

INAUGURAL-DISSERTATION
zur
Erlangung der Doktorwürde
der
Gesamtfakultät für Mathematik, Ingenieur- und Naturwissenschaften
der
Ruprecht–Karls–Universität
Heidelberg

vorgelegt von
Marcelo Fonseca Faraj
aus Belo Horizonte, Brasilien

Tag der mündlichen Prüfung:

Streaming, Local, and Multi-Level (Hyper)Graph Decomposition

Betreuer: Prof. Dr. Christian Schulz

Abstract

(Hyper)Graph decomposition is a family of problems that aim to break down large (hyper)graphs into smaller sub(hyper)graphs for easier analysis. The importance of this lies in its ability to enable efficient computation on large and complex (hyper)graphs, such as social networks, chemical compounds, and computer networks. This dissertation explores several types of (hyper)graph decomposition problems, including graph partitioning, hypergraph partitioning, local graph clustering, process mapping, and signed graph clustering. Our main focus is on streaming algorithms, local algorithms and multilevel algorithms. In terms of streaming algorithms, we make contributions with highly efficient and effective algorithms for (hyper)graph partitioning and process mapping. In terms of local algorithms, we propose sub-linear algorithms which are effective in detecting high-quality local communities around a given seed node in a graph based on the distribution of a given motif. In terms of multilevel algorithms, we engineer high-quality multilevel algorithms for process mapping and signed graph clustering. We provide a thorough discussion of each algorithm along with experimental results demonstrating their superiority over existing state-of-the-art techniques. The results show that the proposed algorithms achieve improved performance and better solutions in various metrics, making them highly promising for practical applications. Overall, this dissertation showcases the effectiveness of advanced combinatorial algorithmic techniques in solving challenging (hyper)graph decomposition problems.

Zusammenfassung

(Hyper-)Graphenzerlegung fasst mehrere Probleme zusammen, bei denen große Graphen oder Hypergraphen zur einfacheren Analyse in kleinere Subgraphen oder Subhypergraphen zerlegt werden. Dies ist bedeutsam, denn es befähigt dazu Berechnungen auf großen und komplexen Graphen und Hypergraphen, wie soziale Netzwerke, chemische Verbindungen und Computernetzwerke, effizient zu machen. Diese Dissertation befasst sich mit der Untersuchung verschiedener Arten von (Hyper-)Graphenzerlegungsproblemen, darunter Graphpartitionierung, Hypergraphpartitionierung, Prozesszuweisung, Signed-Graph-Clustering und Local-Graph-Clustering. Unser Hauptaugenmerk liegt auf Streaming-Algorithmen, lokalen Algorithmen und Multi-Level-Algorithmen. Im Bereich der Streaming-Algorithmen leisten wir Beiträge mit hocheffizienten und effektiven Algorithmen für die (Hyper-)Graphpartitionierung und Prozesszuweisung. Bei den lokalen Algorithmen schlagen wir sub-lineare Algorithmen vor, die auf Grundlage der Verteilung eines bestimmten Motivs hochwertige lokale Cluster um einen gegebenen Startknoten in einem Graphen bestimmen. Was die Multi-Level-Algorithmen betrifft, so entwickeln wir hochwertige Multi-Level-Algorithmen für Prozesszuweisung und Signed-Graph-Clustering. Wir liefern zu jedem Algorithmus eine ausführliche Diskussion zusammen mit experimentellen Ergebnissen, die ihre Überlegenheit gegenüber bestehenden Stand der Technik zeigen. Die Ergebnisse legen dar, dass die vorgeschlagenen Algorithmen eine bessere Laufzeit und bessere Lösungen in verschiedenen Metriken erzielen, womit sie für praktische Anwendungen sehr vielversprechend sind. Insgesamt zeigt diese Dissertation die Effektivität fortschrittlicher kombinatorischer algorithmischer Techniken bei der Lösung anspruchsvoller (Hyper-)Graphenzerlegungsprobleme.

Acknowledgements

I feel incredibly grateful and indebted to my advisor, Prof. Dr. Christian Schulz. Thanks to him, I was introduced to the fascinating world of algorithm engineering and (hyper)graph decomposition problems, which quickly became a passion. His guidance and support throughout my Ph.D. journey were invaluable, and I couldn't have asked for a more nurturing mentor. Under his supervision, I learned the habits and behaviors that make a scientist successful and a leader effective. Working at the University of Vienna and Heidelberg University, thanks to him, allowed me to publish and present many papers at top conferences across the globe, an opportunity I cherish. Most importantly, Christian is not just an excellent mentor but also a kind and friendly person who always motivated and inspired me to pursue my ambitions. Additionally, I am grateful for the chance to meet and interact with some of the giants of Computer Science, such as Peter Sanders, Monika Henzinger, and Robert Tarjan, all thanks to him.

