

Shape Optimizing Load Balancing for Parallel Adaptive Numerical Simulations Using MPI*

Henning Meyerhenke

Institute of Theoretical Informatics
Karlsruhe Institute of Technology
Am Fasanengarten 5, 76131 Karlsruhe, Germany

meyerhenke@kit.edu

Abstract. Load balancing is an important requirement for the efficient execution of numerical simulations on parallel computers. In particular when the simulation domain changes over time, the mapping of computational tasks to processors needs to be modified accordingly. Most state-of-the-art libraries addressing this problem are based on graph repartitioning with a parallel variant of the Kernighan-Lin (KL) heuristic. The KL approach has a number of drawbacks, including the optimized metric and solutions with undesirable properties.

Here we further explore the promising diffusion-based multilevel graph partitioning algorithm DIBAP. We describe the evolution of the algorithm and report on its MPI implementation PDIBAP for parallelism with distributed memory. PDIBAP is targeted at small to medium scale parallelism with dozens of processors. The presented experiments use graph sequences that imitate adaptive numerical simulations. They demonstrate the applicability and quality of PDIBAP for load balancing by repartitioning on this scale. Compared to the faster PARMETIS, PDIBAP's solutions often have partitions with fewer external edges and a smaller communication volume in an underlying numerical simulation.

Keywords: Dynamic load balancing, graph partitioning and repartitioning, parallel adaptive numerical simulations, disturbed diffusion.

1 Introduction

Numerical simulations are very important tools in science and engineering for the analysis of physical processes modeled by partial differential equations (PDEs). To make the PDEs solvable on a computer, they are discretized within the simulation domain, e. g., by the finite element method (FEM). Such a discretization yields a mesh, which can be regarded as a graph with geometric (and possibly

* Parts of this work have been published in preliminary form in the proceedings of the 15th International Conference on Parallel and Distributed Systems 2009 [21]. Also, some parts have been presented at the SIAM Conference on Computational Science and Engineering 2011.

other) information. Application areas of such simulations are fluid dynamics, structural mechanics, nuclear physics, and many others [10].

The solutions of discretized PDEs are usually computed by iterative numerical solvers, which have become classical applications for parallel computers. For efficiency reasons the computational tasks, represented by the mesh elements, must be distributed onto the processors evenly. Moreover, neighboring elements of the mesh need to exchange their values in every iteration to update their own value. Due to the high cost of inter-processor communication, neighboring mesh elements should reside on the same processor. A good initial assignment of subdomains to processors can be found by solving the graph partitioning problem (GPP) [29]. The most common GPP formulation for an undirected graph $G = (V, E)$ asks for a division of V into k pairwise disjoint subsets (*parts*) such that all parts are no larger than $(1 + \epsilon) \cdot \lceil \frac{|V|}{k} \rceil$ (for small $\epsilon \geq 0$) and the *edge-cut*, i. e., the total number of edges having their incident nodes in different subdomains, is minimized.

In many numerical simulations some areas of the mesh are of higher interest than others. For instance, during the simulation of the interaction of a gas bubble with a surrounding liquid, one is interested in the conditions close to the boundary of the fluids. Another application among many others is the simulation of the dynamic behavior of biomolecular systems [2]. To obtain an accurate solution, a high resolution of the mesh is required in the areas of interest. To use the available memory efficiently, one has to work with different resolutions in different areas. Moreover, the areas of interest may change during the simulation, which requires *adaptations* in the mesh and may result in undesirable load imbalances. Hence, after the mesh has been adapted, its elements need to be redistributed such that every processor has a similar computational effort again. While this can be done by solving the GPP for the new mesh, the *repartitioning* process not only needs to find new partitions of high quality. Also as few nodes as possible should be moved to other processors since this *migration* causes high communication costs and changes in the local mesh data structure.

Motivation. The most popular graph partitioning and repartitioning libraries (for details see Section 2) use local node-exchanging heuristics like Kernighan-Lin (KL) [17] within a multilevel improvement process to compute solutions with low edge cuts very quickly. Yet, their deployment can have certain drawbacks. First of all, minimizing the edge-cut with these tools does not necessarily mean to minimize the total running time of parallel numerical simulations [35, 12]. While the total communication volume can be minimized by hypergraph partitioning [4], synchronous parallel applications need to wait for the processor computing longest. Hence, the *maximum norm* (i. e., the worst part in a partition) of the simulation's communication costs is of higher importance. Moreover, for some applications, the *shape* of the subdomains plays a significant role. It can be assessed by various measures such as aspect ratio [8], maximum diameter [25], connectedness, or smooth boundaries. Optimizing partition shapes, however, requires additional techniques (e. g., [8, 25, 22]), which are far from being mature. Finally, due to their sequential nature, the most popular repartitioning heuris-

tics are difficult to parallelize—although significant progress has been made (see Section 2).

Our previously developed partitioning algorithm DIBAP aims at computing well-shaped partitions and uses disturbed diffusive schemes to decide not only *how many* nodes move to other parts, but also *which* ones. It is inherently parallel and overcomes many of the above mentioned difficulties, as could be shown experimentally for static graph partitioning [22]. While it is much slower than state-of-the-art partitioners, it often obtains better results.

