

Lawrence Berkeley National Laboratory

Recent Work

Title

Recent advances in graph partitioning

Permalink

<https://escholarship.org/uc/item/6br8x3gz>

ISBN

9783319494869

Authors

Buluç, A
Meyerhenke, H
Safro, I
et al.

Publication Date

2016-11-01

DOI

10.1007/978-3-319-49487-6_4

Peer reviewed

Recent Advances in Graph Partitioning

Aydin Buluç¹, Henning Meyerhenke², Ilya Safro³, Peter Sanders², and Christian Schulz²

¹ Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, USA

² Institute of Theoretical Informatics, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany

³ School of Computing, Clemson University, Clemson SC, USA

Abstract. We survey recent trends in practical algorithms for balanced graph partitioning, point to applications and discuss future research directions.

1 Introduction

Graphs are frequently used by computer scientists as abstractions when modeling an application problem. Cutting a graph into smaller pieces is one of the fundamental algorithmic operations. Even if the final application concerns a different problem (such as traversal, finding paths, trees, and flows), partitioning large graphs is often an important subproblem for complexity reduction or parallelization. With the advent of ever larger instances in applications such as scientific simulation, social networks, or road networks, *graph partitioning* (GP) therefore becomes more and more important, multifaceted, and challenging. The purpose of this paper is to give a structured overview of the rich literature, with a clear emphasis on explaining key ideas and discussing recent work that is missing in other overviews. For a more detailed picture on how the field has evolved previously, we refer the interested reader to a number of surveys. Bichot and Siarry [BS11] cover studies on GP within the area of numerical analysis. This includes techniques for GP, hypergraph partitioning and parallel methods. The book discusses studies from a combinatorial viewpoint as well as several applications of GP such as the air traffic control problem. Schloegel et al. [SKK03] focus on fast graph partitioning techniques for scientific simulations. In their account of the state of the art in this area around the turn of the millennium, they describe geometric, combinatorial, spectral, and multilevel methods and how to combine them for static partitioning. Load balancing of dynamic simulations, parallel aspects, and problem formulations with multiple objectives or constraints are also considered. Monien et al. [MPS07] discuss heuristics and approximation algorithms used in the multilevel GP framework. In their description they focus mostly on coarsening by matching and local search by node-swapping heuristics. Kim et al. [KHKM11] cover genetic algorithms.

Our survey is structured as follows. First, Section 2 introduces the most important variants of the problem and their basic properties such as NP-hardness. Then Section 3 discusses exemplary applications including parallel processing, road networks, image processing, VLSI design, social networks, and bioinformatics. The core of this overview concerns the solution methods explained in Sections 4–7. They involve a surprising variety of techniques. We begin in Section 4 with basic, global methods that “directly” partition the graph. This ranges from very simple algorithms based on breadth first search to sophisticated combinatorial optimization methods that find exact solutions for small instances. Also methods from computational

geometry and linear algebra are being used. Solutions obtained in this or another way can be improved using a number of heuristics described in Section 5. Again, this ranges from simple-minded but fast heuristics for moving individual nodes to global methods, e.g., using flow or shortest path computations. The most successful approach to partitioning large graphs – the multilevel method – is presented in Section 6. It successively contracts the graph to a more manageable size, solves the base instance using one of the techniques from Section 4, and – using techniques from Section 5 – improves the obtained partition when uncontracting to the original input. Metaheuristics are also important. In Section 7 we describe evolutionary methods that can use multiple runs of other algorithms (e.g., multilevel) to obtain high quality solutions. Thus, the best GP solvers orchestrate multiple approaches into an overall system. Since all of this is very time consuming and since the partitions are often used for parallel computing, parallel aspects of GP are very important. Their discussion in Section 8 includes parallel solvers, mapping onto a set of parallel processors, and migration minimization when repartitioning a dynamic graph. Section 9.2 describes issues of implementation, benchmarking, and experimentation. Finally, Section 10 points to future challenges.

2 Preliminaries

Given a number $k \in \mathbb{N}_{>1}$ and an undirected graph $G = (V, E)$ with *non-negative* edge weights, $\omega : E \rightarrow \mathbb{R}_{>0}$, the *graph partitioning problem* (GPP) asks for a partition Π of V with *blocks* of nodes $\Pi = (V_1, \dots, V_k)$:

1. $V_1 \cup \dots \cup V_k = V$
2. $V_i \cap V_j = \emptyset \quad \forall i \neq j.$

A *balance constraint* demands that all blocks have about equal weights. More precisely, it requires that, $\forall i \in \{1, \dots, k\} : |V_i| \leq L_{\max} := (1 + \epsilon) \lceil |V|/k \rceil$ for some imbalance parameter $\epsilon \in \mathbb{R}_{\geq 0}$. In the case of $\epsilon = 0$, one also uses the term *perfectly balanced*. Sometimes we also use weighted nodes with node weights $c : V \rightarrow \mathbb{R}_{>0}$. Weight functions on nodes and edges are extended to sets of such objects by summing their weights. A block V_i is *overloaded* if $|V_i| > L_{\max}$. A *clustering* is also a partition of the nodes. However, k is usually not given in advance, and the balance constraint is removed. Note that a partition is also a clustering of a graph. In both cases, the *goal* is to minimize or maximize a particular objective function. We recall well-known objective functions for GPP in Section 2.1. A node v is a *neighbor* of node u if there is an edge $\{u, v\} \in E$. If a node $v \in V_i$ has a neighbor $w \in V_j$, $i \neq j$, then it is called *boundary node*. An edge that runs between blocks is also called *cut edge*. The set $E_{ij} := \{\{u, v\} \in E : u \in V_i, v \in V_j\}$ is the set of cut edges between two blocks V_i and V_j . An abstract view of the partitioned graph is the so called *quotient graph* or *communication graph*, where nodes represent blocks, and edges are induced by connectivity between blocks. There is an edge in the quotient graph between blocks V_i and V_j if and only if there is an edge between a node in V_i and a node in V_j in the original, partitioned graph. The *degree* $d(v)$ of a node v is the number of its neighbors. An *adjacency matrix* of a graph is a $|V| \times |V|$ matrix describing node connectivity. The element $a_{u,v}$ of the matrix specifies the weight of

the edge from node u to node v . It is set to zero if there is no edge between these nodes. The *Laplacian matrix* of a graph G is defined as $L = D - A$, where D is the diagonal matrix expressing node degrees, and A is the adjacency matrix. A cycle in a directed graph with negative weight is also called *negative cycle*. A *matching* $M \subseteq E$ is a set of edges that do not share any common nodes, i. e., the graph (V, M) has maximum degree one.

2.1 Objective Functions

In practice, one often seeks to find a partition that minimizes (or maximizes) an objective. Probably the most prominent objective function is to minimize the *total cut*

$$\sum_{i < j} \omega(E_{ij}). \quad (2.1)$$

Other formulations of GPP exist. For instance when GP is used in parallel computing to map the graph nodes to different processors, the *communication volume* is often more appropriate than the cut [HK00]. For a block V_i , the communication volume is defined as $\text{comm}(V_i) := \sum_{v \in V_i} c(v)D(v)$, where $D(v)$ denotes the number of different blocks in which v has a neighbor node, excluding V_i . The *maximum communication volume* is then defined as $\max_i \text{comm}(V_i)$, whereas the *total communication volume* is defined as $\sum_i \text{comm}(V_i)$. The maximum communication volume was used in one subchallenge of the 10th DIMACS Challenge on Graph Partitioning and Graph Clustering [BMSW13]. Although some applications profit from other objective functions such as the communication volume or block shape (formalized by the block’s aspect ratio [DPSW00a]), minimizing the cut size has been adopted as a kind of standard. One reason is that cut optimization seems to be easier in practice. Another one is that for graphs with high structural locality the cut often correlates with most other formulations but other objectives make it more difficult to use a multilevel approach.

There are also GP formulations in which balance is not directly encoded in the problem description but integrated into the objective function. For example, the *expansion* of a non-trivial cut (V_1, V_2) is defined as $\omega(E_{12}) / \min(c(V_1), c(V_2))$. Similarly, the *conductance* of such a cut is defined as $\omega(E_{12}) / \min(\text{vol}(V_1), \text{vol}(V_2))$, where $\text{vol}(S) := \sum_{v \in S} d(v)$ denotes the volume of the set S .

As an extension to the problem, when the application graph changes over time, *repartitioning* becomes necessary. Due to changes in the underlying application, a graph partition may become gradually imbalanced due to the introduction of new nodes (and edges) and the deletion of others. Once the imbalance exceeds a certain threshold, the application should call the repartitioning routine. This routine is to compute a new partition Π' from the old one, Π . In many applications it is favorable to keep the changes between Π and Π' small. Minimizing these changes simultaneously to optimizing Π' with respect to the cut (or a similar objective) leads to multiobjective optimization. To avoid the complexity of the latter, a linear combination of both objectives seems feasible in practice [SKK00].

2.2 Hypergraph Partitioning

A *hypergraph* $H = (V, E)$ is a generalization of a graph in which an edge (usually called *hyperedge* or *net*) can connect any number of nodes. As with graphs, partitioning a hypergraph also means to find an assignment of nodes to different blocks of (mostly) equal size. The objective function, however, is usually expressed differently. A straightforward generalization of the edge cut to hypergraphs is the *hyperedge cut*. It counts the number of hyperedges that connect different blocks. In widespread use for hypergraph partitioning, however, is the so-called $(\lambda - 1)$ metric, $CV(H, \Pi) = \sum_{e \in E} (\lambda(e, \Pi) - 1)$, where $\lambda(e, \Pi)$ denotes the number of distinct blocks connected by the hyperedge e and Π the partition of H 's vertex set.

One drawback of hypergraph partitioning compared to GP is the necessity of more complex algorithms—in terms of implementation and running time, not necessarily in terms of worst-case complexity. Paying this price seems only worthwhile if the underlying application profits significantly from the difference between the graph and the hypergraph model.

To limit the scope, we focus in this paper on GP and forgo a more detailed treatment of hypergraph partitioning. Many of the techniques we describe, however, can be or have been transferred to hypergraph partitioning as well [CA01, PM07, CB⁺07, TK08, FB13]. One important application area of hypergraph partitioning is VLSI design (see Section 3.5).

2.3 Hardness Results and Approximation

Partitioning a graph into k blocks of roughly equal size such that the cut metric is minimized is NP-complete (as decision problem) [HR73, GJS74]. Andreev and Räcke [AR06] have shown that there is no constant-factor approximation for the perfectly balanced version ($\epsilon = 0$) of this problem on general graphs. If $\epsilon \in (0, 1]$, then an $O(\log^2 n)$ factor approximation can be achieved. In case an even larger imbalance $\epsilon > 1$ is allowed, an approximation ratio of $O(\log n)$ is possible [ENRS99]. The minimum weight k -cut problem asks for a partition of the nodes into k non-empty blocks without enforcing a balance constraint. Goldschmidt et al. [GH94] proved that, for a fixed k , this problem can be solved optimally in $O(n^{k^2})$. The problem is NP-complete [GH94] if k is not part of the input.

For the unweighted minimum bisection problem, Feige et al. [FK02] have shown that there is an $O(\log^{1.5} n)$ approximation algorithm and an $O(\log n)$ approximation for minimum bisection on planar graphs. The bisection problem is efficiently solvable if the balance constraint is dropped – in this case it is the minimum cut problem. Wagner et al. [WW93] have shown that the minimum bisection problem becomes harder the more the balance constraint is tightened towards the perfectly balanced case. More precisely, if the block weights are bounded from below by a constant, i. e., $|V_i| \geq C$, then the problem is solvable in polynomial time. The problem is NP-hard if the block weights are constrained by $|V_i| \geq \alpha n^\delta$ for some $\alpha, \delta > 0$ or if $|V_i| = \frac{n}{2}$. The case $|V_i| \geq \alpha \log n$ for some $\alpha > 0$ is open. Note that the case $|V_i| \geq \alpha n^\delta$ also implies that the general GPP with similar lower bounds on the block weights is NP-hard.

If the balance constraint of the problem is dropped and one uses a different objective function such as sparsest cut, then there are better approximation algorithms. The sparsest cut objective combines cut and balance into a single objective function. For general graphs

and the sparsest cut metric, Arora et al. [ARV04,AHK10] achieve an approximation ratio of $O(\sqrt{\log n})$ in $\tilde{O}(n^2)$ time.

Being of high theoretical importance, most of the approximation algorithms are not implemented, and the approaches that implement approximation algorithms are too slow to be used for large graphs or are not able to compete with state-of-the-art GP solvers. Hence, mostly heuristics are used in practice.

3 Applications of Graph Partitioning

We now describe some of the applications of GP. For brevity this list is not exhaustive.

3.1 Parallel Processing

Perhaps the canonical application of GP is the distribution of work to processors of a parallel machine. Scientific computing applications such as sparse direct and iterative solvers extensively use GP to ensure load balance and minimize communication. When the problem domain does not change as the computation proceeds, GP can be applied once in the beginning of the computation. This is known as static partitioning.

Periodic repartitioning, explained in Section 2.1, proved to be useful for scientific computing applications with evolving computational domains such as Adaptive Mesh Refinement (AMR) or volume rendering [ACFK07]. The graph model can be augmented with additional edges and nodes to model the migration costs, as done for parallel direct volume rendering of unstructured grids [ACFK07], an important problem in scientific visualization.

Parallel Graph Computations. GP is also used to partition graphs for parallel processing, for problems such as graph eigenvalue computations [BDR13], breadth-first search [BM12], triangle listing [CC11], PageRank and connected components [SW13]. In computationally intensive graph problems, such as finding the eigenvectors and eigenvalues of graphs, multilevel methods that are tailored to the characteristics of real graphs are suitable [ARK06].

Mesh Partitioning. A *mesh* or *grid* approximates a geometric domain by dividing it into smaller subdomains. Hendrickson defines it as “the scaffolding upon which a function is decomposed into smaller pieces” [Hen98]. Mesh partitioning involves mapping the subdomains of the mesh to processors for parallel processing, with the objective of minimizing communication and load imbalance. A partial differential equation (PDE) that is discretized over a certain grid can be solved by numerous methods such as the finite differences method or the finite elements method. The discretization also defines a system of linear equations that can be represented by a sparse matrix. While it is always possible to use that sparse matrix to do the actual computation over the mesh or grid, sometimes this can be wasteful when the matrix need not be formed explicitly. In the absence of an explicit sparse matrix, the GP solvers first define a graph from the mesh. The right mesh entity to use as the nodes of the graph can be ambiguous and application dependent. Common choices are mesh nodes, groups of mesh nodes that need to stay together, and the dual of mesh nodes. Choosing groups of mesh nodes (such as small regular meshes [FKS⁺12]) with appropriate weighting as graph nodes makes GP cost effective for large problem sizes when the overhead for per-node partitioned

graphs would be too big. Recent work by Zhou et al. [ZS⁺10] gives a thorough treatment of extreme-scale mesh partitioning and dynamic repartitioning using graph models. A variety of solution methodologies described in Section 6, such as the multilevel and geometric methods, has been successfully applied to mesh partitioning.

