Karlsruher Institut für Technologie

Engineering Multilevel Graph Partitioning Algorithms

Peter Sanders, Christian Schulz

Institute for Theoretical Computer Science, Algorithmics II

Overview

- Introduction
- Multilevel Algorithms
- Local Improvement using Flows
- Experimental Evaluation
- Future Work

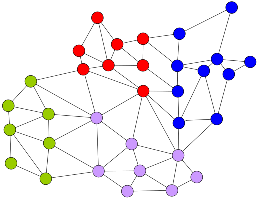

Simulation - FEM

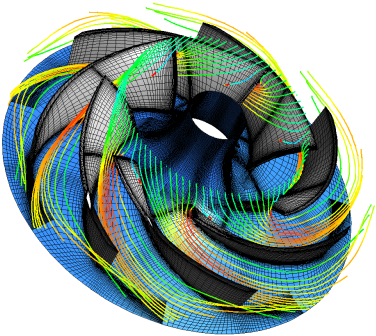

- Simulation space is discretized into a mesh
- Solution of partitial differential equations are approximated by linear equations
- Number of vertices can become quite large $\rightarrow$ time and memory
- Parallel processing required

The Common Parallel Approach

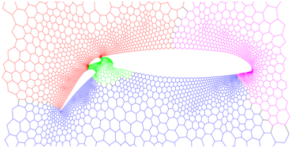

- Mesh partitioned via dual graph
  1. Each volume (data, calculation) is represented by a vertex
  2. Interdependencies are represented by edges
- All PE's get same amount of work
- Communication is expensive

Graph Partitioning Problem:
Partition a graph into (almost) equally sized blocks, such that the number of edges connecting vertices from different blocks is minimal.

ϵ-Balanced Graph Partitioning

Partition graph $G = (V, E, c : V \to \mathbf{R}_{>0}, \omega : E \to \mathbf{R}_{>0})$ into k disjoint blocks s.t.

- total node weight of each block $\leq \dfrac{1+\epsilon}{k}$ total node weight
- total weight of cut edges as small as possible

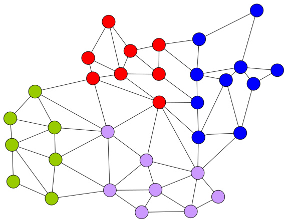

Applications:
finite element simulations, VLSI design, route planning, ...

Multi-Level Graph Partitioning

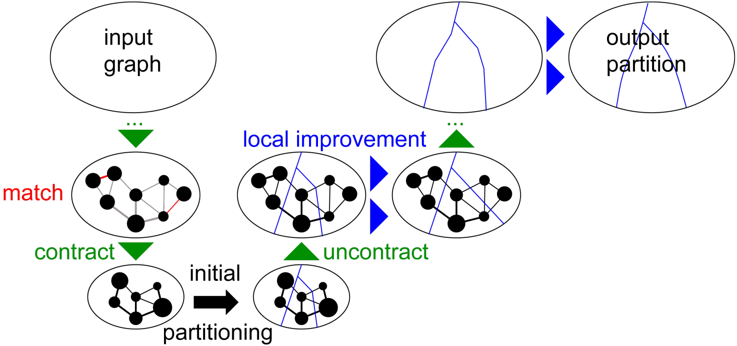

Successful in existing systems:
DiBaP, Chaco, Jostle, Metis, Scotch,. . . **KaPPa**, **KaSPar**

Flows as Local Improvement
Construction

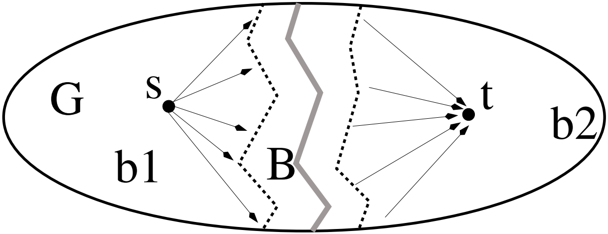

- corridor B, such that each (s, t)-min cut is ϵ-balanced cut in G
- e.g. 2 times BFS (left, right)
- stop the BFS, if size would exceed $(1+\epsilon)\frac{n}{2} - \omega(b2)$
- $\Rightarrow \omega(b2_{\text{new}}) \leq \omega(b2) + (1+\epsilon)\frac{n}{2} - \omega(b2)$