Contribution. In this work we further explore the disturbed diffusive approach and focus on repartitioning for load balancing. First we present how the implementation of PDIBAP has been improved and adapted for MPI parallel repartitioning. With this implementation we perform various repartitioning experiments with benchmark graph sequences. These experiments are the first using PDIBAP for repartitioning and show the suitability of the disturbed diffusive approach. The average quality of the partitions computed by PDIBAP is clearly better than that of the state-of-the-art repartitioners PARMETIS and parallel JOSTLE, while PDIBAP’s migration volume is usually comparable. It is important to note that PDIBAP’s improvement concerning the partition quality for the graph sequences is even higher than in the case of static partitioning.

2 Related Work

We give a short introduction to the state-of-the-art of practical graph repartitioning algorithms and libraries which only require the adjacency information about the graph and no additional problem-related information. For a broader overview the reader is referred to Schloegel et al. [29]. Some recent advances in related topics can also be found in Boman et al. [3].

2.1 Graph Partitioning

To employ local improvement heuristics effectively, they need to start with a reasonably good initial solution. If such a solution is not provided as input, the multilevel approach [11] is a very powerful technique. It consists of three phases: First, one computes a hierarchy of graphs G_0, \dots, G_l by recursive coarsening in the first phase. G_l ought to be very small in size, but similar in structure to the input graph G_0 . A very good initial solution for G_l is computed in the second phase. After that, the solution is extrapolated to the next-finer graph recursively. In this final phase each extrapolated solution is refined using the desired local improvement algorithm. A very common local improvement algorithm for the third phase of the multilevel process is based on the method by Fiduccia and Mattheyses (FM) [9], a variant of the well-known local search heuristic by Kernighan and Lin (KL) [17] with improved running time. The main idea of both is to exchange nodes between parts in the order of the cost reductions possible, while maintaining balanced partition sizes. After every node has been moved

once, the solution with the best gain is chosen. This is repeated several times until no further improvements are found.

State-of-the-art graph partitioning libraries such as METIS [15, 16] and JOSTLE [36] use KL/FM for local improvement and edge-contractions based on matchings for coarsening. Recently, Holtgrewe et al. [13] presented a parallel library for static partitioning called KAPPA. It attains very good edge cut results, mainly by controlling the multilevel process using so-called edge ratings for approximate matchings. Recently Sanders and Osipov [24] and Sanders and Schulz [26] present new sequential approaches based on a radical multilevel strategy and flow-based local improvement, respectively.

2.2 Load Balancing by Repartitioning

In order to consider both a small edge-cut *and* small migration costs when repartitioning dynamic graphs, different strategies have been explored in the literature. To overcome the limitations of simple scratch-remap and rebalance approaches, Schloegel et al. [30, 31] combine both methods. They propose a multilevel algorithm with three main features. In the local improvement phase, two algorithms are used. On the coarse hierarchy levels, a diffusive scheme takes care of balancing the subdomain sizes. Since this might affect the partition quality negatively, a refinement algorithm is employed on the finer levels. It aims at edge-cut minimization by profitable swaps of boundary vertices.

To address the load balancing problem in parallel applications, distributed versions of the partitioners METIS, JOSTLE, and SCOTCH [32, 37, 6] have been developed. Also, the tools PARKWAY [34], a parallel hypergraph partitioner, and ZOLTAN [5], a suite of load balancing algorithms with focus on hypergraph partitioning, need to be mentioned although they concentrate (mostly) on hypergraphs. An efficient parallelization of the KL/FM heuristic that these parallel (hyper)graph partitioners use is complex due to inherently sequential parts in this heuristic. For example, one needs to ensure that during the KL/FM improvement no two neighboring vertices change their partition simultaneously and destroy data consistency. A coloring of the graph's vertices is used by the parallel libraries PARMETIS [30] and KAPPA [13] for this purpose.

2.3 Diffusive Methods for Shape Optimization

Some applications profit from good partition shapes. As an example, the convergence rate of certain iterative linear solvers can depend on the geometric shape of a partition [8]. That is why in previous work [20, 23] we have developed shape-optimizing algorithms based on diffusion. Before that, repartitioning methods employed diffusion mostly for computing *how much* load needs to be migrated between subdomains [28], not *which* elements should be migrated. Generally speaking, a diffusion problem consists of distributing load from some given seed vertex (or vertices) into the whole graph by iterative load exchanges between neighbor vertices. Typical diffusion schemes have the property to result in the balanced load distribution, in which every node has the same amount of load.

This is one reason why diffusion has been studied extensively for load balancing [38]. In order to distinguish dense from sparse graph regions, our algorithms BUBBLE-FOS/C [23] and the much faster DIBAP [22] (also see Section 3) as well as a combination of KL/FM and diffusion by Pellegrini [25] exploit that diffusion sends load entities faster into densely connected subgraphs.