3.2 Complex Networks

In addition to the previously mentioned task of network data distribution across a cluster of machines for fast parallel computations, complex networks introduced numerous further applications of GPP. A common task in these applications is to identify groups of similar entities whose similarity and connectivity is modeled by the respective networks. The quality of the localizations is quantified with different domain-relevant objectives. Many of them are based on the principle of finding groups of entities that are weakly connected to the rest of the network. In many cases such connectivity also represents similarity. In the context of optimization problems on graphs, by complex networks we mean weighted graphs with non-trivial structural properties that were created by real-life or modelling processes [New10]. Often, models and real-life network generation processes are not well understood, so designing optimization algorithms for such graphs exhibit a major bottleneck in many applications.

Power Grids. Disturbances and cascading failures are among the central problems in power grid systems that can cause catastrophic blackouts. Splitting a power network area into self-sufficient islands is an approach to prevent the propagation of cascading failures [LRJL05]. Often the cut-based objectives of the partitioning are also combined with the load shedding schemes that enhance the robustness of the system and minimize the impact of cascading events [LL09]. Finding vulnerabilities of power systems by GPP has an additional difficulty. In some applications, one may want to find more than one (nearly) minimum partitioning because of the structural difference between the solutions. Spectral GP (see Section 4.2) is also used to detect contingencies in power grid vulnerability analysis by splitting the network into regions with excess generation and excess load [DLL⁺05].

Geographically Embedded Networks. Recent advances of location-aware devices (such as GPS) stimulated a rapid growth of streaming spatial network data that has to be analyzed by extremely fast algorithms. These networks model entities (nodes) tied to geographic places and links that represent flows such as migrations, vehicle trajectories, and activities of people [DG12]. In problems related to spatial data and geographical networks, the cut-based objective of GP (and clustering) is often reinforced by the spatial contiguity constraints.

Biological Networks. Many complex biological systems can be modeled by graph-theoretic representations. Examples include protein-protein interactions, and gene co-expression networks. In these networks nodes are biological entities (such as genes and proteins) and edges correspond to their common participation in some biological process. Such processes can vary from simple straightforward interactions (such as protein-protein interaction and gene-gene co-expression) to more complex relationships in which more than two entities are involved. Partitioning and clustering of such networks may have several goals. One of them is related to data reduction given an assumption that clustered nodes behave biologically similarly to

each other. Another one is the detection of some biological processes by finding clusters of involved nodes. For details see [JS08, Mon10].

Social Networks. Identification of community structure is among the most popular topics in social network science. In contrast to the traditional GPP, community detection problems rarely specify the number of clusters *a priori*. Notwithstanding this difference, GP methods contributed a lot of their techniques to the community detection algorithms [For09]. Moreover, GP solvers are often used as first approximations for them. We refer the reader to examples of methods where GP is used for solving the community detection problem [New13].

3.3 Road Networks

GP is a very useful technique to speed up route planning [Lau04, MSS⁺07, KLSV10, DGPW11, LS12, DW13]. For example, edges could be road segments and nodes intersections.⁴

Lauther [Lau04] introduced the arc-flags algorithm, which uses a geometric partitioning approach as a preprocessing step to reduce the search space of Dijkstra’s algorithm. Möhring et al. [MSS⁺07] improved this method in several ways. Using high quality graph partitions turns out to be one key improvement here since this reduces the preprocessing cost drastically. One reason is that road networks can be partitioned using surprisingly small cuts but these are not easy to find.

Schulz et al. [SWZ02] propose a multilevel algorithm for routing based on precomputing connections between border nodes of a graph partition. This was one of the first successful speedup technique for shortest paths. It was outclassed later by other hierarchy based methods, and, somewhat surprisingly resurfaced after Delling et al. [DGPW11, DW13] did thorough algorithm engineering for this approach. Again, a key improvement was to use high quality graph partitions. Since the approach excels at fast recomputation of the preprocessing information when the edge weights change, the method is now usually called *customizable route planning*. Luxen and Schieferdecker [LS12] use GP to efficiently compute candidate sets for alternative routes in road networks and Kieritz et al. [KLSV10] parallelize shortest-path preprocessing and query algorithms. Maue et al. [MSM09] show how to use precomputed distances between blocks of a partition to make the search goal directed. Here, block diameter seems more relevant than cut size, however.

3.4 Image Processing

Image segmentation is a fundamental task in computer vision for which GP and clustering methods have become among the most attractive solution techniques. The goal of image segmentation is to partition the pixels of an image into groups that correspond to objects. Since the computations preceding segmentation are often relatively cheap and since the computations after segmentation work on a drastically compressed representation of the image (objects rather than pixels), segmentation is often the computationally most demanding part in an image processing pipeline. The image segmentation problem is not well-posed and can

⁴ Sometimes more complex models are used to model lanes, turn costs etc.

usually imply more than one solution. During the last two decades, graph-based representations of an image became very popular and gave rise to many cut-based approaches for several problems including image segmentation. In this representation each image pixel (or in some cases groups of pixels) corresponds to a node in a graph. Two nodes are connected by a weighted edge if some similarity exists between them. Usually, the criteria of similarity is a small geodesic distance which can result in mesh-like graphs with four or more neighbors for each node. The edge weights represent another measure of (dis)similarity between nodes such as the difference in the intensity between the connected pixels (nodes).

GP can be formulated with different objectives that can explicitly reflect different definitions of the segmented regions depending on the applications. The classical minimum cut formulation of the GP objective (2.1) can lead in practice to finding too small segmented objects. One popular modification of the objective that was adopted in image segmentation, called *normalized cut*, is given by $\text{ncut}(A, B) = \omega(E_{AB})/\text{vol}(A) + \omega(E_{AB})/\text{vol}(B)$. This objective is similar to the conductance objective described in Section 2.1. Many efficient algorithms were proposed for solving GPP with the normalized cut objective. Among the most successful are spectral and multilevel approaches. Another relevant formulation of the partitioning objective which is useful for image segmentation is given by optimizing the isoperimetric ratio for sets [GS06]. For more information on graph partitioning and image segmentation see [PZZ13, CG12].

3.5 VLSI Physical Design

Physical design of digital circuits for very large-scale integration (VLSI) systems has a long history of being one of the most important customers of graph and hypergraph partitioning, often reinforced by several additional domain relevant constraints. The partitioning should be accomplished in a reasonable computation time, even for circuits with millions of modules, since it is one of the bottlenecks of the design process. The goal of the partitioning is to reduce the VLSI design complexity by partitioning it into smaller components (that can range from a small set of field-programmable gate arrays to fully functional integrated circuits) as well as to keep the total length of all the wires short. The typical optimization objective (see (2.1)) is to minimize the total weight of connections between subcircuits (blocks), where nodes are the cells, i.e., small logical or functional units of the circuit (such as gates), and edges are the wires. Because the gates are connected with wires with more than two endpoints, hypergraphs model the circuit more accurately. Examples of additional constraints for the VLSI partitioning include information on the I/O of the circuit, sets of cells that must belong to the same blocks, and maximum cut size between two blocks. For more information about partitioning of VLSI circuits see [KLMH11, CS03].

4 Global Algorithms

We begin our discussion of the wide spectrum of GP algorithms with methods that work with the entire graph and compute a solution directly. These algorithms are often used for smaller graphs or are applied as subroutines in more complex methods such as local search or multilevel algorithms. Many of these methods are restricted to bipartitioning but can be generalized to k -partitioning for example by recursion.

After discussing exact methods in Section 4.1 we turn to heuristic algorithms. Spectral partitioning (Section 4.2) uses methods from linear algebra. Graph growing (Section 4.3) uses breadth first search or similar ways to directly add nodes to a block. Flow computations are discussed in Section 4.4. Section 4.5 summarizes a wide spectrum of geometric techniques. Finally, Section 4.5 introduces *streaming* algorithms which work with a very limited memory footprint.

4.1 Exact Algorithms

There is a large amount of literature on methods that solve GPP optimally. This includes methods dedicated to the bipartitioning case [BCR97,KRC00,Fel05,SST03,Arm07,AFHM08,DGRW12,DW12,FW11,HPZ13,HK99,LR03] and some methods that solve the general GPP [FMDS⁺98,Sen01]. Most of the methods rely on the branch-and-bound framework [LD60].

Bounds are derived using various approaches: Karisch et al. [KRC00] and Armbruster [Arm07] use semi-definite programming, and Sellman et al. [SST03] and Sensen [Sen01] employ multi-commodity flows. Linear programming is used by Brunetta et al. [BCR97], Ferreira et al. [FMDS⁺98], Lisser et al. [LR03] and by Armbruster et al. [AFHM08]. Hager et al. [HPZ13,HK99] formulate GPP in form of a continuous quadratic program on which the branch and bound technique is applied. The objective of the quadratic program is decomposed into convex and concave components. The more complicated concave component is then tackled by an SDP relaxation. Felner et al. [Fel05] and Delling et al. [DGRW12,DW12] utilize combinatorial bounds. Delling et al. [DGRW12,DW12] derive the bounds by computing minimum s - t cuts between partial assignments (A, B) , i. e., $A, B \subseteq V$ and $A \cap B = \emptyset$. The method can partition road networks with more than a million nodes, but its running time highly depends on the bisection width of the graph.

In general, depending on the method used, two alternatives can be observed. Either the bounds derived are very good and yield small branch-and-bound trees but are hard to compute. Or the bounds are somewhat weaker and yield larger trees but are faster to compute. The latter is the case when using combinatorial bounds. On finite connected subgraphs of the two dimensional grid without holes, the bipartitioning problem can be solved optimally in $O(n^4)$ time [FW11]. Recent work by Bevern et al. [vBFSS13] looks at the parameterized complexity for computing balanced partitions in graphs.

All of these methods can typically solve only very small problems while having very large running times, or if they can solve large bipartitioning instances using a moderate amount of time [DGRW12,DW12], highly depend on the bisection width of the graph. Methods that solve the general GPP [FMDS⁺98,Sen01] have immense running times for graphs with up to

a few hundred nodes. Moreover, the experimental evaluation of these methods only considers small block numbers $k \leq 4$.

4.2 Spectral Partitioning

One of the first methods to split a graph into two blocks, spectral bisection, is still in use today. Spectral techniques were first used by Donath and Hoffman [DH72,DH73] and Fiedler [Fie75], and have been improved subsequently by others [Bop87,PSL90,Sim91,HL95b,BS93]. Spectral bisection infers global information of the connectivity of a graph by computing the eigenvector corresponding to the second smallest eigenvalue of the Laplacian matrix L of the graph. This eigenvector z_2 is also known as *Fiedler vector*; it is the solution of a relaxed integer program for cut optimization. A partition is derived by determining the median value \bar{m} in z_2 and assigning all nodes with an entry smaller or equal to \bar{m} to V_1 and all others to V_2 .

The second eigenvector can be computed using a modified Lanczos algorithm [Lan50]. However, this method is expensive in terms of running time. Barnard and Simon [BS93] use a multilevel method to obtain a fast approximation of the Fiedler vector. The algorithmic structure is similar to the multilevel method explained in Section 6, but their method coarsens with independent node sets and performs local improvement with Rayleigh quotient iteration. Hendrickson and Leland [HL95b] extend the spectral method to partition a graph into more than two blocks by using multiple eigenvectors; these eigenvectors are computationally inexpensive to obtain. The method produces better partitions than recursive bisection, but is only useful for the partitioning of a graph into four or eight blocks. The authors also extended the method to graphs with node and edge weights.

4.3 Graph Growing

A very simple approach for obtaining a bisection of a graph is called graph growing [KK98a, GL81]. Most of its variants are based on breadth-first search. Its simplest version works as follows. Starting from a random node v , the nodes are assigned to block V_1 using a breadth-first search (BFS) starting at v . The search is stopped after half of the original node weights are assigned to this block and V_2 is set to $V \setminus V_1$. This method can be combined with a local search algorithm to improve the partition. Multiple restarts of the algorithm are important to get a good solution. One can also try to find a good starting node by looking at a node that has maximal distance from a random seed node [GL81]. Variations of the algorithm always add the node to the block that results in the smallest increase in the cut [KK98a]. An extension to $k > 2$ blocks and with iterative improvement is described in Section 5.5.

4.4 Flows

The well-known max-flow min-cut theorem [FF56] can be used to separate two node sets in a graph by computing a maximum flow and hence a minimum cut between them. This approach completely ignores balance, and it is not obvious how to apply it to the balanced GPP. However, at least for random regular graphs with small bisection width this can be

done [BCLS87]. Maximum flows are also often used as a subroutine. Refer to Section 5.4 for applications to improve a partition and to Section 6.4 for coarsening in the context of the multilevel framework. There are also applications of flow computations when quality is measured by expansion or conductance [LR04, AL08].

4.5 Geometric Partitioning

Partitioning can utilize the coordinates of the graph nodes in space, if available. This is especially useful in finite element models and other geometrically-defined graphs from traditional scientific computing. Here, geometrically “compact” regions often correspond to graph blocks with small cut. Partitioning using nodal coordinates comes in many flavors, such as recursive coordinate bisection (RCB) [Sim91] and inertial partitioning [Wil91, FL93]. In each step of its recursion, RCB projects graph nodes onto the coordinate axis with the longest expansion of the domain and bisects them through the median of their projections. The bisecting plane is orthogonal to the coordinate axis, which can create partitions with large separators in case of meshes with skewed dimensions. Inertial partitioning can be interpreted as an improvement over RCB in terms of worst case performance because its bisecting plane is orthogonal to a plane L that minimizes the moments of inertia of nodes. In other words, the projection plane L is chosen such that it minimizes the sum of squared distances to all nodes.

The random spheres algorithm of Miller et al. [MTV91, GMT98] generalizes the RCB algorithm by stereographically projecting the d dimensional nodes to a random $d + 1$ dimensional sphere which is bisected by a plane through its center point. This method gives performance guarantees for planar graphs, k -nearest neighbor graphs, and other “well-behaved” graphs.

Other representatives of geometry-based partitioning algorithms are space-filling curves [PB94, HW02, Zum03, Bad13] which reduce d -dimensional partitioning to the one-dimensional case. Space filling curves define a bijective mapping from V to $\{1, \dots, |V|\}$. This mapping aims at the preservation of the nodes’ locality in space. The partitioning itself is simpler and cheaper than RCB once the bijective mapping is constructed. A generalization of space-filling curves to general graphs can be done by so-called graph-filling curves [SW04].

A recent work attempts to bring information on the graph structure into the geometry by embedding arbitrary graphs into the coordinate space using a multilevel graph drawing algorithm [KR13]. For a more detailed, albeit not very recent, treatment of geometric methods, we refer the interested reader to Schloegel et al. [SKK03].

4.6 Streaming Graph Partitioning (SGP)

Streaming data models are among the most popular recent trends in big data processing. In these models the input arrives in a data stream and has to be processed on the fly using much less space than the overall input size. SGP algorithms are very fast. They are even faster than multilevel algorithms but give lower solution quality. Nevertheless, many applications that require extremely fast repartitioning methods (such as those that deal with dynamic networks) can still greatly benefit from the SGP algorithms when an initial solution obtained

by a stronger (static data) algorithm is supplied as an initial ordering. For details on SGP we refer the reader to [SK12, TGRV00, NU13].

5 Iterative Improvement Heuristics

Most high quality GP solvers iteratively improve starting solutions. We outline a variety of methods for this purpose, moving from very fine-grained localized approaches to more global techniques.

5.1 Node-swapping Local Search

Local search is a simple and widely used metaheuristic for optimization that iteratively changes a solution by choosing a new one from a neighborhood. Defining the neighborhood and the selection strategy allows a wide variety of techniques. Having the improvement of paging properties of computer programs in mind, Kernighan and Lin [KL70] were probably the first to define GPP and to provide a local search method for this problem. The selection strategy finds the swap of node assignments that yields the largest decrease in the total cut size. Note that this “decrease” is also allowed to be negative. A round ends when all nodes have been moved in this way. The solution is then reset to the best solution encountered in this round. The algorithm terminates when a round has not found an improvement.

