
-  M. Halappanavar, F. Dobrian, A. Pothen  
*Matchings in Massive Graphs of Terrascale Computers via Approximation*  
To appear, 2009

## References IV

-  F. Pellegrini, J. Roman  
*Scotch: A Software Package for Static Mapping by Dual Recursive Bipartitioning of Process and Architecture Graphs*  
High-Performance Computing and Networking, p. 493-498
-  F. Manne and R. H. Bisseling  
A Parallel Approximation Algorithm for the Weighted Maximum Matching Problem  
7th International Conference on Parallel Processing and Applied Mathematics (PPAM), 2007,
-  Fiduccia, C. M. and Mattheyses, R. M.  
A Linear-Time Heuristic for Improving Network Partitions  
19th Conference on Design Automation, 1982