3 Diffusion-based Repartitioning with DibaP

The algorithm DIBAP, which we have developed and implemented with shared memory parallelism previously [22], is a hybrid multilevel combination of the two (re)partitioning methods BUBBLE-FOS/C and TRUNCCONS, which are both based on disturbed diffusion. We call a diffusion scheme *disturbed* if it is modified such that its steady state does not result in the balanced distribution. Disturbed diffusion schemes can be helpful to determine if two graph nodes or regions are densely connected to each other, i. e., if they are connected by many paths of small length. Before we explain the whole algorithm DIBAP, we describe its two main components for (re-)partitioning in more detail.

3.1 Bubble-FOS/C

In contrast to Lloyd’s related k -means algorithm [18], BUBBLE-FOS/C partitions or clusters graphs instead of geometric inputs. Given a graph $G = (V, E)$ and $k \geq 2$, initial partition representatives (centers) are chosen in the first step of the algorithm, one center for each of the k parts. All remaining vertices are assigned to their closest center vertex. While for k -means one usually uses Euclidean distance, BUBBLE-FOS/C employs the disturbed diffusion scheme FOS/C [23] as distance measure (or, more precisely, as similarity measure). The similarity of a node v to a non-empty node subset S is computed by solving the linear system $\mathbf{L}w = d$ for w , where \mathbf{L} is the Laplacian matrix of the graph and d a suitably chosen vector that disturbs the underlying diffusion system. After the assignment step, each part computes its new center for the next iteration – again using FOS/C, but with a different right-hand side vector d . The two operations *assigning vertices to parts* and *computing new centers* are repeated alternately a fixed number of times or until a stable state is reached. Each operation requires the solution of k linear systems, one for each partition.

It turns out that this iteration of two alternating operations yields very good partitions. Apart from the distinction of dense and sparse regions, the final partitions are very compact and have short boundaries. However, the repeated solution of linear systems makes BUBBLE-FOS/C slow.

3.2 TruncCons

The algorithm TRUNCCONS [22] (for *truncated consolidations*) is also an iterative method for the diffusion-based local improvement of partitions, but it is

much faster than BUBBLE-FOS/C. Within each TRUNCCONS iteration, the following is performed independently for each partition π_c : First, the initial load vector $w^{(0)}$ is set. Nodes of π_c receive an equal amount of initial load $|V|/|\pi_c|$, while the other nodes' initial load is set to 0. Then, this load is distributed within the graph by performing a small number ψ of FOS (first order diffusion scheme) [7] iterations. The final load vector w is computed as $w = \mathbf{M}^\psi w^{(0)}$, where $\mathbf{M} = \mathbf{I} - \alpha \mathbf{L}$ denotes the diffusion matrix [7] of G . A common choice for α is $\alpha := \frac{1}{(1+\deg(G))}$. The computation $w = \mathbf{M}^\psi w^{(0)}$ could be realized by ψ matrix-vector products. A more localized view of its realization is given by iterative load exchanges on each vertex v with its.

After the load vectors have been computed this way independently for all k parts, each node v is assigned to the partition it has obtained the highest load from. This completes one TRUNCCONS iteration, which can be repeated several times (the total number is denoted by Λ subsequently) to facilitate sufficiently large movements of the parts. A node with the same amount of load as all its neighbors does not change its load in the next FOS iteration. Due to the choice of initial loads, such an *inactive* node is a certain distance away from the partition boundary. By avoiding load computations for inactive nodes, we can restrict the computational effort to areas close to the partition boundaries.

3.3 The Hybrid Algorithm DibaP

The main components of DIBAP are depicted in Figure 1. To build a multilevel hierarchy, the fine levels are coarsened (1) by approximate maximum weight matchings. Once the graphs are sufficiently small, the construction mechanism can be changed. In our sequential implementation, we switch the construction mechanism (2) to the more expensive coarsening based on algebraic multigrid (AMG) – for an overview on AMG cf. [33]. This is advantageous regarding running time because, after computing an initial partition (3), BUBBLE-FOS/C is used as local improvement algorithm on the coarse levels (4). Since BUBBLE-FOS/C uses AMG as linear solver, such a hierarchy needs to be built anyway. In our parallel implementation PDIBAP (cf. Section 4), however, due to ease of programming, we decided to coarsen by matchings, use a conjugate gradient solver, and leave the use of AMG for future work. Eventually, the partitions on the fine levels are improved by the local improvement

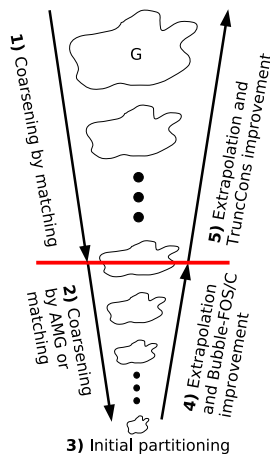


Fig. 1. Sketch of the combined multilevel hierarchy and the corresponding repartitioning algorithms used within DIBAP.

scheme TRUNCCONS. DIBAP includes additional components, e. g., for balancing partition sizes and smoothing partition boundaries, see Section 4.3.