A major drawback of the KL method is that it is expensive in terms of asymptotic running time. The implementation assumed in [KL70] takes time $O(n^2 \log n)$ and can be improved to $O(m \max(\log n, \Delta))$ where Δ denotes the maximum degree [Dut93]. A major breakthrough is the modification by Fiduccia and Mattheyses [FM82]. Their carefully designed data structures and adaptations yield the KL/FM local search algorithm, whose asymptotic running time is $O(m)$. Bob Darrow was the first who implemented the KL/FM algorithm [FM82].

Karypis and Kumar [KK98b] further accelerated KL/FM by only allowing boundary nodes to move and by stopping a round when the edge cut does not decrease after x node moves. They improve quality by random tie breaking and by allowing additional rounds even when no improvements have been found.

A highly localized version of KL/FM is considered in [OS10]. Here, the search spreads from a single boundary node. The search stops when a stochastic model of the search predicts that a further improvement has become unlikely. This strategy has a better chance to climb out of local minima and yields improved cuts for the GP solvers KaSPar [OS10] and KaHIP [SS11].

Rather than swapping nodes, Holtgrewe et al. move a single node at a time allowing more flexible tradeoffs between reducing the cut or improving balance [HSS10].

Helpful Sets by Diekmann et al. [DMP95, MS04] introduce a more general neighborhood relation in the bipartitioning case. These algorithms are inspired by a proof technique of Hromkovič and Monien [HM91] for proving upper bounds on the bisection width of a graph. Instead of migrating single nodes, whole sets of nodes are exchanged between the blocks to improve the cut. The running time of the algorithm is comparable to the KL/FM algorithm, while solution quality is often better than other methods [MS04].

5.2 Extension to k -way Local Search

It has been shown by Simon and Teng [ST97] that, due to the lack of global knowledge, recursive bisection can create partitions that are very far away from the optimal partition so that there is a need for k -way local search algorithms. There are multiple ways of extending the KL/FM algorithm to get a local search algorithm that can improve a k -partition.

One early extension of the KL/FM algorithm to k -way local search uses $k(k-1)$ priority queues, one for each type of move (source block, target block) [San89, HL95a]. For a single movement one chooses the node that maximizes the gain, breaking ties by the improvement in balance.

Karypis and Kumar [KK98b] present a k -way version of the KL/FM algorithm that runs in linear time $O(m)$. They use a single global priority queue for all types of moves. The priority used is the maximum local gain, i. e., the maximum reduction in the cut when the node is moved to one of its neighboring blocks. The node that is selected for movement yields the maximum improvement for the objective and maintains or improves upon the balance constraint.

Most current local search algorithms exchange nodes between blocks of the partition trying to decrease the cut size while also maintaining balance. This highly restricts the set of possible improvements. Sanders and Schulz [SS13b, Sch13] relax the balance constraint for node movements but globally maintain (or improve) balance by combining multiple local searches. This is done by reducing the combination problem to finding negative cycles in a graph, exploiting the existence of efficient algorithms for this problem.

5.3 Tabu Search

A more expensive k -way local search algorithm is based on tabu search [Glo89, Glo90], which has been applied to GP by [RPG96, BH11b, BH10, BH11a, GBF11]. We briefly outline the method reported by Galinier et al. [GBF11]. Instead of moving a node exactly once per round, as in the traditional versions of the KL/FM algorithms, specific types of moves are excluded only for a number of iterations. The number of iterations that a move (v , block) is excluded depends on an aperiodic function f and the current iteration i . The algorithm always moves a non-excluded node with the highest gain. If the node is in block A , then the move (v , A) is excluded for $f(i)$ iterations after the node is moved to the block yielding the highest gain, i. e., the node cannot be put back to block A for $f(i)$ iterations.

5.4 Flow Based Improvement

Sanders and Schulz [SS11, SS13a] introduce a max-flow min-cut based technique to improve the edge cut of a given bipartition (and generalize this to k -partitioning by successively looking at pairs of blocks that are adjacent in the quotient graph). The algorithm constructs an s - t flow problem by growing an area around the given boundary nodes/cut edges. The area is chosen such that each s - t cut in this area corresponds to a feasible bipartition of the original graph, i. e., a bipartition that fulfills the balance constraint. One can then apply a max-flow min-cut algorithm to obtain a min-cut in this area and hence a nondecreased cut

between the blocks. There are multiple improvements to extend this method, for example, by iteratively applying the method, searching in larger areas for feasible cuts, or applying a heuristic to output better balanced minimum cuts by using the given max-flow.

5.5 Bubble Framework

Diekmann et al. [DPSW00b] extend graph growing and previous ideas [WCE95] to obtain an iterative procedure called *Bubble framework*, which is capable of partitioning into $k > 2$ *well-shaped* blocks. Some applications profit from good geometric block shapes, e.g., the convergence rate of certain iterative linear solvers.

Graph growing is extended first by carefully selecting k seed nodes that are evenly distributed over the graph. The key property for obtaining a good quality, however, is an iterative improvement within the second and the third step – analogous to Lloyd’s k -means algorithm [Llo82]. Starting from the k seed nodes, k breadth-first searches grow the blocks analogous to Section 4.3, only that the breadth-first searches are scheduled such that the smallest block receives the next node. Local search algorithms are further used within this step to balance the load of the blocks and to improve the cut of the resulting partition, which may result in unconnected blocks. The final step of one iteration computes new seed nodes for the next round. The new center of a block is defined as the node that minimizes the sum of the distances to all other nodes within its block. To avoid their expensive computation, approximations are used. The second and the third step of the algorithm are iterated until either the seed nodes stop changing or no improved partition was found for more than 10 iterations. Figure 5.1 illustrates the three steps of the algorithm. A drawback of the algorithm is its computational complexity $O(km)$.

Subsequently, this approach has been improved by using distance measures that better reflect the graph structure [Sch04, MS05, MMS06]. For example, Schamberger [Sch04] introduced the usage of diffusion as a growing mechanism around the initial seeds and extended the method to weighted graphs. More sophisticated diffusion schemes, some of which have been employed within the Bubble framework, are discussed in Section 5.6.




Fig. 5.1: The three steps of the Bubble framework. Black nodes indicate the seed nodes. On the left hand side, seed nodes are found. In the middle, a partition is found by performing breadth-first searches around the seed nodes and on the right hand side new seed nodes are found.

5.6 Random Walks and Diffusion

A *random walk* on a graph starts on a node v and then chooses randomly the next node to visit from the set of neighbors (possibly including v itself) based on transition probabilities. The latter can for instance reflect the importance of an edge. This iterative process can be repeated an arbitrary number of times. It is governed by the so-called *transition matrix* \mathbf{P} , whose entries denote the edges’ transition probabilities. More details can be found in Lovasz’s random walk survey [Lov93].

Diffusion, in turn, is a natural process describing a substance’s desire to distribute evenly in space. In a discrete setting on graphs, diffusion is an iterative process which exchanges splittable entities between neighboring nodes, usually until all nodes have the same amount. Diffusion is a special random walk; thus, both can be used to identify dense graph regions: Once a random walk reaches a dense region, it is likely to stay there for a long time, before leaving it via one of the relatively few outgoing edges. The relative size of $\mathbf{P}_{u,v}^t$, the probability of a random walk that starts in u to be located on v after t steps, can be exploited for assigning u and v to the same or different clusters. This fact is used by many authors for graph clustering, cf. Schaeffer’s survey [Sch07].

Due to the difficulty of enforcing balance constraints, works employing these approaches for partitioning are less numerous. Meyerhenke et al. [MMS09b] present a similarity measure based on diffusion that is employed within the Bubble framework. This diffusive approach bears some conceptual resemblance to spectral partitioning, but with advantages in quality [MS12]. Balancing is enforced by two different procedures that are only loosely coupled to the actual partitioning process. The first one is an iterative procedure that tries to adapt the amount of diffusion load in each block by multiplying it with a suitable scalar. Underloaded blocks receive more load, overloaded ones less. It is then easier for underloaded blocks to “flood” other graph areas as well. In case the search for suitable scalars is unsuccessful, the authors employ a second approach that extends previous work [WCE97]. It computes a migrating flow on the quotient graph of the partition. The flow value f_{ij} between blocks i and j specifies how many nodes have to be migrated from i to j in order to balance the partition. As a key and novel property for obtaining good solutions, to determine *which* nodes should be migrated in which order, the diffusive similarity values computed before within the Bubble framework are used [MMS09b, Mey08].

Diffusion-based partitioning has been subsequently improved by Pellegrini [Pel07a], who combines KL/FM and diffusion for bipartitioning in the tool Scotch. He speeds up previous approaches by using *band graphs* that replace unimportant graph areas by a single node. An extension of these results to k -way partitioning with further adaptations has been realized within the tools DibaP [MMS09a] and PDibaP for repartitioning [Mey12]. Integrated into a multilevel method, diffusive partitioning is able to compute high-quality solutions, in particular with respect to communication volume and block shape. It remains further work to devise a faster implementation of the diffusive approach without running time dependence on k .

6 Multilevel Graph Partitioning

Clearly the most successful heuristic for partitioning large graphs is the *multilevel graph partitioning* approach. It consists of the three main phases outlined in Figure 6.1: coarsening, initial partitioning, and uncoarsening. The main goal of the coarsening (in many multilevel approaches implemented as *contraction*) phase is to gradually approximate the original problem and the input graph with fewer degrees of freedom. In multilevel GP solvers this is achieved by creating a hierarchy of successively coarsened graphs with decreasing sizes in such a way that cuts in the coarse graphs reflect cuts in the fine graph. There are multiple possibilities to create graph hierarchies. Most methods used today *contract* sets of nodes on the fine level. Contracting $U \subset V$ amounts to replacing it with a single node u with $c(u) := \sum_{w \in U} c(w)$. Contraction (and other types of coarsening) might produce parallel edges which are replaced by a single edge whose weight accumulates the weights of the parallel edges (see Figure 6.2). This implies that balanced partitions on the coarse level represent balanced partitions on the fine level with the same cut value.

Coarsening is usually stopped when the graph is sufficiently small to be *initially partitioned* using some (possibly expensive) algorithm. Any of the basic algorithms from Section 4 can be used for initial partitioning as long as they are able to handle general node and edge weights. The high quality of more expensive methods that can be applied at the coarsest level does not necessarily translate into quality at the finest level, and some GP multilevel solvers rather run several faster but diverse methods repeatedly with different random tie breaking instead of applying expensive global optimization techniques.

Uncoarsening consists of two stages. First, the solution obtained on the coarse level graph is mapped to the fine level graph. Then the partition is improved, typically by using some variants of the improvement methods described in Section 5. This process of uncoarsening and local improvement is carried on until the finest hierarchy level has been processed. One run of this simple coarsening-uncoarsening scheme is also called a *V-cycle* (see Figure 6.1).

There are at least three intuitive reasons why the multilevel approach works so well: First, at the coarse levels we can afford to perform a lot of work per node without increasing the overall execution time by a lot. Furthermore, a single node move at a coarse level corresponds to a big change in the final solution. Hence, we might be able to find improvements easily that would be difficult to find on the finest level. Finally, fine level local improvements are




Fig. 6.1: The multilevel approach to GP. The left figure shows a two-level contraction-based scheme. The right figure shows different chains of coarsening-uncoarsening in the multilevel frameworks.

expected to run fast since they already start from a good solution inherited from the coarse level. Also multilevel methods can benefit from their iterative application (such as chains of V-cycles) when the previous iteration’s solution is used to improve the quality of coarsening. Moreover, (following the analogy to multigrid schemes) the inter-hierarchical coarsening-uncoarsening iteration can also be reconstructed in such way that more work will be done at the coarser levels (see F-, and W-cycles in Figure 6.1, and [Wal04,SS11]). An important technical advantage of multilevel approaches is related to parallelization. Because multilevel approaches achieve a global solution by local processing only (though applied at different levels of coarseness) they are naturally parallelization-schemes friendly.

6.1 Contracting a Single Edge

A minimalistic approach to coarsening is to contract only two nodes connected by a single edge in the graph. Since this leads to a hierarchy with (almost) n levels, this method is called n -level GP [OS10]. Together with a k -way variant of the highly localized local search from Section 5.1, this leads to a very simple way to achieve high quality partitions. Compared to other techniques, n -level partitioning has some overhead for coarsening, mainly because it needs a priority queue and a dynamic graph data structure. On the other hand, for graphs with enough locality (e.g. from scientific computing), the n -level method empirically needs only sublinear work for local improvement.

6.2 Contracting a Matching

The most widely used contraction strategy contracts (large) matchings, i. e., the contracted sets are pairs of nodes connected by edges and these edges are not allowed to be incident to each other. The idea is that this leads to a geometrically decreasing size of the graph and hence a logarithmic number of levels, while subsequent levels are “similar” so that local improvement can quickly find good solutions. Assuming linear-time algorithms on all levels, one then gets linear overall execution time. Conventional wisdom is that a good matching contains many high weight edges since this decreases the weight of the edges in the coarse graph and will eventually lead to small cuts. However, one also wants a certain uniformity in the node weights so that it is not quite clear what should be the objective of the matching algorithm. A successful recent approach is to delegate this tradeoff between edge weights and uniformity to an *edge rating* function [ARK06,HSS10]. For example, the function $f(u, v) = \frac{\omega(\{u,v\})}{c(v)c(u)}$ works very well [HSS10,SS11] (also for n -level partitioning [OS10]). The concept of algebraic distance yields further improved edge ratings [SSS12].

The weighted matching problem itself has attracted a lot of interest motivated to a large extent by its application for coarsening. Although the maximum weight matching problem can be solved optimally in polynomial time, optimal algorithms are too slow in practice.




Fig. 6.2: An example matching and contraction of the matched edges.

There are very fast heuristic algorithms like (Sorted) Heavy Edge Matching, Light Edge Matching, Random Matching, etc. [KK98a,SKK03] that do not give any quality guarantees however. On the other hand, there are (near) linear time matching algorithms that are slightly more expensive but give approximation guarantees and also seem to be more robust in practice. For example, a greedy algorithm considering the edges in order of descending edge weight guarantees half of the optimal edge weight. Preis’ algorithm [Pre99] and the Path Growing Algorithm [DH03] have a similar flavor but avoid sorting and thus achieve linear running time for arbitrary edge weights. The Global Path Algorithm (GPA) [MS07] is a synthesis of Greedy and Path Growing achieving somewhat higher quality in practice and is not a performance bottleneck in many cases. GPA is therefore used in KaHIP [SS11, SS,SS13b]. Linear time algorithms with better approximation guarantee are available [PS04, DH05, MS07, DPS11] and the simplest of them seem practical [MS07]. However, it has not been tried yet whether they are worth the additional effort for GP.

6.3 Coarsening for Scale-free Graphs

Matching-based graph coarsening methods are well-suited for coarsening graphs arising in scientific computing. On the other hand, matching-based approaches can fail to create good hierarchies for graphs with irregular structure. Consider the extreme example that the input graph is a star. In this case, a matching algorithm can contract only one edge per level, which leads to a number of levels that is undesirable in most cases.