Flows as Local Improvement

- cut are optimal in corridor *B*
- ϵ-balanced partitioning NP-hard

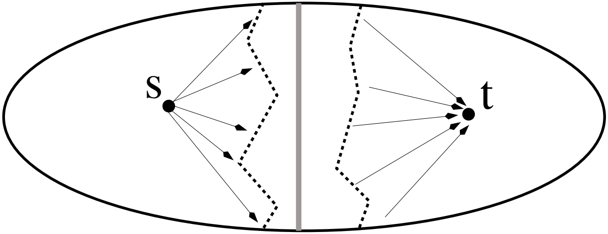

Flows as Local Improvement
Adaptive Search

- search in larger areas for feasable cuts
- adaptively control the size of corridor B
- Most Balanced Minimum Cuts (NP-hard)
- We need: SCC's, DAGs, Closed Vertex Sets, Topological Sorting

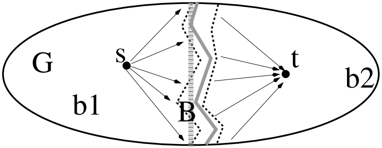

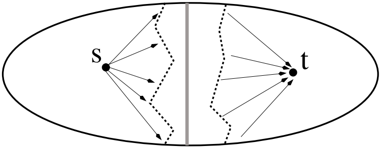

- the maximal upper bound factor is called α'

Local Improvement for *k*-partitions
Using Flows?

on each pair of blocks

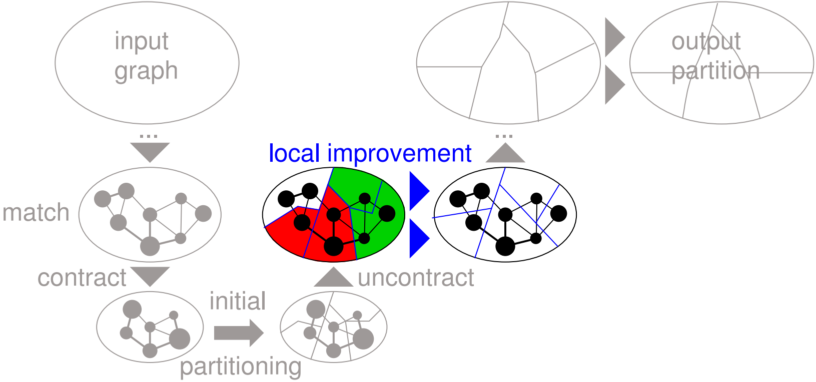

Local Improvement for *k*-partitions

Using Flows?

on each pair of blocks

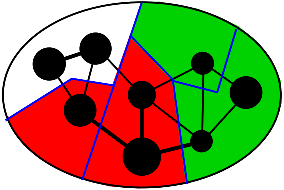

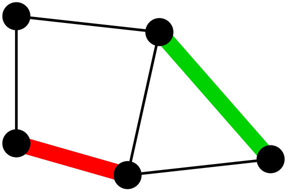

Experiments

Today's Testset

Medium sized instances		
graph	n	m
rgg17	2^{17}	1 457 506
rgg18	2^{18}	3 094 566
Delaunay17	2^{17}	786 352
Delaunay18	2^{18}	1 572 792
bcsstk29	13 992	605 496
4elt	15 606	91 756
fesphere	16 386	98 304
cti	16 840	96 464
memplus	17 758	108 384
cs4	33 499	87 716
pwt	36 519	289 588
bcsstk32	44 609	1 970 092
body	45 087	327 468
t60k	60 005	178 880
wing	62 032	243 088
brack2	62 631	733 118
finan512	74 752	522 240
bel	463 514	1 183 764
nld	893 041	2 279 080
af_shell9	504 855	17 084 020