The rationale behind DIBAP can be explained as follows. While BUBBLE-FOS/C computes high-quality graph partitions with good shapes, its similarity measure FOS/C is very expensive to compute compared to established partitioning heuristics. To overcome this problem, we use the simpler process TRUNCCONS, a truly local algorithm to improve partitions generated in a multilevel process. It exploits the observation that, once a reasonably good solution has been found, alterations during a local improvement step take place mostly at the partition boundaries. The disturbing truncation within TRUNCCONS allows for a concentration of the computations around the partition boundaries, where the changes in subdomain affiliation occur. Moreover, since TRUNCCONS is also based on disturbed diffusion, the good properties of the partitions generated by BUBBLE-FOS/C are mostly preserved.

4 PDibAP: Parallel DibaP for Repartitioning

In this section we describe our parallel implementation of DIBAP using MPI. In particular we highlight some differences to the sequential (and thread-parallel) version used for static partitioning [22].

4.1 Distributed Memory Parallelism

The foundation of our PDIBAP implementation (data structure, linear algebra routines, auxiliary functions) is to a large extent based on the code described in more detail in our previous work [23] and in Schamberger’s thesis [27].

As graph data structure PDIBAP employs the standard distributed compressed sparse row (CSR) format with ghost (or halo) vertices. The linear systems within BUBBLE-FOS/C are solved with a conjugate gradient (CG) solver using the traditional domain decomposition approach for distributed parallelism. That means that each system is distributed over all processors and solved by all of them in parallel at the same time, which requires three communication operations per iteration within CG. The TRUNCCONS process is executed in a similar manner. To account for the inactive nodes, however, we do not perform complete matrix-vector multiplications, but perform local load exchanges only if an active vertex is involved. Both CG and TRUNCCONS require a halo update after each iteration. This communication routine is rather expensive, so that the number of iterations should be kept small. The linear algebra routines within PDIBAP do not make use of external libraries. This is due to the fact that the solution process in BUBBLE-FOS/C is very specialized [23, 27].

4.2 Repartitioning

So far, PDIBAP is targeted at repartitioning dynamic graphs. The option for parallel static partitioning is still in its infancy due to a limitation in the multilevel process, which we explain later on in this section.

When PDIBAP is used for repartitioning instead of partitioning, one part of its input is an initial partition. Based on this partition, the graph is distributed onto the processors. We can assume that this partition is probably more unbalanced than advisable. It might also contain some undesirable artifacts. Nevertheless, its quality is not likely to be extremely bad. It is therefore reasonable to improve the initial partition instead of starting from scratch. Moreover, a refinement limits the number of migrated vertices as well, an important feature of dynamic repartitioning methods.

In particular if the imbalance is higher than allowed, it is advisable to employ the multilevel paradigm. Local improvements on the input graph would not result in sufficiently large movements to a high quality solution. Therefore, a matching hierarchy is constructed until only a few thousand nodes remain in the coarsest graph. So far, only edges whose endpoints lie in the same part are considered to be part of the matching. This simplifies the parallel implementation and is a viable approach when repartitioning. For static partitioning, however, edges in the cut between parts on different processors should be considered as well. Otherwise, the multilevel hierarchy contains only a few levels after which no more edges are found for the matching. The development and/or integration of a more general matching is part of future work.

After constructing the hierarchy, the initial partition is projected downwards the hierarchy onto the coarsest level. On the coarsest level the graph is repartitioned with BUBBLE-FOS/C, starting with the projected initial solution. Going up the multilevel hierarchy recursively, the result is then improved with either BUBBLE-FOS/C or TRUNCCONS, depending on the size of the level. After the refinement, the current solution is extrapolated to the next level until the process stops at the input level. Sometimes the matching algorithm has hardly coarsened a level. This happens for example to avoid star-like subgraphs with strongly varying node degrees. Limited coarsening results in two very similar adjacent levels. Local improvement with TRUNCCONS on both of these levels would result in similar solutions with an unnecessary running time investment. That is why in such a case TRUNCCONS is skipped on the finer level of the two.

4.3 Balancing Procedures

In general the diffusive processes employed by DIBAP do not guarantee the nearly perfect balance required by numerical simulations (say, for example, no part should be larger than the average part size plus 3%). That is why we employ two balancing procedures within PDIBAP. The first one called **ScaleBalance** is an iterative procedure that tries to determine for every part $1 \leq p \leq k$ a scalar β_p such that the assignment of vertices to parts based on the load vector entries $\beta_p w_p$ results in a balanced partition. More details can be found in Meyerhenke et al. [23, p. 554]. While **ScaleBalance** works surprisingly well in many cases, it also happens that it is not fully effective even after a fairly large number of iterations. Then we employ a second approach, called **FlowBalance**, whose basic idea is described in previous work as well [23, p. 554]. Here we highlight recent changes necessary to adapt to the distributed parallelism in PDIBAP.