Abou-Rjeili and Karypis [ARK06] modify a large set of matching algorithms such that an unmatched node can potentially be matched with one of its neighbors even if it is already matched. Informally speaking, instead of matchings, whole groups of nodes are contracted to create the graph hierarchies. These approaches significantly improve partition quality on graphs having a power-law degree distribution.

Another approach has been presented by Fagginger Auer and Bisseling [AB13]. The authors create graph hierarchies for social networks by allowing pairwise merges of nodes that have the same neighbors and by merging multiple nodes, i. e., collapsing multiple neighbors of a high degree node with this node.

Meyerhenke et al. [MSS14, MSS15] presented an approach that uses a modification of the original label propagation algorithm [RAK07] to compute size-constrained clusterings which are then contracted to compute good multilevel hierarchies for such graphs. The same algorithm is used as a very simple greedy local search algorithm.

Glantz et al. [GMS14] introduce an edge rating based on how often an edge appears in relatively balanced light cuts induced by spanning trees. Intriguingly, this cut-based approach yields partitions with very low communication volume for scale-free graphs.

6.4 Flow Based Coarsening

Using max-flow computations, Delling et al. [DG⁺11] find “natural cuts” separating heuristically determined regions from the remainder of the graph. Components cut by none of these cuts are then contracted reducing the graph size by up to two orders of magnitude. They

use this as the basis of a two-level GP solver that quickly gives very good solutions for road networks.

6.5 Coarsening with Weighted Aggregation

Aggregation-based coarsening identifies nodes on the fine level that survive in the coarsened graph. All other nodes are assigned to these coarse nodes. In the general case of *weighted aggregation*, nodes on a fine level belong to nodes on the coarse level with some probability. This approach is derived from a class of hierarchical linear solvers called Algebraic Multigrid (AMG) methods [CS09, MMS06]. First results on the bipartitioning problem were obtained by Ron et al. in [RWSB05]. As AMG linear solvers have shown, weighted aggregation is important in order to express the likelihood of nodes to belong together. The accumulated likelihoods “smooth the solution space” by eliminating from it local minima that will be detected instantaneously by the local processing at the uncoarsening phase. This enables a relaxed formulation of coarser levels and avoids making hardened local decisions, such as edge contractions, before accumulating relevant global information about the graph.

Weighted aggregation can lead to significantly denser coarse graphs. Hence, only the most efficient AMG approaches can be adapted to graph partitioning successfully. Furthermore one has to avoid unbalanced node weights. In [SSS12] *algebraic distance* [CS11] is used as a measure of connectivity between nodes to obtain sparse and balanced coarse levels of high quality. These principles and their relevance to AMG are summarized in [RSB11].

Lafon and Lee [LL06] present a related coarsening framework whose goal is to retain the spectral properties of the graph. They use matrix-based arguments using random walks (for partitioning methods based on random walks see Section 5.6) to derive approximation guarantees on the eigenvectors of the coarse graph. The disadvantage of this approach is the rather expensive computation of eigenvectors.

7 Evolutionary Methods and Further Metaheuristics

In recent years a number of metaheuristics have been applied to GPP. Some of these works use concepts that have already been very popular in other application domains such as genetic or evolutionary algorithms. For a general overview of genetic/evolutionary algorithms tackling GPP, we refer the reader to the overview paper by Kim et al. [KHKM11]. In this section we focus on the description of hybrid evolutionary approaches that combine evolutionary ideas with the multilevel GP framework [SWC04, BH10, BH11a]. Other well-known metaheuristics such as multi-agent and ant-colony optimization [KSR04, CS06], and simulated annealing [JS98] are not covered here. Neither do we discuss the recently proposed metaheuristics PROBE by Chardaire et al. [CBM07] (a genetic algorithm without selection) and Fusion-Fission by Bichot [Bic07] (inspired by nuclear processes) in detail. Most of these algorithms are able to produce solutions of a very high quality, but only if they are allowed to run for a very long time. Hybrid evolutionary algorithms are usually able to compute partitions with considerably better quality than those that can be found by using a single execution of a multilevel algorithm.

The first approach that combined evolutionary ideas with a multilevel GP solver was by Soper et al. [SWC04]. The authors define two main operations, a combine and a mutation operation. Both operations modify the edge weights of the graph depending on the input partitions and then use the multilevel partitioner Jostle, which uses the modified edge weights to obtain a new partition of the original graph. The combine operation first computes node weight biases based on the two input partitions/parents of the population and then uses those to compute random perturbations of the edge weights which help to mimic the input partitions. While producing partitions of very high quality, the authors report running times of up to one week. A similar approach based on edge weight perturbations is used by Delling et al. [DG⁺11].

A multilevel memetic algorithm for the perfectly balanced graph partition problem, i. e., $\epsilon = 0$, was proposed by Benlic and Hao [BH10, BH11a]. The main idea of their algorithm is that among high quality solutions a large number of nodes will always be grouped together. In their work the partitions represent the individuals. We briefly sketch the combination operator for the case that two partitions are combined. First the algorithm selects two individuals/partitions from the population using a λ -tournament selection rule, i. e., choose λ random individuals from the population and select the best among those if it has not been selected previously. Let the selected partitions be $P_1 = (V_1, \dots, V_k)$ and $P_2 = (W_1, \dots, W_k)$. Then sets of nodes that are grouped together, i. e.,

$$\mathcal{U} := \{\{V_1 \cap W_{\sigma(1)}\}, \dots, \{V_k \cap W_{\sigma(k)}\}\}$$

are computed. This is done such that the number of nodes that are grouped together, i. e., $\sum_{j=1}^k |V_j \cap W_{\sigma(j)}|$, is maximum among all permutations σ of $\{1, \dots, k\}$. An offspring is created as follows. Sets of nodes in \mathcal{U} will be grouped within a block of the offspring. That means if a node is in one of the sets of \mathcal{U} , then it is assigned to the same block to which it was assigned to in P_1 . Otherwise, it is assigned to a random block, such that the balance constraint remains fulfilled. Local search is then used to improve the computed offspring before it is inserted into the population. Benlic and Hao [BH11a] combine their approach with tabu search. Their algorithms produce partitions of very high quality, but cannot guarantee that the output partition fulfills the desired balance constraint.

Sanders and Schulz introduced a distributed evolutionary algorithm, KaFFPaE (KaFF-PaEvolutionary) [SS12]. They present a general combine operator framework, which means that a partition \mathcal{P} can be combined with another partition or an arbitrary clustering of the graph, as well as multiple mutation operators to ensure diversity in the population. The combine operation uses a modified version of the multilevel GP solver within KaHIP [SS11] that will not contract edges that are cut in one of the input partitions/clustering. In contrast to the other approaches, the combine operation can ensure that the resulting offspring/partition is at least as good as the input partition \mathcal{P} . The algorithm is equipped with a scalable communication protocol similar to randomized rumor spreading and has been able to improve the best known partitions for many inputs.

8 Parallel Aspects of Graph Partitioning

In the era of stalling CPU clock speeds, exploiting parallelism is probably the most important way to accelerate computer programs from a hardware perspective. When executing parallel graph algorithms without shared memory, a good distribution of the graph onto the PEs is very important. Since parallel computing is a major purpose for GP, we discuss in this section several techniques beneficial for parallel scenarios. (i) Parallel GP algorithms are often necessary due to memory constraints: Partitioning a huge distributed graph on a single PE is often infeasible. (ii) When different PEs communicate with different speeds with each other, techniques for mapping the blocks communication-efficiently onto the PEs become important. (iii) When the graph changes over time (as in dynamic simulations), so does its partition. Once the imbalance becomes too large, one should find a new partition that unifies three criteria for this purpose: balance, low communication, and low migration.

8.1 Parallel Algorithms

Parallel GP algorithms are becoming more and more important since parallel hardware is now ubiquitous and networks grow. If the underlying application is in parallel processing, finding the partitions in parallel is even more compelling. The difficulty of parallelization very much depends on the circumstances. It is relatively easy to run sequential GP solvers multiple times with randomized tie breaking in all available decisions. Completely independent runs quickly lead to a point of diminishing return but are a useful strategy for very simple initial partitioners as the one described in Section 4.3. Evolutionary GP solvers are more effective (thanks to very good combination operators) and scale very well, even on loosely coupled distributed machines [SS12].

Most of the geometry-based algorithms from Section 4.5 are parallelizable and perhaps this is one of the main reasons for using them. In particular, one can use them to find an initial distribution of nodes to processors in order to improve the locality of a subsequent graph based parallel method [HSS10]. If such a “reasonable” distribution of a large graph over the local memories is available, distributed memory multilevel partitioners using MPI can be made to scale [KK96, WC00, CP08, HSS10]. However, loss of quality compared to the sequential algorithms is a constant concern. A recent parallel matching algorithm allows high quality coarsening, though [BOS⁺13]. If k coincides with the number of processors, one can use parallel edge coloring of the quotient graph to do pairwise refinement between neighboring blocks. At least for mesh-like graphs this scales fairly well [HSS10] and gives quality comparable to sequential solvers. This comparable solution quality also holds for parallel JOSTLE as described by Walshaw and Cross [WC02].

Parallelizing local search algorithms like KL/FM is much more difficult since local search is inherently sequential and since recent results indicate that it achieves best quality when performed in a highly localized way [OS10, SS11]. When restricting local search to improving moves, parallelization is possible, though [KK99b, LK13, MSS15, ASS15]. In a shared memory context, one can also use speculative parallelism [SNBP11]. The diffusion-based improvement methods described in Section 5.6 are also parallelizable without loss of quality since they are formulated in a naturally data parallel way [Mey12, Pel12].

8.2 Mapping Techniques

Fundamentals. Parallel computing on graphs is one major application area of GP, see Section 3.1. A partition with a small communication volume translates directly into an efficient application if the underlying hardware provides uniform communication speed between each pair of processing elements (PEs). Most of today’s leading parallel systems, however, are built as a hierarchy of PEs, memory systems, and network connections [MSSD13]. Communicating data between PEs close to each other is thus usually less expensive than between PEs with a high distance. On such architectures it is important to extend GPP by a flexible assignment of blocks to PEs [TBFS00].

Combining partitioning and mapping to PEs is often done in two different ways. In the first one, which we term *architecture-aware partitioning*, the cost of communicating data between a pair of PEs is directly incorporated into the objective function during the partitioning process. As an example, assuming that block (or process) i is run on PE i , the communication-aware edge cut function is $\sum_{i < j} \omega(E_{ij}) \cdot \omega_p(i, j)$, where $\omega_p(i, j)$ specifies the cost of communicating a unit item from PE i to PE j [WC01]. This approach uses a network cost matrix (NCM) to store the distance function ω_p [WC01, p. 603ff.]. Since the entries are queried frequently during partitioning, a recomputation of the matrix would be too costly. For large systems one must find a way around storing the full NCM on each PE, as the storage size scales quadratically with the number of PEs. A similar approach with emphasis on modeling heterogeneous communication costs in grid-based systems is undertaken by the software PaGrid [HAB06]. Moulitsas and Karypis [MK08] perform architecture-aware partitioning in two phases. Their so-called *predictor-corrector approach* concentrates in the first phase only on the resources of each PE and computes an according partition. In the second phase the method corrects previous decisions by modifying the partition according to the interconnection network characteristics, including heterogeneity.

An even stronger decoupling takes place for the second problem formulation, which we refer to as the *mapping problem*. Let $G_c = (V_c, E_c, \omega_c)$ be the *communication graph* that models the application’s communication, where $(u, v) \in E_c$ denotes how much data process u sends to process v . Let furthermore $G_p = (V_p, E_p, \omega_p)$ be the *processor graph*, where $(i, j) \in E_p$ specifies the bandwidth (or the latency) between PE i and PE j . We now assume that a partition has already been computed, inducing a communication graph G_c . The task after partitioning is then to find a communication-optimal *mapping* $\pi : V_c \mapsto V_p$.

Different objective functions have been proposed for this mapping problem. Since it is difficult to capture the deciding hardware characteristics, most authors concentrate on simplified cost functions – similar to the simplification of the edge cut for graph partitioning. Apparently small variations in the cost functions rarely lead to drastic variations in application running time. For details we refer to Pellegrini’s survey on static mapping [Pel11] (which we wish to update with this section, not to replace) and the references therein. Global sum type cost functions do not have the drawback of requiring global updates. Moreover, discontinuities in their search space, which may inhibit metaheuristics to be effective, are usually less pronounced than for maximum-based cost functions. Commonly used is the sum, for

all edges of G_c , of their weight multiplied by the cost of a unit-weight communication in G_p [Pel11]: $f(G_c, G_p, \pi) := \sum_{(u,v) \in E_c} \omega_c(u, v) \cdot \omega_p(\pi(u), \pi(v))$.

The accuracy of the distance function ω_p depends on several factors, one of them being the *routing algorithm*, which determines the paths a message takes. The maximum length over all these paths is called the *dilation* of the embedding π . One simplifying assumption can be that the routing algorithm is oblivious [HS11] and, for example, uses always shortest paths. When multiple messages are exchanged at the same time, the same communication link may be requested by multiple messages. This *congestion* of edges in G_p can therefore be another important factor to consider and whose maximum (or average) over all edges should be minimized. Minimizing the maximum congestion is NP-hard, cf. Garey and Johnson [GJ79] or more recent work [HS11, KL91].

Algorithms. Due to the problem’s complexity, exact mapping methods are only practical in special cases. Leighton’s book [Lei92] discusses embeddings between arrays, trees, and hypercubic topologies. One can apply a wide range of optimization techniques to the mapping problem, also multilevel algorithms. Their general structure is very similar to that described in Section 6. The precise differences of the single stages are beyond our scope. Instead we focus on very recent results – some of which also use hierarchical approaches. For pointers to additional methods we refer the reader to Pellegrini [Pel11] and Aubanel’s short summary [Aub09] on resource-aware load balancing.

Greedy approaches such as the one by Brandfass et al. [BAG13] map the node v_c of G_c with the highest total communication cost w.r.t. to the already mapped nodes onto the node v_p of G_p with the smallest total distance w.r.t. to the already mapped nodes. Some variations exist that improve this generic approach in certain settings [HS11, GMN15].

Hoefler and Snir [HS11] employ the reverse Cuthill-McKee (RCM) algorithm as a mapping heuristic. Originally, RCM has been conceived for the problem of minimizing the bandwidth of a sparse matrix [GL81]. In case both G_c and G_p are sparse, the simultaneous optimization of both graph layouts can lead to reasonable mapping results, also cf. Pellegrini [Pel07b].

Many metaheuristics have been used to solve the mapping problem. Uçar et al. [UAKI06] implement a large variety of methods within a clustering approach, among them genetic algorithms, simulated annealing, tabu search, and particle swarm optimization. Brandfass et al. [BAG13] present local search and evolutionary algorithms. Their experiments confirm that metaheuristics are significantly slower than problem-specific heuristics, but obtain high-quality solutions [BAG13, UAKI06].

Another common approach is to partition G_c – or the application graph itself – simultaneously together with G_p into the same number of blocks k' . This is for example done in SCOTCH [Pel94]. For this approach k' is chosen small enough so that it is easy to test which block in G_c is mapped onto which block in G_p . Since this often implies $k' < k$, the partitioning is repeated recursively. When the number of nodes in each block is small enough, the mapping within each block is computed by brute force. If $k' = 2$ and the two graphs to be partitioned are the application graph and G_p , the method is called *dual recursive bipartitioning*. Recently, schemes that model the processor graph as a tree have emerged [CLA12] in this algorithmic context and in similar ones [JMT13].