Experiments

Flows as Local Improvement

Variant	(+Flow, -MB, -FM)				(+Flow, +MB, -FM)				(+Flow, -MB, +FM)				(+Flow, +MB, +FM)			
α'	Avg.	Best.	Bal.	t	Avg.	Best.	Bal.	t	Avg.	Best.	Bal.	t	Avg.	Best.	Bal.	t
16	3 031	2 888	1,025	4,17	2 950	2 841	1,023	3,92	2 802	2 704	1,025	4,30	2 774	2 688	1,023	5,01
8	3 044	2 905	1,025	2,11	2 962	2 855	1,023	2,07	2 806	2 705	1,025	2,41	2 778	2 693	1,023	2,72
4	3 126	2 963	1,024	1,24	3 041	2 933	1,021	1,29	2 825	2 723	1,025	1,62	2 800	2 706	1,022	1,76
2	3 374	3 180	1,022	0,90	3 274	3 107	1,018	0,96	2 869	2 758	1,024	1,31	2 855	2 746	1,021	1,39
1	3 698	3 488	1,018	0,76	3 587	3 410	1,016	0,80	2 926	2 804	1,024	1,19	2 923	2 802	1,023	1,22
(-Flow, -MB, +FM)	2 974	2 851	1,025	1,13												

Effectiveness (+Flow, +MB, -FM)	Avg.	Best.	Bal.
$\alpha' = 1$	3 389	3 351	1,016
2	3 088	3 049	1,017
4	2 922	2 892	1,022
8	2 865	2 841	1,023
16	2 870	2 839	1,023
(-Flow, -MB, +FM)	2 833	2 803	1,025

Effectiveness (+Flow,-MB, +FM)	Avg.	Best.	Bal.
$\alpha' = 1$	2 786	2 759	1,024
2	2 748	2 724	1,024
4	2 721	2 698	1,025
8	2 718	2 690	1,025
16	2 730	2 697	1,025
(-Flow, -MB, +FM)	2 831	2 801	1,025

Effectiveness (+Flow,+MB,+FM)	Avg.	Best.	Bal.
$\alpha' = 1$	2 781	2 754	1,023
2	2 735	2 711	1,021
4	2 702	2 682	1,022
8	2 699	2 675	1,023
16	2 711	2 682	1,022
(-Flow, -MB, +FM)	2 827	2 799	1,025

Outlook

- Global Search (Iterated Multilevel Algorithms)
- incorporate more localized local searches
- → KaFFPa

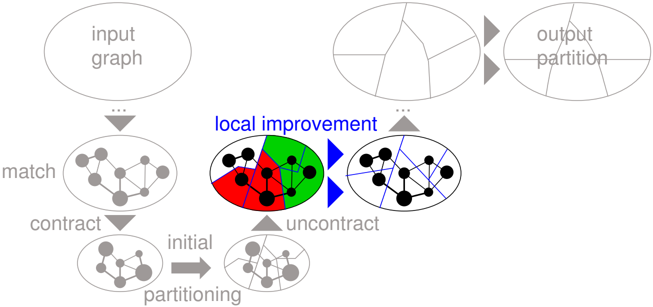

Future Work

Karlsruher Institut für Technologie

1. better initial partitioning for large k
2. integration into meta-heuristics (even parallel)
3. exploit shared memory parallelism
4. output vertex separators
5. graph clustering

10th DIMACS Implementation Challenge
Announcement

- Graph Partitioning and Clustering
- Organizers:
  1. David Bader Georgia Institute of Technology
  2. Peter Sanders and Dorothea Wagner, KIT
- 1st phase: call for instances and applications deadline ”soon”
- 2nd phase: paper submission
- 3rd phase: challenge workshop
- 4th phase: proceedings

Thank you!

Contact: **christian.schulz@kit.edu**