First, we solve a load balancing problem on the quotient graph of the partition Π . The quotient graph Q contains a vertex for each part in Π and two vertices are connected by an edge in Q if and only if their corresponding parts share a common boundary in Π . The load balancing problem can be solved with diffusion [14]. The solution yields the migrating flow that balances the partition. Hence, we know *how many* vertices have to be moved from π_i to π_j , let us call this number n_{ij} . It remains to be determined *which* vertices take this move. For quality reasons, this decision should be based on the diffusion values in the respective load vectors. That is why we want to migrate the n_{ij} vertices with the highest values in the load vector w_j .

In our sequential and thread-parallel version of DIBAP, we use a binary heap as priority queue to perform the necessary selection, migration, and resulting updates to the partition. Since parallel priority queues require a considerable effort to obtain good scalability, we opt for a different approach here. For ease of implementation (and because the amount of computation and communication is relatively small), each processor preselects its local vertices with the highest n_{ij} load values in w_j . These preselected load values are sent to processor p_j , which performs a sequential selection. The threshold value found this way is broadcast back to all processors. Finally, all processors assign their vertices whose diffusion loads in w_j is higher than the threshold to part π_j .

This approach might experience problems when the selected threshold value occurs multiple times among the preselected candidate values. In such a case, the next larger candidate value is chosen as threshold. Another problem could be the scheduled order in which migration takes place. It could happen that a processor needs to move a number of vertices that it is about to obtain by a later move. To address this, we employ a conservative approach and move rather fewer vertices than too many. As a compensation, the whole procedure is repeated iteratively until a balanced partition is found.

5 Experiments

Here we present some of our experimental results comparing our PDIBAP implementation to the KL/FM-based load balancers PARMETIS and parallel JOSTLE.

5.1 Benchmark Data

Our benchmark set comprises two types of graph sequences. The first one consists of three smaller graph sequences with 51 frames each, having between approximately $1M$ and $3M$ vertices, respectively. The second group contains two larger sequences of 36 frames each. Each frame in this group has approximately $4.5M$ to $16M$ vertices. These sequences result in 50 and 35 repartitioning steps, respectively. We choose to (re)partition the smaller sequences into $k = 36$ and $k = 60$ parts, while the larger ones are divided into $k = 60$ and $k = 84$ parts. These values have been chosen as multiples of 12 because one of our test machines has 12 cores per node (the other one contains quad-core CPUs).

All graphs of these five sequences have a two-dimensional geometry and have been generated to resemble adaptive numerical simulations such as those occurring in computational fluid dynamics. A visual impression of some of the data (in smaller versions) is available in previous work [23, p. 562f.]. The graph of frame $i + 1$ in a given sequence is obtained from the graph of frame i by changes restricted to local areas. As an example, some areas are coarsened, whereas others are refined. These changes are in most cases due to the movement of an object in the simulation domain and often result in unbalanced subdomain sizes. For more details the reader is referred to Marquardt and Schamberger [19], who have provided the generator for the sequence data.¹ Some of these frames are also part of the archive of the 10th DIMACS Implementation Challenge [1].

5.2 Hardware and Software Settings

We have conducted our experiments on a cluster with 60 Fujitsu RX200S6 nodes each having 2 Intel Xeon X5650 processors at 2.66 GHz (results in 12 compute cores per node). Moreover, each node has 36 GB of main memory. The interconnect is InfiniBand HCA 4x SDR HCA PCI-e, the operating system Cent OS 5.4. PDIBAP is implemented in C/C++. PDIBAP as well as PARMETIS and parallel JOSTLE have been compiled with Intel C/C++ compiler 11.1 and MVAPICH2 1.5.1 as MPI library.

The main parameters controlling the running time and quality of the DIBAP algorithm are the number of iterations in the (re)partitioning algorithms BUBBLE-FOS/C and TRUNCCONS. For our experiments we perform 3 iterations within BUBBLE-FOS/C, with one `AssignPartition` and one `ComputeCenters` operation, respectively. The faster local approach TRUNCCONS is used on all multilevel hierarchy levels with graph sizes above 12,000 vertices. For TRUNCCONS, the parameter settings $\lambda = 9$ and $\psi = 14$ for the outer and inner iteration, respectively. These settings provide a good trade-off between running time and quality. The allowed imbalance is set to the default value 3% for all tools.

5.3 Results

In addition to the graph partitioning metrics edge-cut and communication volume (of the underlying application based on the computed partition), we are here also interested in migration costs. These costs result from data changing their processor after repartitioning. We count the number of nodes that change their subdomain from one frame to the next as a measure of these costs. One could also assign cost weights to the partitioning objectives and the migration volume to evaluate the linear combination of both. Since these weights depend both on the underlying application and the parallel architecture, we have not pursued this here. We compare PDIBAP to the state-of-the-art repartitioning tools PARMETIS and parallel JOSTLE. Both competitors are mainly based on

¹ Some of the input data can be downloaded from the website <http://www.upb.de/cs/henningm/graph.html>.

the node-exchanging KL heuristic for local improvement. The load balancing toolkit ZOLTAN [5], whose integrated KL/FM partitioner is based on the hypergraph concept, is not included in the detailed presentation. Our experiments with it indicate that it is not as suitable for our benchmark set of FEM graphs, in particular because it yields disconnected parts which propagate and worsen in the course of the sequence. We conclude that currently the dedicated graph (as opposed to hypergraph) partitioners seem more suitable for this problem type.