Hoefler and Snir [HS11] compare the greedy, RCM, and dual recursive (bi)partitioning mapping techniques experimentally. On a 3D torus and two other real architectures, their results do not show a clear winner. However, they confirm previous studies [Pel11] in that performing mapping at all is worthwhile. Bhatele et al. [BJGK11] discuss topology-aware mappings of different communication patterns to the physical topology in the context of MPI on emerging architectures. Better mappings avoid communication hot spots and reduce communication times significantly. Geometric information can also be helpful for finding good mappings on regular architectures such as tori [BK11].

8.3 Migration Minimization during Repartitioning

Repartitioning involves a tradeoff between the quality of the new partition and the migration volume. Larger changes between the old partition Π and the new one Π' , necessary to obtain a small communication volume in Π' , result in a higher migration volume. Different strategies have been explored in the literature to address this tradeoff. Two simple ones and their limitations are described by Schloegel et al. [SKK97]. One approach is to compute a new partition Π' from scratch and determine a migration-minimal mapping between Π and Π' . This approach delivers good partitions, but the migration volume is often very high. Another strategy simply migrates nodes from overloaded blocks to underloaded ones, until a new balanced partition is reached. While this leads to optimal migration costs, it often delivers poor partition quality. To improve these simple schemes, Schloegel et al. [SKK00] combine the two and get the best of both in their tool ParMetis.

Migration minimization with virtual nodes has been used in the repartitioning case by, among others, Hendrickson et al. [HLD96]. For each block, an additional node is added, which may not change its affiliation. It is connected to each node v of the block by an edge whose weight is proportional to the migration cost for v . Thus, one can account for migration costs and partition quality at the same time. A detailed discussion of this general technique was made by Walshaw [Wal10]. Recently, this technique has been extended to heterogeneous architectures by Fourestier and Pellegrini [FP11].

Diffusion-based partitioning algorithms are particularly strong for repartitioning. PDibaP yields about 30 – 50% edge cut improvement compared to ParMetis and about 15% improvement on parallel Jostle with a comparable migration volume [Mey12] (a short description of these tools can be found in Section 9.3). Hypergraph-based repartitioning is particularly important when the underlying problem has a rather irregular structure [CB⁺07].

9 Implementation and Evaluation Aspects

The two major factors that make up successful GP algorithms are speed and quality. It depends on the application if one of them is favored over the other and what quality means. Speed requires an appropriate implementation, for which we discuss the most common graph data structures in practice first in this section. Then, we discuss GP benchmarks to assess different algorithms and implementations, some widely used, others with potential. Finally, relevant software tools for GP are presented.

9.1 Sparse Graph Data Structures

The graph data structure used by most partitioning software is the Compressed Sparse Rows (CSR) format, also known as adjacency arrays. CSR is a cache and storage efficient data structure for representing static graphs. The CSR representation of a graph can be composed of two, three, or four arrays, depending upon whether edges or nodes are weighted. The node array (\mathbf{V}) is of size $n + 1$ and holds the node pointers. The edge array and the edge weights array, if present, are of size m each. Each entry in the edge array (\mathbf{E}) holds the node id of the target node, while the corresponding entry in the edge weights array (\mathbf{W}) holds the weight of the edge. The node array holds the offsets to the edge array, meaning that the target nodes of the outgoing edges of the i th node are accessible from $\mathbf{E}(\mathbf{V}(i))$ to $\mathbf{E}(\mathbf{V}(i + 1) - 1)$ and their respective weights are accessible from $\mathbf{W}(\mathbf{V}(i))$ to $\mathbf{W}(\mathbf{V}(i + 1) - 1)$. Both Metis and Scotch use a CSR-like data structure. Since nodes can also be weighted in graph partitioning, an additional vector of size n is often used to store node weights in that case. The CSR format can further be improved and reinforced by rearranging the nodes with one of the cache-oblivious layouts such as the minimum logarithmic arrangement [ST11, CKL⁺09].

Among distributed-memory GP solvers, ParMetis and PT-Scotch use a 1D node distribution where each processor owns approximately n/p nodes and their corresponding edges. By contrast, Zoltan uses a 2D edge distribution that has lower communication requirements in theory.

9.2 Benchmarking

The Walshaw benchmark⁵ was created in 2000 by Soper, Walshaw and Cross [SWC04]. This public domain archive, maintained by Chris Walshaw, contains 34 real-world graphs stemming from applications such as finite element computations, matrix computations, VLSI Design and shortest path computations. More importantly, it also contains for each graph the partitions with the smallest cuts found so far. Submissions are sought that achieve improved cut values for $k \in \{2, 4, 8, 16, 32, 64\}$ and balance parameters $\epsilon \in \{0, 0.01, 0.03, 0.05\}$, while running time is not an issue. Currently, solutions of over 40 algorithms have been submitted to the archive. It is the most popular GP benchmark in the literature.

There are many other very valuable sources of graphs for experimental evaluations: the 10th DIMACS Implementation Challenge [BMS⁺ar, BMSW13], the Florida Sparse Matrix Collection [Dav], the Laboratory of Web Algorithms [LoWA], the Koblenz Network Collection [Kun13], and the Stanford Large Network Dataset Collection [Les]. Many of the graphs are available at the website of the 10th DIMACS Implementation Challenge [BMS⁺ar, BMSW13] in the graph format that is used by many GP software tools.

Aubanel et al. [GSAG08] present a different kind of partitioning benchmark. Instead of measuring the edge cut of the partitions, the authors evaluate the execution time of a parallel PDE solver to benchmark the partitions produced by different GP solvers. The crucial module of the benchmark is parallel matrix-vector multiplication, which is meaningful for other numerical routines as well.

⁵ <http://staffweb.cms.gre.ac.uk/~wc06/partition/>

Many fast methods for GPP are based on approaches in which finding a global solution is done by local operations only. Testing if such methods are robust against falling into local optima obtained by the local processing is a very important task. In [SSS12] a simple strategy for checking the quality of such methods was presented. To construct a potentially hard instance, one may consider a mixture of graphs with very different structures that are weakly connected with each other. For example, in multilevel algorithms these graphs can force the algorithm to contract incorrect edges that lead to uneven coarsening; also, they can attract a “too strong” refinement to reach a local optimum, which can contradict better optima at finer levels. Examples of real graphs that contain such mixtures of structures include multi-mode networks [TLZN08] and logistics multi-stage system networks [Sto06]. Hardness of particular structures for GP solvers is confirmed by generating graphs that are similar to the given ones at both coarse and/or fine resolutions [GMS12].

9.3 Software Tools

There are a number of software packages that implement the described algorithms. One of the first publicly available software packages called Chaco is due to Hendrickson and Leland [Hen]. As most of the publicly available software packages, Chaco implements the multilevel approach outlined in Section 6 and basic local search algorithms. Moreover, they implement spectral partitioning techniques. Probably the fastest and best known system is the Metis family by Karypis and Kumar [KK98a, KK98b]. kMetis [KK98b] is focused on partitioning speed and hMetis [KK99a], which is a hypergraph partitioner, aims at partition quality. PaToH [cA11] is also a widely used hypergraph partitioner that produces high quality partitions. ParMetis is a widely used parallel implementation of the Metis GP algorithm [KK96]. Scotch [Pel, CP06, CP08] is a GP framework by Pellegrini. It uses recursive multilevel bisection and includes sequential as well as parallel partitioning techniques. Jostle [WC07, WC00] is a well-known sequential and parallel GP solver developed by Chris Walshaw. The commercialised version of this partitioner is known as NetWorks. It has been able to hold most of the records in the Walshaw Benchmark for a long period of time. If a model of the communication network is available, then Jostle and Scotch are able to take this model into account for the partitioning process. Party [DPSW00b, MS04] implements the Bubble/shape-optimized framework and the Helpful Sets algorithm. The software packages DibaP and its MPI-parallel variant PDibaP by Meyerhenke [MMS09a, Mey12] implement the Bubble framework using diffusion; DibaP also uses AMG-based techniques for coarsening and solving linear systems arising in the diffusive approach. Recently, Sanders and Schulz [SS, SS13b] released the GP package KaHIP (Karlsruhe High Quality Partitioning) which implements for example flow-based methods, more-localized local searches and several parallel and sequential meta-heuristics. KaHIP scored most of the points in the GP subchallenge of the 10th DIMACS Implementation Challenge [BMSW13] and currently holds most of the entries in the Walshaw Benchmark.

To address the load balancing problem in parallel applications, distributed versions of the established sequential partitioners Metis, Jostle and Scotch [SKK02, WC07, Pel12] have been developed. The tools Parkway by Trifunovic and Knottenbelt [TK08] as well as Zoltan by

Devine et al. [DBH⁺06] focus on hypergraph partitioning. Recent results of the 10th DIMACS Implementation Challenge [BMSW13] suggest that scaling current hypergraph partitioners to very large systems is even more challenging than graph partitioners.

10 Future Challenges

It is an interesting question to what extent the multitude of results sketched above have reached a state of maturity where future improvements become less and less likely. On the one hand, if you consider the Walshaw benchmark with its moderately sized static graphs with mostly regular structure, the quality obtained using the best current systems is very good and unlikely to improve much in the future. One can already get very good quality with a careful application of decade old techniques like KL/FM local search and the multilevel approach. On the other hand, as soon as you widen your view in some direction, there are plenty of important open problems.

Bridging Gaps Between Theory and Practice. We are far from understanding why (or when) the heuristic methods used in practice produce solutions very close to optimal. This is particularly striking for bipartitioning, where recent exact results suggest that heuristics often find the optimal solution. In contrast, theoretical results state that we cannot even find constant-factor approximations in polynomial time. On the other hand, the sophisticated theoretical methods developed to obtain approximation guarantees are currently not used in the most successful solvers. It would be interesting to see to what extent these techniques can yield a practical contribution. There is a similar problem for exact solvers, which have made rapid progress for the case $k = 2$. However, it remains unclear how to use them productively for larger graphs or in case $k > 2$, for example as initial partitioners in a multilevel system or for pair-wise local improvement of subgraphs. What *is* surprisingly successful, is the use of solvers with performance guarantees for subproblems that are easier than partitioning. For example, KaHIP [SS] uses weighted matching, spanning trees, edge coloring, BFS, shortest paths, diffusion, maximum flows, and strongly connected components. Further research into this direction looks promising.

Difficult Instances. The new “complex network” applications described in Section 3.2 result in graphs that are not only very large but also difficult to handle for current graph partitioners. This difficulty results from an uneven degree distribution and much less locality than observed in traditional inputs. Here, improved techniques within known frameworks (e.g., better coarsening schemes) and even entirely different approaches can give substantial improvements in speed or quality.

Another area where large significant quality improvements are possible are for large k . Already for the largest value of k considered in the Walshaw benchmark (64), the spread between different approaches is considerable. Considering graphs with billions of nodes and parallel machines reaching millions of processors, $k \leq 64$ increasingly appears like a special case. The multilevel method loses some of its attractiveness for large k since even initial partitioning must solve quite large instances. Hence new ideas are required.

Multilevel Approach. While the multilevel paradigm has been extremely successful for GP, there are still many algorithmic challenges ahead. The variety of continuous systems mul-

tilelevel algorithms (such as various types of multigrid) turned into a separate field of applied mathematics, and optimization. Yet, multilevel algorithms for GPP still consist in practice of a very limited number of multilevel techniques. The situation with other combinatorial optimization problems is not significantly different. One very promising direction is bridging the gaps between the theory and practice of multiscale computing and multilevel GP such as introducing nonlinear coarsening schemes. For example, a novel multilevel approach for the minimum vertex separator problem was recently proposed using the continuous bilinear quadratic program formulation [HHS14], and a *hybrid of the geometric multigrid, and full approximation scheme* for continuous problem was used for graph drawing, and VLSI placement problems [RSB10, CS03]. Development of more sophisticated coarsening schemes, edge ratings, and metrics of nodes' similarity that can be propagated throughout the hierarchies are among the future challenges for graph partitioning as well as any attempt of their rigorous analysis.

Parallelism and Other Hardware Issues. Scalable high quality GP (with quality comparable to sequential partitioners) remains an open problem. With the advent of exascale machines with millions of processors and possibly billions of threads, the situation is further aggravated. Traditional “flat” partitions of graphs for processing on such machines implies a huge number of blocks. It is unclear how even sequential partitioners perform for such instances. Resorting to recursive partitioning brings down k and also addresses the hierarchical nature of such machines. However, this means that we need parallel partitioners where the number of available processors is much bigger than k . It is unclear how to do this with high quality. Approaches like the band graphs from PT-Scotch are interesting but likely to fail for complex networks.

Efficient implementation is also a big issue since complex memory hierarchies and heterogeneity (e.g., GPUs or FPGAs) make the implementation complicated. In particular, there is a mismatch between the fine-grained discrete computations predominant in the best sequential graph partitioners and the massive data parallelism (SIMD-instructions, GPUs, . . .) in high performance computing which better fits highly regular numeric computations. It is therefore likely that high quality GP will only be used for the higher levels of the machine hierarchy, e.g., down to cluster nodes or CPU sockets. At lower levels of the architectural hierarchy, we may use geometric partitioning or even regular grids with dummy values for non-existing cells (e.g. [FKS⁺12]).

While exascale computing is a challenge for high-end applications, many more applications can profit from GP in cloud computing and using tools for high productivity such as Map/Reduce [DG04], Pregel [MAB⁺10], GraphLab [LGK⁺12], Combinatorial BLAS [BG11], or Parallel Boost Graph Library [GL05]. Currently, none of these systems uses sophisticated GP software.

These changes in architecture also imply that we are no longer interested in algorithms with little computations but rather in data access with high locality and good energy efficiency.

Beyond Balanced k -partitioning with Cut Minimization. We have intentionally fixed our basic model assumptions above to demonstrate that even the classical setting has a lot of

open problems. However, these assumption become less and less warranted in the context of modern massively parallel hardware and huge graphs with complex structure. For example, it looks like the assumptions that low total cut is highly correlated with low bottleneck cut or communication volume (see Section 2.1) is less warranted for complex network [BM12]. Eventually, we would like a dynamic partition that adapts to the communication requirements of a computation such as PageRank or BFS with changing sets of active nodes and edges. Also, the fixed value for k becomes questionable when we want to tolerate processor failures or achieve “malleable” computations that adapt their resource usage to the overall situation, e.g., to the arrival or departure of high priority jobs. Techniques like overpartitioning, repartitioning (with changed k), and (re)mapping will therefore become more important. Even running time as the bottom-line performance goal might be replaced by energy consumption [SDM11].

Acknowledgements

We express our gratitude to Bruce Hendrickson, Dominique LaSalle, and George Karypis for many valuable comments on a preliminary draft of the manuscript.

References

- [AB13] B. F. Auer and R. H. Bisseling. Graph Coarsening and Clustering on the GPU. In Bader et al. [BMSW13], pages 19–36.
- [ACFK07] C. Aykanat, B. B. Cambazoglu, F. Findik, and T. Kurc. Adaptive Decomposition and Remapping Algorithms for Object-space-parallel Direct Volume Rendering of Unstructured Grids. *J. Parallel Distrib. Comput.*, 67(1):77–99, January 2007.
- [AFHM08] M. Armbruster, M. Fügenschuh, C. Helmberg, and A. Martin. A Comparative Study of Linear and Semidefinite Branch-and-Cut Methods for Solving the Minimum Graph Bisection Problem. In *13th International Conference on Integer Programming and Combinatorial Optimization (IPCO)*, volume 5035 of *LNCS*, pages 112–124. Springer, 2008.
- [AHK10] S. Arora, E. Hazan, and S. Kale. $O(\sqrt{\log n})$ Approximation to Sparsest Cut in $\tilde{O}(n^2)$ Time. *SIAM Journal on Computing*, 39(5):1748–1771, 2010.
- [AL08] R. Andersen and K. J. Lang. An Algorithm for Improving Graph Partitions. In *19th ACM-SIAM Symposium on Discrete Algorithms*, pages 651–660, 2008.
- [AR06] K. Andreev and H. Räcke. Balanced Graph Partitioning. *Theory of Computing Systems*, 39(6):929–939, 2006.
- [ARK06] A. Abou-Rjeili and G. Karypis. Multilevel Algorithms for Partitioning Power-Law Graphs. In *20th International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 2006.
- [Arm07] M. Armbruster. *Branch-and-Cut for a Semidefinite Relaxation of Large-Scale Minimum Bisection Problems*. PhD thesis, U. Chemnitz, 2007.
- [ARV04] S. Arora, S. Rao, and U. Vazirani. Expander Flows, Geometric Embeddings and Graph Partitioning. In *36th ACM Symposium on the Theory of Computing (STOC)*, pages 222–231, 2004.
- [ASS15] Y. Akhremtsev, P. Sanders, and C. Schulz. (Semi-)External Algorithms for Graph Partitioning and Clustering. In *15th Workshop on Algorithm Engineering and Experimentation (ALENEX)*, pages 33–43, 2015.
- [Aub09] E. Aubanel. Resource-Aware Load Balancing of Parallel Applications. In E. Udoh and F. Z. Wang, editors, *Handbook of Research on Grid Technologies and Utility Computing: Concepts for Managing Large-Scale Applications*, pages 12–21. Information Science Reference - Imprint of: IGI Publishing, May 2009.
- [Bad13] M. Bader. *Space-Filling Curves*. Springer, 2013.

- [BAG13] B. Brandfass, T. Alrutz, and T. Gerhold. Rank Reordering for MPI Communication Optimization. *Computers & Fluids*, 80(0):372–380, 2013.
- [BCLS87] T. Bui, S. Chaudhuri, F. Leighton, and M. Sipser. Graph Bisection Algorithms with Good Average Case Behavior. *Combinatorica*, 7:171–191, 1987.
- [BCR97] L. Brunetta, M. Conforti, and G. Rinaldi. A Branch-and-Cut Algorithm for the Equicut Problem. *Mathematical Programming*, 78(2):243–263, 1997.
- [BDR13] E. G. Boman, K. D. Devine, and S. Rajamanickam. Scalable Matrix Computations on Large Scale-Free Graphs Using 2D Graph Partitioning. In *ACM/IEEE Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, 2013.
- [BG11] A. Buluç and J. R. Gilbert. The Combinatorial BLAS: Design, Implementation, and Applications. *International Journal of High Performance Computing Applications*, 25(4):496–509, 2011.
- [BH10] U. Benlic and J. K. Hao. An Effective Multilevel Memetic Algorithm for Balanced Graph Partitioning. In *22nd IEEE Int. Conference on Tools with Artificial Intelligence (ICTAI)*, pages 121–128, 2010.
- [BH11a] U. Benlic and J. K. Hao. A Multilevel Memetic Approach for Improving Graph k -Partitions. *IEEE Transactions on Evolutionary Computation*, 15(5):624–642, 2011.
- [BH11b] U. Benlic and J. K. Hao. An Effective Multilevel Tabu Search Approach for Balanced Graph Partitioning. *Computers & Operations Research*, 38(7):1066–1075, 2011.
- [Bic07] C. E. Bichot. A New Method, the Fusion Fission, for the Relaxed k -Way Graph Partitioning Problem, and Comparisons with some Multilevel Algorithms. *Journal of Mathematical Modelling and Algorithms*, 6(3):319–344, 2007.
- [BJGK11] A. Bhatele, N. Jain, W. D. Gropp, and L. V. Kale. Avoiding Hot-Spots on Two-Level Direct Networks. In *ACM/IEEE Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, pages 76:1–76:11. ACM, 2011.
- [BK11] A. Bhatele and L. Kale. Heuristic-Based Techniques for Mapping Irregular Communication Graphs to Mesh Topologies. In *13th Conference on High Performance Computing and Communications (HPCC)*, pages 765–771, 2011.
- [BM12] A. Buluç and K. Madduri. Graph Partitioning for Scalable Distributed Graph Computations. In Bader et al. [BMSW13], pages 83–102.
- [BMS⁺ar] D. A. Bader, H. Meyerhenke, P. Sanders, C. Schulz, A. Kappes, and D. Wagner. Benchmarking for Graph Clustering and Graph Partitioning. In *Encyclopedia of Social Network Analysis and Mining*, to appear.
- [BMSW13] D. A. Bader, H. Meyerhenke, P. Sanders, and D. Wagner, editors. *Graph Partitioning and Graph Clustering – 10th DIMACS Impl. Challenge*, volume 588 of *Contemporary Mathematics*. AMS, 2013.
- [Bop87] R. B. Boppana. Eigenvalues and Graph Bisection: An Average-Case Analysis). In *28th Symposium on Foundations of Computer Science (FOCS)*, pages 280–285, 1987.
- [BOS⁺13] M. Birn, V. Osipov, P. Sanders, C. Schulz, and N. Sitchinava. Efficient Parallel and External Matching. In *Euro-Par*, volume 8097 of *LNCIS*, pages 659–670. Springer, 2013.
- [BS93] S. T. Barnard and H. D. Simon. A Fast Multilevel Implementation of Recursive Spectral Bisection for Partitioning Unstructured Problems. In *6th SIAM Conference on Parallel Processing for Scientific Computing*, pages 711–718, 1993.
- [BS11] C. Bichot and P. Siarry, editors. *Graph Partitioning*. Wiley, 2011.
- [CA01] U. Catalyurek and C. Aykanat. A Hypergraph-Partitioning Approach for Coarse-Grain Decomposition. In *ACM/IEEE Conference on Supercomputing (SC)*. ACM, 2001.
- [cA11] Ü. Çatalyürek and C. Aykanat. *PaToH: Partitioning Tool for Hypergraphs*, 2011.
- [CB⁺07] U. Catalyurek, E. Boman, et al. Hypergraph-based Dynamic Load Balancing for Adaptive Scientific Computations. In *21st Int. Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 2007.
- [CBM07] P. Chardaire, M. Barake, and G. P. McKeown. A PROBE-Based Heuristic for Graph Partitioning. *IEEE Transactions on Computers*, 56(12):1707–1720, 2007.
- [CC11] S. Chu and J. Cheng. Triangle Listing in Massive Networks and its Applications. In *17th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 672–680, 2011.
- [CG12] K. S. Camilus and V. K. Govindan. A Review on Graph Based Segmentation. *IJIGSP*, 4:1–13, 2012.
- [CKL⁺09] F. Chierichetti, R. Kumar, S. Lattanzi, M. Mitzenmacher, A. Panconesi, and P. Raghavan. On Compressing Social Networks. In *15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 219–228, 2009.
- [CLA12] S. Y. Chan, T. C. Ling, and E. Aubanel. The Impact of Heterogeneous Multi-Core Clusters on Graph Partitioning: An Empirical Study. *Cluster Computing*, 15(3):281–302, 2012.

- [CP06] C. Chevalier and F. Pellegrini. Improvement of the Efficiency of Genetic Algorithms for Scalable Parallel Graph Partitioning in a Multi-level Framework. In *12th International Conference on Parallel Processing*, volume 4128 of *LNCS*, pages 243–252. Springer, 2006.
- [CP08] C. Chevalier and F. Pellegrini. PT-Scotch: A Tool for Efficient Parallel Graph Ordering. *Parallel Computing*, 34(6):318–331, 2008.
- [CS03] J. Cong and J. Shinnerl. *Multilevel Optimization in VLSICAD*. Springer, 2003.
- [CS06] F. Comellas and E. Sapena. A Multiagent Algorithm for Graph Partitioning. In *Applications of Evolutionary Computing (EvoWorkshops)*, volume 3907 of *LNCS*, pages 279–285. Springer, 2006.
- [CS09] C. Chevalier and I. Safro. Comparison of Coarsening Schemes for Multi-Level Graph Partitioning. In *Proceedings Learning and Intelligent Optimization*, 2009.
- [CS11] J. Chen and I. Safro. Algebraic Distance on Graphs. *SIAM J. Scientific Computing*, 33(6):3468–3490, 2011.
- [Dav] T. Davis. The University of Florida Sparse Matrix Collection, <http://www.cise.ufl.edu/research/sparse/matrices>, 2008.
- [DBH⁺06] K. D. Devine, E. G. Boman, R. T. Heaphy, R. H. Bisseling, and U. V. Catalyurek. Parallel hypergraph partitioning for scientific computing. In *Proceedings of the IEEE International Parallel and Distributed Processing Symposium, IPDPS'06*, pages 124–124, Washington, DC, USA, 2006. IEEE Computer Society.
- [DG04] J. Dean and S. Ghemawat. MapReduce: Simplified Data Processing on Large Clusters. In *6th Symposium on Operating System Design and Implementation (OSDI)*, pages 137–150. USENIX, 2004.
- [DG⁺11] D. Delling, A. V. Goldberg, et al. Graph Partitioning with Natural Cuts. In *25th International Parallel and Distributed Processing Symposium (IPDPS)*, pages 1135–1146, 2011.
- [DG12] H. J. Diasheng Guo, Ke Liao. Power System Reconfiguration based on Multi-Level Graph Partitioning. In *7th International Conference, GIScience 2012*, 2012.
- [DGPW11] D. Delling, A. V. Goldberg, T. Pajor, and R. F. Werneck. Customizable Route Planning. In *10th Symp. on Experimental Algorithms (SEA)*, volume 6630 of *LNCS*, pages 376–387. Springer, 2011.
- [DGRW12] D. Delling, A. V. Goldberg, I. Razenshteyn, and R. F. Werneck. Exact Combinatorial Branch-and-Bound for Graph Bisection. In *12th Workshop on Algorithm Engineering and Experimentation (ALENEX)*, pages 30–44, 2012.
- [DH72] W. E. Donath and A. J. Hoffman. Algorithms for Partitioning of Graphs and Computer Logic Based on Eigenvectors of Connection Matrices. *IBM Technical Disclosure Bulletin*, 15(3):938–944, 1972.
- [DH73] W. E. Donath and A. J. Hoffman. Lower Bounds for the Partitioning of Graphs. *IBM Journal of Research and Development*, 17(5):420–425, 1973.
- [DH03] D. Drake and S. Hougardy. A Simple Approximation Algorithm for the Weighted Matching Problem. *Information Processing Letters*, 85:211–213, 2003.
- [DH05] D. E. Drake Vinkemeier and S. Hougardy. A Linear-Time Approximation Algorithm for Weighted Matchings in Graphs. *ACM Transactions Algorithms*, 1(1):107–122, 2005.
- [DLL⁺05] V. Donde, V. Lopez, B. Lesieutre, A. Pinar, C. Yang, and J. Meza. Identification of Severe Multiple Contingencies in Electric Power Networks. In *37th N. A. Power Symposium*, pages 59–66. IEEE, 2005.
- [DMP95] R. Diekmann, B. Monien, and R. Preis. Using Helpful Sets to Improve Graph Bisections. *Interconnection Networks and Mapping and Scheduling Parallel Computations*, 21:57–73, 1995.
- [DPS11] R. Duan, S. Pettie, and H.-H. Su. Scaling Algorithms for Approximate and Exact Maximum Weight Matching. *CoRR*, abs/1112.0790, 2011.
- [DPSW00a] 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.
- [DPSW00b] R. Diekmann, R. Preis, F. Schlimbach, and C. Walshaw. Shape-optimized Mesh Partitioning and Load Balancing for Parallel Adaptive FEM. *Parallel Computing*, 26(12):1555–1581, 2000.
- [Dut93] S. Dutt. New Faster Kernighan-Lin-type Graph-Partitioning Algorithms. In *4th IEEE/ACM Conference on Computer-Aided Design*, pages 370–377, 1993.
- [DW12] D. Delling and R. F. Werneck. Better Bounds for Graph Bisection. In *20th European Symposium on Algorithms*, volume 7501 of *LNCS*, pages 407–418, 2012.
- [DW13] D. Delling and R. F. Werneck. Faster Customization of Road Networks. In *12th Symposium on Experimental Algorithms*, volume 7933 of *LNCS*, pages 30–42. Springer, 2013.
- [ENRS99] G. Even, J. S. Naor, S. Rao, and B. Schieber. Fast Approximate Graph Partitioning Algorithms. *SIAM Journal on Computing*, 28(6):2187–2214, 1999.
- [FB13] B. O. Fagginger Auer and R. H. Bisseling. Abusing a Hypergraph Partitioner for Unweighted Graph Partitioning. In Bader et al. [BMSW13], pages 19–35.