The partitioning quality is measured in our experiments by the edge cut (EC, a summation norm) and the maximum communication volume (CV_{\max}). CV_{\max} is the sum of the maximum incoming communication volume and the maximum outgoing communication volume, taken over all parts, respectively. The values are displayed in Table 1, averaged over the whole sequence and aggregated by the different k . Very similar results are obtained for the geometric mean in nearly all cases, which is why we do not show these data as well. The migration costs are recorded in both norms and shown for each sequence (again aggregated) in Table 2. Missing values for parallel JOSTLE (—) indicate program crashes on the corresponding instance(s).

Table 1. Average edge cut and communication volume (max norm) for repartitionings computed by PARMETIS, JOSTLE, and DIBAP. Lower values are better, best values per instance are written in bold.

Sequence	PARMETIS		Par. JOSTLE		PDIBAP	
	EC	CV_{\max}	EC	CV_{\max}	EC	CV_{\max}
biggerlowtric	11873.5	1486.7	9875.1	1131.9	8985.5	981.8
biggerbubbles	16956.8	2205.3	14113.2	1638.7	12768.3	1443.5
biggertrace	17795.6	2391.1	14121.3	1687.0	12229.2	1367.5
hugetic	34168.5	2903.0	28208.3	2117.6	24974.4	1766.2
hugetrace	54045.8	5239.7	—	—	34147.4	2459.4

The aggregated graph partitioning metrics show that DIBAP is able to compute the best partitions consistently. DIBAP’s advance is highest for the communication volume. With about 12–19% on parallel JOSTLE and about 34–53% on PARMETIS these improvements are clearly higher than the approximately 7% obtained for static partitioning [22], which is due to the fact that parallel KL (re)partitioners often compute worse solutions than their serial counterparts for static partitioning.

The results for the migration volume are not consistent. All tools have a similar amount of best values. The fact that PARMETIS is competitive is slightly surprising when compared to previous results [21], where it compared worse. Also unexpected, PDIBAP shows significantly higher migration costs for the instance biggerbubbles. Figure 2 displays the migration volumes for

Table 2. Average migration volume in the ℓ_1 - and ℓ_∞ -norm for repartitionings computed by PARMETIS, JOSTLE, and DIBAP. Lower values are better, best values per instance are written in bold.

Sequence	PARMETIS		Par. JOSTLE		PDIBAP	
	ℓ_∞	ℓ_1	ℓ_∞	ℓ_1	ℓ_∞	ℓ_1
biggerslowtric	60314.3	606419.1	64252.2	557608.7	65376.1	550427.0
biggerbubbles	77420.0	1249424.3	68865.1	791723.6	93767.5	1328116.1
biggertrace	54131.2	733750.4	49997.8	533809.2	46620.4	613071.2
hugetric	231072.8	2877441.8	244082.5	2932607.6	232382.6	2875302.5
hugetrace	175795.8	3235984.1	–	–	189085.3	3308461.4

each frame within the *slowrot* sequence in the ℓ_∞ -norm. One gets an impression of the different strategies employed by the three programs. While DIBAP has a more constant migration volume, the values for parallel JOSTLE and PARMETIS show a higher amplitude. It depends on the instance which strategy pays off.

These results lead to the conclusion that DIBAP’s implicit optimization with the iterative algorithms BUBBLE-FOS/C and TRUNCCONS focusses more on good partitions than on small migration costs. In some cases the latter objective should receive more attention. As currently no explicit mechanisms for migration optimization are integrated, such mechanisms could be implemented if one finds in other experiments that the migration costs become too high with DIBAP.

It is interesting to note that further experiments indicate a multilevel approach to be indeed necessary in order to produce sufficiently large partition movements that keep up with the movements of the simulation. Partitions generated by multilevel DIBAP are of a noticeably higher quality regarding the graph partitioning metrics than those computed by TRUNCCONS without multilevel approach. Also, using a multilevel hierarchy results in steadier migration costs, which rarely deviate much from the mean.

The running time of the tools for the dynamic graph instances used in this study can be characterized as follows. PARMETIS is the fastest, taking from a fraction of a second up to a few seconds for each frame. Parallel JOSTLE is approximately a factor of 2-3 slower than PARMETIS. PDIBAP, however, is significantly slower than both tools, with an average slowdown of about 25-50 compared to PARMETIS. It requires from a few seconds up to a few minutes for each frame.

We would like to stress that a high repartitioning quality is often very important. Usually, the most time consuming parts of numerical simulations are the numerical solvers. Hence, a reduced communication volume provided by an excellent partitioning can pay off unless the repartitioning time is extremely high. Nevertheless, a further acceleration of shape-optimizing load balancing is of utmost importance. Minutes for each repartitioning step seem to be problematic for some targeted applications.