- [Fel05] A. Felner. Finding Optimal Solutions to the Graph Partitioning Problem with Heuristic Search. *Annals of Mathematics and Artificial Intelligence*, 45:293–322, 2005.
- [FF56] L. R. Ford and D. R. Fulkerson. Maximal Flow through a Network. *Canadian Journal of Mathematics*, 8(3):399–404, 1956.
- [Fie75] M. Fiedler. A Property of Eigenvectors of Nonnegative Symmetric Matrices and its Application to Graph Theory. *Czechoslovak Mathematical Journal*, 25(4):619–633, 1975.
- [FK02] U. Feige and R. Krauthgamer. A Polylogarithmic Approximation of the Minimum Bisection. *SIAM Journal on Computing*, 31(4):1090–1118, 2002.
- [FKS⁺12] J. Fietz, M. Krause, C. Schulz, P. Sanders, and V. Heuveline. Optimized Hybrid Parallel Lattice Boltzmann Fluid Flow Simulations on Complex Geometries. In *Euro-Par 2012 Parallel Processing*, volume 7484 of *LNCS*, pages 818–829. Springer, 2012.
- [FL93] C. Farhat and M. Lesoinne. Automatic Partitioning of Unstructured Meshes for the Parallel Solution of Problems in Computational Mechanics. *Journal for Numerical Methods in Engineering*, 36(5):745–764, 1993.
- [FM82] C. M. Fiduccia and R. M. Mattheyses. A Linear-Time Heuristic for Improving Network Partitions. In *19th Conference on Design Automation*, pages 175–181, 1982.
- [FMDS⁺98] C. E. Ferreira, A. Martin, C. C. De Souza, R. Weismantel, and L. A. Wolsey. The Node Capacitated Graph Partitioning Problem: A Computational Study. *Mathematical Programming*, 81(2):229–256, 1998.
- [For09] S. Fortunato. Community Detection in Graphs. *CoRR*, abs/0906.0612, 2009.
- [FP11] S. Fourestier and F. Pellegrini. Adaptation au repartitionnement de graphes d’une méthode d’optimisation globale par diffusion. In *RenPar’20*, 2011.
- [FW11] A. Feldmann and P. Widmayer. An $O(n^4)$ Time Algorithm to Compute the Bisection Width of Solid Grid Graphs. In *19th European Symposium on Algorithms*, volume 6942 of *LNCS*, pages 143–154. Springer, 2011.
- [GBF11] P. Galinier, Z. Boujbel, and M. C. Fernandes. An Efficient Memetic Algorithm for the Graph Partitioning Problem. *Annals of Operations Research*, 191(1):1–22, 2011.
- [GH94] O. Goldschmidt and D. S. Hochbaum. A Polynomial Algorithm for the k -Cut Problem for Fixed k . *Mathematics of Operations Research*, 19(1):24–37, 1994.
- [GJ79] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., 1979.
- [GJS74] M. R. Garey, D. S. Johnson, and L. Stockmeyer. Some Simplified NP-Complete Problems. In *6th ACM Symposium on Theory of Computing*, STOC, pages 47–63. ACM, 1974.
- [GL81] A. George and J. W. H. Liu. *Computer Solution of Large Sparse Positive Definite Systems*. Prentice-Hall, 1981.
- [GL05] D. Gregor and A. Lumsdaine. The Parallel BGL: A Generic Library for Distributed Graph Computations. *Parallel Object-Oriented Scientific Computing (POOSC)*, 2005.
- [Glo89] F. Glover. Tabu Search — Part I. *ORSA Journal on Computing*, 1(3):190–206, 1989.
- [Glo90] F. Glover. Tabu Search — Part II. *ORSA Journal on Computing*, 2(1):4–32, 1990.
- [GMN15] R. Glantz, H. Meyerhenke, and A. Noe. Algorithms for mapping parallel processes onto grid and torus architectures. In *Proc. 23rd Euromicro Intl. Conf. on Parallel, Distributed and Network-based Processing*, 2015. To appear. Preliminary version: <http://arxiv.org/abs/1411.0921>.
- [GMS12] A. Gutfraind, L. A. Meyers, and I. Safro. Multiscale Network Generation. *CoRR*, abs/1207.4266, 2012.
- [GMS14] R. Glantz, H. Meyerhenke, and C. Schulz. Tree-based coarsening and partitioning of complex networks. In J. Gudmundsson and J. Katajainen, editors, *Proc. 13th Intl. Symp. on Experimental Algorithms (SEA 2014)*, volume 8504 of *Lecture Notes in Computer Science*, pages 364–375. Springer International Publishing, 2014.
- [GMT98] J. R. Gilbert, G. L. Miller, and S. H. Teng. Geometric Mesh Partitioning: Implementation and Experiments. *SIAM Journal on Scientific Computing*, 19(6):2091–2110, 1998.
- [GS06] L. Grady and E. L. Schwartz. Isoperimetric Graph Partitioning for Image Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28:469–475, 2006.
- [GSAG08] K. Ghazinour, R. E. Shaw, E. E. Aubanel, and L. E. Garey. A Linear Solver for Benchmarking Partitioners. In *22nd IEEE International Symposium on Parallel and Distributed Processing (IPDPS)*, pages 1–8, 2008.
- [HAB06] S. Huang, E. Aubanel, and V. C. Bhavsar. PaGrid: A Mesh Partitioner for Computational Grids. *Journal of Grid Computing*, 4(1):71–88, 2006.
- [Hen] B. Hendrickson. Chaco: Software for Partitioning Graphs. <http://www.cs.sandia.gov/~bahendr/chaco.html>.

- [Hen98] B. Hendrickson. Graph Partitioning and Parallel Solvers: Has the Emperor No Clothes? In *5th Symposium on Solving Irregularly Structured Problems in Parallel*, volume 1457 of *LNCS*, pages 218–225. Springer, 1998.
- [HHS14] W. W. Hager, J. T. Hungerford, and I. Safro. A multilevel bilinear programming algorithm for the vertex separator problem. <http://arxiv-web3.library.cornell.edu/abs/1410.4885>, (arXiv:1410.4885), 2014.
- [HK99] W. W. Hager and Y. Krylyuk. Graph Partitioning and Continuous Quadratic Programming. *SIAM Journal on Discrete Mathematics*, 12(4):500–523, 1999.
- [HK00] B. Hendrickson and T. G. Kolda. Graph Partitioning Models for Parallel Computing. *Parallel Computing*, 26(12):1519–1534, 2000.
- [HL95a] B. Hendrickson and R. Leland. A Multilevel Algorithm for Partitioning Graphs. In *ACM/IEEE Conference on Supercomputing'95*, 1995.
- [HL95b] B. Hendrickson and R. Leland. An Improved Spectral Graph Partitioning Algorithm for Mapping Parallel Computations. *SIAM Journal on Scientific Computing*, 16(2):452–469, 1995.
- [HLD96] B. Hendrickson, R. Leland, and R. V. Driessche. Enhancing Data Locality by Using Terminal Propagation. In *29th Hawaii International Conference on System Sciences (HICSS'9) Volume 1: Software Technology and Architecture*, page 565, 1996.
- [HM91] J. Hromkovič and B. Monien. The Bisection Problem for Graphs of Degree 4 (Configuring Transputer Systems). In *16th Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 520 of *LNCS*, pages 211–220, 1991.
- [HPZ13] W. W. Hager, D. T. Phan, and H. Zhang. An Exact Algorithm for Graph Partitioning. *Mathematical Programming*, 137(1-2):531–556, 2013.
- [HR73] L. Hyafil and R. Rivest. Graph Partitioning and Constructing Optimal Decision Trees are Polynomial Complete Problems. Technical Report 33, IRIA – Laboratoire de Recherche en Informatique et Automatique, 1973.
- [HS11] T. Hoeffler and M. Snir. Generic Topology Mapping Strategies for Large-scale Parallel Architectures. In *ACM International Conference on Supercomputing (ICS'11)*, pages 75–85. ACM, 2011.
- [HSS10] M. Holtgrewe, P. Sanders, and C. Schulz. Engineering a Scalable High Quality Graph Partitioner. *24th IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 1–12, 2010.
- [HW02] J. Hungershofer and J.-M. Wierum. On the Quality of Partitions based on Space-Filling Curves. In *International Conference on Computational Science—ICCS*, pages 36–45. Springer, 2002.
- [JMT13] E. Jeannot, G. Mercier, and F. Tessier. Process Placement in Multicore Clusters: Algorithmic Issues and Practical Techniques. *IEEE Transactions on Parallel and Distributed Systems*, PP(99):1–1, 2013.
- [JS98] M. Jerrum and G. B. Sorkin. The Metropolis Algorithm for Graph Bisection. *Discrete Applied Mathematics*, 82(1-3):155–175, 1998.
- [JS08] B. Junker and F. Schreiber. *Analysis of Biological Networks*. Wiley, 2008.
- [KHKM11] J. Kim, I. Hwang, Y.-H. Kim, and B.-R. Moon. Genetic Approaches for Graph Partitioning: A Survey. In *13th Genetic and Evolutionary Computation (GECCO)*, pages 473–480. ACM, 2011.
- [KK96] G. Karypis and V. Kumar. Parallel Multilevel k -way Partitioning Scheme for Irregular Graphs. In *ACM/IEEE Supercomputing'96*, 1996.
- [KK98a] G. Karypis and V. Kumar. A Fast and High Quality Multilevel Scheme for Partitioning Irregular Graphs. *SIAM Journal on Scientific Computing*, 20(1):359–392, 1998.
- [KK98b] G. Karypis and V. Kumar. Multilevel k -way Partitioning Scheme for Irregular Graphs. *Journal on Parallel and Distributed Computing*, 48(1):96–129, 1998.
- [KK99a] G. Karypis and V. Kumar. Multilevel k -Way Hypergraph Partitioning. In *36th ACM/IEEE Design Automation Conference*, pages 343–348. ACM, 1999.
- [KK99b] G. Karypis and V. Kumar. Parallel Multilevel series k -Way Partitioning Scheme for Irregular Graphs. *Siam Review*, 41(2):278–300, 1999.
- [KL70] B. W. Kernighan and S. Lin. An Efficient Heuristic Procedure for Partitioning Graphs. *The Bell System Technical Journal*, 49(1):291–307, 1970.
- [KL91] Y. M. Kim and T.-H. Lai. The Complexity of Congestion-1 Embedding in a Hypercube. *Journal of Algorithms*, 12(2):246 – 280, 1991.
- [KLMH11] A. B. Kahng, J. Lienig, I. L. Markov, and J. Hu. *VLSI Physical Design - From Graph Partitioning to Timing Closure*. Springer, 2011.
- [KLSV10] T. Kieritz, D. Luxen, P. Sanders, and C. Vetter. Distributed Time-Dependent Contraction Hierarchies. In *9th Symposium on Experimental Algorithms*, volume 6049 of *LNCS*, pages 83–93. Springer, 2010.
- [KR13] S. Kirmani and P. Raghavan. Scalable Parallel Graph Partitioning. In *High Performance Computing, Networking, Storage and Analysis, SC'13*. ACM, 2013.

- [KRC00] S. E. Karisch, F. Rendl, and J. Clausen. Solving Graph Bisection Problems with Semidefinite Programming. *INFORMS Journal on Computing*, 12(3):177–191, 2000.
- [KSR04] P. Korosec, J. Silc, and B. Robic. Solving the Mesh-Partitioning Problem with an Ant-Colony Algorithm. *Parallel Computing*, 30(5-6):785–801, 2004.
- [Kun13] J. Kunegis. KONECT – the koblenz network collection. In *Web Observatory Workshop*, pages 1343–1350, 2013.
- [Lan50] C. Lanczos. An Iteration Method for the Solution of the Eigenvalue Problem of Linear Differential and Integral Operators. *Journal of Research of the National Bureau of Standards*, 45(4):255–282, 1950.
- [Lau04] U. Lauther. An Extremely Fast, Exact Algorithm for Finding Shortest Paths in Static Networks with Geographical Background. In *Münster GI-Days*, 2004.
- [LD60] A. H. Land and A. G. Doig. An Automatic Method of Solving Discrete Programming Problems. *Econometrica*, 28(3):497–520, 1960.
- [Lei92] F. T. Leighton. *Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes*. Morgan Kaufmann Publishers, 1992.
- [Les] J. Lescovec. Stanford Network Analysis Package (SNAP). <http://snap.stanford.edu/index.html>.
- [LGK⁺12] Y. Low, J. Gonzalez, A. Kyrola, D. Bickson, C. Guestrin, and J. M. Hellerstein. Distributed GraphLab: A Framework for Machine Learning in the Cloud. *PVLDB*, 5(8):716–727, 2012.
- [LK13] D. Lasalle and G. Karypis. Multi-threaded Graph Partitioning. In *27th International Parallel and Distributed Processing Symposium (IPDPS)*, pages 225–236, 2013.
- [LL06] S. Lafon and A. B. Lee. Diffusion Maps and Coarse-Graining: A Unified Framework for Dimensionality Reduction, Graph Partitioning and Data Set Parametrization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(9):1393–1403, 2006.
- [LL09] J. Li and C.-C. Liu. Power System Reconfiguration based on Multilevel Graph Partitioning. In *PowerTech*, pages 1–5, 2009.
- [Llo82] S. Lloyd. Least Squares Quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.
- [Lov93] L. Lovász. Random Walks on Graphs: A Survey. *Combinatorics, Paul Erdős is Eighty*, 2:1–46, 1993.
- [LoWA] U. o. M. Laboratory of Web Algorithms. Datasets, <http://law.dsi.unimi.it/datasets.php>.
- [LR03] A. Lisser and F. Rendl. Graph Partitioning using Linear and Semidefinite Programming. *Mathematical Programming*, 95(1):91–101, 2003.
- [LR04] K. Lang and S. Rao. A Flow-Based Method for Improving the Expansion or Conductance of Graph Cuts. In *10th Integer Programming and Combinatorial Optimization*, volume 3064 of *LNCS*, pages 383–400. Springer, 2004.
- [LRJL05] H. Li, G. Rosenwald, J. Jung, and C.-C. Liu. Strategic Power Infrastructure Defense. *Proceedings of the IEEE*, 93(5):918–933, 2005.
- [LS12] D. Luxen and D. Schieferdecker. Candidate Sets for Alternative Routes in Road Networks. In *11th Symposium on Experimental Algorithms (SEA’12)*, volume 7276 of *LNCS*, pages 260–270. Springer, 2012.
- [MAB⁺10] G. Malewicz, M. H. Austern, A. J. C. Bik, J. C. Dehnert, I. Horn, N. Leiser, and G. Czajkowski. Pregel: a System for Large-Scale Graph Processing. In *ACM SIGMOD Int. Conference on Management of Data (SIGMOD)*, pages 135–146. ACM, 2010.
- [Mey08] H. Meyerhenke. *Disturbed Diffusive Processes for Solving Partitioning Problems on Graphs*. PhD thesis, Universität Paderborn, 2008.
- [Mey12] H. Meyerhenke. Shape Optimizing Load Balancing for MPI-Parallel Adaptive Numerical Simulations. In Bader et al. [BMSW13], pages 67–82.
- [MK08] I. Moulitsas and G. Karypis. Architecture Aware Partitioning Algorithms. In *8th Int. Conference on Algorithms and Architectures for Parallel Processing (ICA3PP)*, pages 42–53, 2008.
- [MMS06] H. Meyerhenke, B. Monien, and S. Schamberger. Accelerating Shape Optimizing Load Balancing for Parallel FEM Simulations by Algebraic Multigrid. In *20th IEEE Int. Parallel and Distributed Processing Symposium (IPDPS)*, page 57 (CD), 2006.
- [MMS09a] H. Meyerhenke, B. Monien, and T. Sauerwald. A New Diffusion-Based Multilevel Algorithm for Computing Graph Partitions. *Journal of Parallel and Distributed Computing*, 69(9):750–761, 2009.
- [MMS09b] H. Meyerhenke, B. Monien, and S. Schamberger. Graph Partitioning and Disturbed Diffusion. *Parallel Computing*, 35(10–11):544–569, 2009.
- [Mon10] R. Mondaini. *Biomat 2009: International Symposium on Mathematical and Computational Biology, Brasilia, Brazil, 1-6 August 2009*. World Scientific, 2010.

- [MPS07] B. Monien, R. Preis, and S. Schamberger. Approximation Algorithms for Multilevel Graph Partitioning. In T. F. Gonzalez, editor, *Handbook of Approximation Algorithms and Metaheuristics*, chapter 60, pages 60–1–60–15. Taylor & Francis, 2007.
- [MS04] B. Monien and S. Schamberger. Graph Partitioning with the Party Library: Helpful-Sets in Practice. In *16th Symposium on Computer Architecture and High Performance Computing*, pages 198–205, 2004.
- [MS05] H. Meyerhenke and S. Schamberger. Balancing Parallel Adaptive FEM Computations by Solving Systems of Linear Equations. In *11th Int. Euro-Par Conference*, volume 3648 of *LNCS*, pages 209–219. Springer-Verlag, 2005.
- [MS07] J. Maue and P. Sanders. Engineering Algorithms for Approximate Weighted Matching. In *6th Workshop on Experimental Algorithms (WEA)*, volume 4525 of *LNCS*, pages 242–255. Springer, 2007.
- [MS12] H. Meyerhenke and T. Sauerwald. Beyond Good Partition Shapes: An Analysis of Diffusive Graph Partitioning. *Algorithmica*, 64(3):329–361, 2012.
- [MSM09] J. Maue, P. Sanders, and D. Matijevic. Goal Directed Shortest Path Queries Using Precomputed Cluster Distances. *ACM Journal of Experimental Algorithmics*, 14:3.2:1–3.2:27, 2009.
- [MSS+07] R. H. Möhring, H. Schilling, B. Schütz, D. Wagner, and T. Willhalm. Partitioning Graphs to Speedup Dijkstra’s Algorithm. *ACM Journal of Experimental Algorithmics*, 11(2006), 2007.
- [MSS14] H. Meyerhenke, P. Sanders, and C. Schulz. Partitioning Complex Networks via Size-constrained Clustering. In *Proc. of the 13th Int. Symp. on Experimental Algorithms*, LNCS. Springer, 2014.
- [MSS15] H. Meyerhenke, P. Sanders, and C. Schulz. Parallel Graph Partitioning for Complex Networks. In *Proc. 29th IEEE Intl. Parallel & Distributed Processing Symp. (IPDPS 2015)*, 2015. To appear. Preliminary version: <http://arxiv.org/abs/1404.4797>.
- [MSSD13] H. Meuer, E. Strohmaier, H. Simon, and J. Dongarra. June 2013 | TOP500 supercomputer sites. <http://top500.org/lists/2013/06/>, June 2013.
- [MTV91] G. Miller, S.-H. Teng, and S. Vavasis. A Unified Geometric Approach to Graph Separators. In *32nd Symposium on Foundations of Computer Science (FOCS)*, pages 538–547, 1991.
- [New10] M. Newman. *Networks: An Introduction*. Oxford University Press, Inc., New York, NY, USA, 2010.
- [New13] M. E. J. Newman. Community Detection and Graph Partitioning. *CoRR*, abs/1305.4974, 2013.
- [NU13] J. Nishimura and J. Ugander. Restreaming Graph Partitioning: Simple Versatile Algorithms for Advanced Balancing. In *19th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining (KDD)*, 2013.
- [OS10] V. Osipov and P. Sanders. n -Level Graph Partitioning. In *18th European Symposium on Algorithms (ESA): Part I*, volume 6346 of *LNCS*, pages 278–289. Springer, 2010.
- [PB94] J. R. Pilkington and S. B. Baden. Partitioning with Space-filling Curves. Technical Report CS94-349, UC San Diego, Dept. of Computer Science and Engr., 1994.
- [Pel] F. Pellegrini. Scotch Home Page. <http://www.labri.fr/pelegrin/scotch>.
- [Pel94] F. Pellegrini. Static Mapping by Dual Recursive Bipartitioning of Process and Architecture Graphs. In *Scalable High-Performance Computing Conference (SHPCC)*, pages 486–493. IEEE, May 1994.
- [Pel07a] F. Pellegrini. A Parallelisable Multi-Level Banded Diffusion Scheme for Computing Balanced Partitions with Smooth Boundaries. In *13th Int. Euro-Par Conference*, volume 4641 of *LNCS*, pages 195–204, 2007.
- [Pel07b] F. Pellegrini. Scotch and libScotch 5.0 User’s Guide. Technical report, LaBRI, Université Bordeaux I, December 2007.
- [Pel11] F. Pellegrini. Static Mapping of Process Graphs. In C.-E. Bichot and P. Siarry, editors, *Graph Partitioning*, chapter 5, pages 115–136. John Wiley & Sons, 2011.
- [Pel12] F. Pellegrini. Scotch and PT-Scotch Graph Partitioning Software: An Overview. In U. Naumann and O. Schenk, editors, *Combinatorial Scientific Computing*, pages 373–406. CRC Press, 2012.
- [PM07] D. A. Papa and I. L. Markov. Hypergraph Partitioning and Clustering. In T. F. Gonzalez, editor, *Handbook of Approximation Algorithms and Metaheuristics*, chapter 61, pages 61–1–61–19. CRC Press, 2007.
- [Pre99] R. Preis. Linear Time $1/2$ -Approximation Algorithm for Maximum Weighted Matching in General Graphs. In *16th Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 1563 of *LNCS*, pages 259–269. Springer, 1999.
- [PS04] S. Pettie and P. Sanders. A Simpler Linear Time $2/3 - \epsilon$ Approximation for Maximum Weight Matching. *Information Processing Letters*, 91(6):271–276, 2004.
- [PSL90] A. Pothen, H. D. Simon, and K. P. Liou. Partitioning Sparse Matrices with Eigenvectors of Graphs. *SIAM Journal on Matrix Analysis and Applications*, 11(3):430–452, 1990.
- [PZZ13] B. Peng, L. Zhang, and D. Zhang. A survey of Graph Theoretical Approaches to Image Segmentation. *Pattern Recognition*, 46(3):1020 – 1038, 2013.

- [RAK07] U. N. Raghavan, R. Albert, and S. Kumara. Near Linear Time Algorithm to Detect Community Structures in Large-Scale Networks. *Physical Review E*, 76(3), 2007.
- [RPG96] E. Rolland, H. Pirkul, and F. Glover. Tabu Search for Graph Partitioning. *Annals of Operations Research*, 63(2):209–232, 1996.
- [RSB10] D. Ron, I. Safro, and A. Brandt. A Fast Multigrid Algorithm for Energy Minimization under Planar Density Constraints. *Multiscale Modeling & Simulation*, 8(5):1599–1620, 2010.
- [RSB11] D. Ron, I. Safro, and A. Brandt. Relaxation-Based Coarsening and Multiscale Graph Organization. *Multiscale Modeling & Simulation*, 9(1):407–423, 2011.
- [RWSB05] D. Ron, S. Wishko-Stern, and A. Brandt. An Algebraic Multigrid based Algorithm for Bisectioning General Graphs. Technical Report MCS05-01, Department of Computer Science and Applied Mathematics, The Weizmann Institute of Science, 2005.
- [San89] L. A. Sanchis. Multiple-Way Network Partitioning. *IEEE Trans. on Computers*, 38(1):62–81, 1989.
- [Sch04] S. Schamberger. On Partitioning FEM Graphs using Diffusion. In *HPGC Workshop of the 18th International Parallel and Distributed Processing Symposium (IPDPS'04)*. IEEE Computer Society, 2004.
- [Sch07] S. E. Schaeffer. Graph Clustering. *Computer Science Review*, 1(1):27–64, August 2007.
- [Sch13] C. Schulz. *High Quality Graph Partitioning*. PhD thesis. epubli GmbH, 2013.
- [SDM11] J. Shalf, S. Dosanjh, and J. Morrison. Exascale Computing Technology Challenges. In *High Performance Computing for Computational Science (VECPAR)*, volume 6449 of LNCS, pages 1–25. Springer Berlin Heidelberg, 2011.
- [Sen01] N. Sensen. Lower Bounds and Exact Algorithms for the Graph Partitioning Problem Using Multicommodity Flows. In *9th European Symposium on Algorithms (ESA)*, volume 2161 of LNCS, pages 391–403. Springer, 2001.
- [Sim91] H. D. Simon. Partitioning of Unstructured Problems for Parallel Processing. *Computing Systems in Engineering*, 2(2):135–148, 1991.
- [SK12] I. Stanton and G. Kliot. Streaming Graph Partitioning for Large Distributed Graphs. In *18th ACM SIGKDD Int. Conference on Knowledge discovery and data mining (KDD)*, pages 1222–1230. ACM, 2012.
- [SKK97] K. Schloegel, G. Karypis, and V. Kumar. Multilevel Diffusion Schemes for Repartitioning of Adaptive Meshes. *Journal of Parallel and Distributed Computing*, 47(2):109–124, 1997.
- [SKK00] K. Schloegel, G. Karypis, and V. Kumar. A Unified Algorithm for Load-Balancing Adaptive Scientific Simulations. In *Supercomputing 2000*, page 59 (CD). IEEE Computer Society, 2000.
- [SKK02] K. Schloegel, G. Karypis, and V. Kumar. Parallel Static and Dynamic Multi-Constraint Graph Partitioning. *Concurrency and Computation: Practice and Experience*, 14(3):219–240, 2002.
- [SKK03] K. Schloegel, G. Karypis, and V. Kumar. Graph Partitioning for High-Performance Scientific Simulations. In J. Dongarra, I. Foster, G. Fox, W. Gropp, K. Kennedy, L. Torczon, and A. White, editors, *Sourcebook of parallel computing*, pages 491–541. Morgan Kaufmann Publishers, 2003.
- [SNBP11] X. Sui, D. Nguyen, M. Burtscher, and K. Pingali. Parallel Graph Partitioning on Multicore Architectures. In *Languages and Compilers for Parallel Computing*, pages 246–260. Springer, 2011.
- [SS] P. Sanders and C. Schulz. KaHIP – Karlsruhe High Quality Partitioning Homepage. <http://algo2.iti.kit.edu/documents/kahip/index.html>.
- [SS11] P. Sanders and C. Schulz. Engineering Multilevel Graph Partitioning Algorithms. In *19th European Symposium on Algorithms (ESA)*, volume 6942 of LNCS, pages 469–480. Springer, 2011.
- [SS12] P. Sanders and C. Schulz. Distributed Evolutionary Graph Partitioning. In *12th Workshop on Algorithm Engineering and Experimentation (ALENEX)*, pages 16–29, 2012.
- [SS13a] P. Sanders and C. Schulz. High Quality Graph Partitioning. In Bader et al. [BMSW13], pages 19–36.
- [SS13b] P. Sanders and C. Schulz. Think Locally, Act Globally: Highly Balanced Graph Partitioning. In *12th International Symposium on Experimental Algorithms (SEA)*, LNCS. Springer, 2013.
- [SSS12] I. Safro, P. Sanders, and C. Schulz. Advanced Coarsening Schemes for Graph Partitioning. In *11th Int. Symposium on Experimental Algorithms (SEA)*, volume 7276 of LNCS, pages 369–380. Springer, 2012.
- [SST03] M. Sellmann, N. Sensen, and L. Timajev. Multicommodity Flow Approximation used for Exact Graph Partitioning. In *11th European Symposium on Algorithms (ESA)*, volume 2832 of LNCS, pages 752–764. Springer, 2003.
- [ST97] H. D. Simon and S. H. Teng. How Good is Recursive Bisection? *SIAM Journal on Scientific Computing*, 18(5):1436–1445, 1997.
- [ST11] I. Safro and B. Temkin. Multiscale Approach for the Network Compression-Friendly Ordering. *J. Discrete Algorithms*, 9(2):190–202, 2011.

- [Sto06] L. Stock. *Strategic Logistics Management*. Cram101 Textbook Outlines. Lightning Source Inc, 2006.
- [SW04] S. Schamberger and J.-M. Wierum. A Locality Preserving Graph Ordering Approach for Implicit Partitioning: Graph-Filling Curves. In *17th Int. Conference on Parallel and Distributed Computing Systems (PDCS)*, pages 51–57. ISCA, 2004.
- [SW13] S. Salihoglu and J. Widom. GPS: A Graph Processing System. In *Proceedings of the 25th International Conference on Scientific and Statistical Database Management, SSDBM*, pages 22:1–22:12. ACM, 2013.
- [SWC04] A. J. Soper, C. Walshaw, and M. Cross. A Combined Evolutionary Search and Multilevel Optimisation Approach to Graph-Partitioning. *Journal of Global Optimization*, 29(2):225–241, 2004.
- [SWZ02] F. Schulz, D. Wagner, and C. D. Zaroliagis. Using Multi-Level Graphs for Timetable Information. In *4th Workshop on Algorithm Engineering and Experiments (ALENEX)*, volume 2409 of *LNCS*, pages 43–59. Springer, 2002.
- [TBFS00] J. Teresco, M. Beall, J. Flaherty, and M. Shephard. A Hierarchical Partition Model for Adaptive Finite Element Computation. *Computer Methods in Applied Mechanics and Engineering*, 184(2–4):269 – 285, 2000.
- [TGRV00] C. E. Tsourakakis, C. Gkantsidis, B. Radunovic, and M. Vojnovic. Fennel: Streaming Graph Partitioning for Massive Scale Graphs. Technical Report MSR-TR-2012-113, Microsoft Research, 2000.
- [TK08] A. Trifunović and W. J. Knottenbelt. Parallel Multilevel Algorithms for Hypergraph Partitioning. *J. Parallel Distrib. Comput.*, 68(5):563–581, 2008.
- [TLZN08] L. Tang, H. Liu, J. Zhang, and Z. Nazeri. Community Evolution in Dynamic Multi-Mode Networks. In *14th ACM SIGKDD Int. Conference on Knowledge discovery and data mining (KDD)*, pages 677–685. ACM, 2008.
- [UAKI06] B. Ucar, C. Aykanat, K. Kaya, and M. Ikinici. Task Assignment in Heterogeneous Computing Systems. *Journal of Parallel and Distributed Computing*, 66(1):32 – 46, 2006.
- [vBFSS13] R. van Bevern, A. E. Feldmann, M. Sorge, and O. Suchý. On the Parameterized Complexity of Computing Balanced Partitions in Graphs. *CoRR*, abs/1312.7014, 2013.
- [Wal04] C. Walshaw. Multilevel Refinement for Combinatorial Optimisation Problems. *Annals of Operations Research*, 131(1):325–372, 2004.
- [Wal10] C. Walshaw. Variable partition inertia: Graph repartitioning and load balancing for adaptive meshes. In M. Parashar and X. Li, editors, *Advanced Computational Infrastructures for Parallel and Distributed Adaptive Applications*, pages 357–380. 2010.
- [WC00] C. Walshaw and M. Cross. Mesh Partitioning: A Multilevel Balancing and Refinement Algorithm. *SIAM Journal on Scientific Computing*, 22(1):63–80, 2000.
- [WC01] C. Walshaw and M. Cross. Multilevel Mesh Partitioning for Heterogeneous Communication Networks. *Future Generation Comp. Syst.*, 17(5):601–623, 2001.
- [WC02] C. Walshaw and M. Cross. Parallel mesh partitioning on distributed memory systems. In B. Topping, editor, *Computational Mechanics Using High Performance Computing*, pages 59–78. Saxe-Coburg Publications, Stirling, 2002. Invited chapter.
- [WC07] C. Walshaw and M. Cross. JOSTLE: Parallel Multilevel Graph-Partitioning Software – An Overview. In *Mesh Partitioning Techniques and Domain Decomposition Techniques*, pages 27–58. Civil-Comp Ltd., 2007.
- [WCE95] C. Walshaw, M. Cross, and M. G. Everett. A Localized Algorithm for Optimizing Unstructured Mesh Partitions. *Journal of High Performance Computing Applications*, 9(4):280–295, 1995.
- [WCE97] C. Walshaw, M. Cross, and M. G. Everett. Dynamic load-balancing for parallel adaptive unstructured meshes. In *Proc. 8th SIAM Conference on Parallel Processing for Scientific Computing (PPSC’97)*, 1997.
- [Wil91] R. D. Williams. Performance of Dynamic Load Balancing Algorithms for Unstructured Mesh Calculations. *Concurrency: Practice and experience*, 3(5):457–481, 1991.
- [WW93] D. Wagner and F. Wagner. Between Min Cut and Graph Bisection. In *18th Symposium on Mathematical Foundations of Computer Science (MFCS)*, pages 744–750. Springer, 1993.
- [ZS⁺10] M. Zhou, O. Sahní, et al. Controlling Unstructured Mesh Partitions for Massively Parallel Simulations. *SIAM Journal on Scientific Computing*, 32(6):3201–3227, 2010.
- [Zum03] G. Zumbusch. *Parallel Multilevel Methods: Adaptive Mesh Refinement and Loadbalancing*. Teubner, 2003.