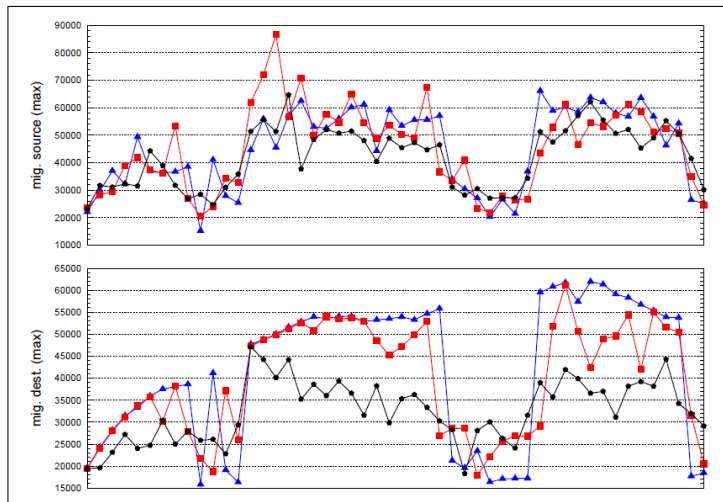


Fig. 2. Number of migrating nodes (ℓ_∞ -norm) in each frame of the biggertrace sequence for PDIBAP (black circle), METIS (blue triangle), and JOSTLE (red square). Lower values are better.

6 Conclusions

With this work we have demonstrated that the shape-optimizing repartitioning algorithm DIBAP based on disturbed diffusion can be a good alternative to traditional KL-based methods for balancing the load in parallel adaptive numerical simulations. In particular, the parallel implementation PDIBAP is very suitable for simulations of small to medium scale, i. e., when the number of vertices and edges in the dynamic graphs are on the order of several millions. While PDIBAP is still significantly slower than the state-of-the-art, it usually computes considerably better solutions w. r. t. edge cut and communication volume. In situations where the quality of the load balancing phase is more important than its running time – e. g., when the computation time between the load balancing phases is relatively high – the use of PDIBAP is expected to pay off.

As part of future work, we aim at an improved multilevel process and faster partitioning methods. It would also be worthwhile to investigate if BUBBLE-FOS/C and TRUNCCONS can be further adapted algorithmically, for example to reduce the dependence on k in the running time.

References

1. David Bader, Henning Meyerhenke, Peter Sanders, and Dorothea Wagner. 10th DIMACS implementation challenge. <http://www.cc.gatech.edu/dimacs10/>, 2012.
2. N. A. Baker, D. Sept, M. J. Holst, and J. A. McCammon. The adaptive multilevel finite element solution of the Poisson-Boltzmann equation on massively parallel computers. *IBM J. of Research and Development*, 45(3.4):427–438, May 2001.

3. Erik G. Boman, Umit V. Catalyurek, Cdric Chevalier, Karen D. Devine, Ilya Safro, and Michael M. Wolf. Advances in parallel partitioning, load balancing and matrix ordering for scientific computing. *J. of Physics: Conference Series*, 180, 2009.
4. U. Catalyurek and C. Aykanat. Hypergraph-partitioning-based decomposition for parallel sparse-matrix vector multiplication. *IEEE Transactions on Parallel and Distributed System*, 10(7):673–693, 1999.
5. U.V. Catalyurek, E.G. Boman, K.D. Devine, D. Bozdog, R.T. Heaphy, and L.A. Riesen. Hypergraph-based dynamic load balancing for adaptive scientific computations. In *Proc. of the 21st Intl. Parallel and Distributed Processing Symposium (IPDPS'07)*. IEEE Computer Society, 2007. Best Algorithms Paper Award.
6. C. Chevalier and F. Pellegrini. Pt-scotch: A tool for efficient parallel graph ordering. *Parallel Comput.*, 34(6-8):318–331, 2008.
7. G. Cybenko. Dynamic load balancing for distributed memory multiprocessors. *Parallel and Distributed Computing*, 7:279–301, 1989.
8. R. Diekmann, R. Preis, F. Schlimbach, and C. Walshaw. Shape-optimized mesh partitioning and load balancing for parallel adaptive FEM. *Parallel Computing*, 26:1555–1581, 2000.
9. C. M. Fiduccia and R. M. Mattheyses. A linear-time heuristic for improving network partitions. In *Proceedings of the 19th Conference on Design automation (DAC'82)*, pages 175–181. IEEE Press, 1982.
10. G. Fox, R. Williams, and P. Messina. *Parallel Computing Works!* Morgan Kaufmann, 1994.
11. B. Hendrickson and R. Leland. A multi-level algorithm for partitioning graphs. In *Proceedings Supercomputing '95*, page 28 (CD). ACM Press, 1995.
12. Bruce Hendrickson and Tamara G. Kolda. Graph partitioning models for parallel computing. *Parallel Comput.*, 26(12):1519–1534, 2000.
13. Manuel Holtgrewe, Peter Sanders, and Christian Schulz. Engineering a scalable high quality graph partitioner. In *IPDPS*, pages 1–12. IEEE, 2010.
14. Y. F. Hu and R. F. Blake. An improved diffusion algorithm for dynamic load balancing. *Parallel Computing*, 25(4):417–444, 1999.
15. George Karypis and Vipin Kumar. *MeTiS: A Software Package for Partitioning Unstructured Graphs, Partitioning Meshes, and Computing Fill-Reducing Orderings of Sparse Matrices, Version 4.0*. Univ. of Minnesota, Minneapolis, MN, 1998.
16. George Karypis and Vipin Kumar. Multilevel k-way partitioning scheme for irregular graphs. *Journal of Parallel and Distributed Computing*, 48(1):96–129, 1998.
17. B. W. Kernighan and S. Lin. An efficient heuristic for partitioning graphs. *Bell Systems Technical Journal*, 49:291–308, 1970.
18. Stuart P. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–136, 1982.
19. O. Marquardt and S. Schamberger. Open benchmarks for load balancing heuristics in parallel adaptive finite element computations. In *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications, (PDPTA '05)*, pages 685–691. CSREA Press, 2005.
20. H. Meyerhenke and S. Schamberger. Balancing parallel adaptive FEM computations by solving systems of linear equations. In *Proceedings of the 11th International Euro-Par Conference*, volume 3648 of *Lecture Notes in Computer Science*, pages 209–219. Springer-Verlag, 2005.
21. Henning Meyerhenke. Dynamic load balancing for parallel numerical simulations based on repartitioning with disturbed diffusion. In *Proc. Internatl. Conference on Parallel and Distributed Systems (ICPADS'09)*, pages 150–157. IEEE Computer Society, 2009.

22. Henning Meyerhenke, Burkhard Monien, and Thomas Sauerwald. A new diffusion-based multilevel algorithm for computing graph partitions. *Journal of Parallel and Distributed Computing*, 69(9):750–761, 2009. Best Paper Awards and Panel Summary: IPDPS 2008.
23. Henning Meyerhenke, Burkhard Monien, and Stefan Schamberger. Graph partitioning and disturbed diffusion. *Parallel Computing*, 35(10–11):544–569, 2009.
24. Vitaly Osipov and Peter Sanders. n -level graph partitioning. In *Proc. 18th Annual European Symposium on Algorithms (ESA'10)*, pages 278–289, 2010.
25. François Pellegrini. A parallelisable multi-level banded diffusion scheme for computing balanced partitions with smooth boundaries. In *Proc. 13th International Euro-Par Conference*, volume 4641 of *LNCS*, pages 195–204. Springer-Verlag, 2007.
26. Peter Sanders and Christian Schulz. Engineering multilevel graph partitioning algorithms. In *Proc. 19th Annual European Symposium on Algorithms (ESA'11)*, pages 469–480, 2011.
27. Stefan Schamberger. *Shape Optimized Graph Partitioning*. PhD thesis, Universität Paderborn, 2006.
28. K. Schloegel, G. Karypis, and V. Kumar. Wavefront diffusion and LMSR: Algorithms for dynamic repartitioning of adaptive meshes. *IEEE Transactions on Parallel and Distributed Systems*, 12(5):451–466, 2001.
29. K. Schloegel, G. Karypis, and V. Kumar. Graph partitioning for high performance scientific simulations. In *The Sourcebook of Parallel Computing*, pages 491–541. Morgan Kaufmann, 2003.
30. Kirk Schloegel, George Karypis, and Vipin Kumar. Multilevel diffusion schemes for repartitioning of adaptive meshes. *Journal of Parallel and Distributed Computing*, 47(2):109–124, 1997.
31. Kirk Schloegel, George Karypis, and Vipin Kumar. A unified algorithm for load-balancing adaptive scientific simulations. In *Proceedings of Supercomputing 2000*, page 59 (CD). IEEE Computer Society, 2000.
32. Kirk Schloegel, George Karypis, and Vipin Kumar. Parallel static and dynamic multi-constraint graph partitioning. *Concurrency and Computation: Practice and Experience*, 14(3):219–240, 2002.
33. Klaus Stüben. An introduction to algebraic multigrid. In U. Trottenberg, C. W. Oosterlee, and A. Schüller, editors, *Multigrid*, pages 413–532. Academic Press, 2000. Appendix A.
34. Aleksandar Trifunović and William J. Knottenbelt. Parallel multilevel algorithms for hypergraph partitioning. *J. Parallel Distrib. Comput.*, 68(5):563–581, 2008.
35. Denis Vanderstraeten, R. Keunings, and Charbel Farhat. Beyond conventional mesh partitioning algorithms and the minimum edge cut criterion: Impact on realistic applications. In *Proceedings of the Seventh SIAM Conference on Parallel Processing for Scientific Computing (PPSC'95)*, pages 611–614. SIAM, 1995.
36. C. Walshaw and M. Cross. Mesh partitioning: a multilevel balancing and refinement algorithm. *SIAM Journal on Scientific Computing*, 22(1):63–80, 2000.
37. C. Walshaw and M. Cross. Parallel optimisation algorithms for multilevel mesh partitioning. *Parallel Computing*, 26(12):1635–1660, 2000.
38. C. Xu and F. C. M. Lau. *Load Balancing in Parallel Computers*. Kluwer, 1997.