I cannot express enough gratitude towards my wife, Yara Dalilla Xavier Faraj, who has been my rock throughout this journey. Yara's unwavering love and support have been instrumental in helping me navigate the challenges of pursuing a Ph.D. abroad. She has been with me every step of the way, from being my girlfriend to becoming my wife and greatest partner in life. Without her, those four years would have been much more arduous and less happy. I am thankful for Yara's caring and nurturing presence, which provided me with a sense of family during my time away from home. Yara's encouragement and unwavering support have helped me stay motivated and focused on my goals. I feel blessed to have such a wonderful person in my life, and I am excited to continue our journey together as we embark on new adventures.

I would like to express my deep gratitude to two wonderful couples who made my time in Europe even more special. Firstly, my cousin Alfred Lenzner and his wife Marília Simão dos Santos, with whom Yara and I shared some great moments during our stay in Germany. Secondly, I would like to thank Alexander Noe and Anique-Marie Cabardos, with whom we had the pleasure of spending unforgettable moments in both Vienna and Mannheim. Their company, warmth, and positive energy made our experiences in Europe truly unique and enriching.

I want to express my gratitude to all the amazing people who have been a part of my academic journey. First, I would like to thank my colleagues from Heidelberg University, Alexander Noe, Konrad von Kirchbach, Stefan Neumann, Gramoz Goranci, Alexander Svozil, Wolfgang Ost, and Ernestine Großmann, for being wonderful co-Ph.D. students and for making my time in Heidelberg so enjoyable. I would also like to express my sincere thanks to Kathrin Hanauer, Rudolf Hürner, Sagar Kale, Ami Paz, Xiaowei Wu (Eric), Kamer Kaya, Ulrike Frolik-Steffan, Christina Licayan, Catherine Proux-Wieland, and Henrik Reinstädter, for being such amazing colleagues and making my Ph.D. journey an enriching experience. I also want to thank all my co-authors, including Ümit V. Çatalyürek, Karen D. Devine, Lars Gottesbüren, Tobias Heuer, Alexander van der Grinten, Henning Meyerhenke, Jesper Larsson Träff, Peter Sanders, Sebastian Schlag, Christian Schulz, Daniel Seemaier, Dorothea Wagner, Adil Chhabra, Felix Hausberger, and Kamal Eyubov, for collaborating with me on joint papers.

Furthermore, I would like to thank all the students I had the pleasure of supervising in their scientific practica or theses. I would also like to express my appreciation to Bora Uçar for the wonderful opportunity to work together on various projects and share an office with him in Lyon and Heidelberg. I would like to extend my thanks to Peter Bastian for the incredible opportunity to participate in the IGCM 2023 conference in Bangalore. Additionally, I would like to thank Monika Henzinger for the opportunity to be a part of her research group in Vienna. I cannot forget to thank Jesper Larsson Träff, who worked with me in Vienna and with whom I had some amazing philosophical discussions. Finally, I want to thank my dear friend and former (co)advisor João Fernando Machry Sarubbi, who has been a source of inspiration and guidance since my Bachelor's thesis. I am grateful for his continued support and advice as I navigate through my academic and personal life.

I want to express my deepest gratitude to my family. First and foremost, I want to thank my wonderful parents, who have been an unwavering source of love and support throughout my life. Even though they were miles away from me, they never failed to show their constant presence and care in a way that only Brazilian parents could do or even imagine. I also want to thank my amazing siblings, Nara Carolina, Priscila, Sarah, Valéria, and Mário, along with their partners and children. I am blessed to have such a loving and caring family, and I am grateful for all the lessons and values they have instilled in me. Additionally, I want to express my appreciation to my aunts, uncles, cousins, and in-laws for their support, encouragement, and inspiration.

In conclusion, I want to express my deepest gratitude to God for His continuous blessings and guidance in my life. As Aristotle said, "To know thyself is the beginning of all wisdom." However, our knowledge of ourselves is limited, and it is only through God's infinite wisdom and love that we can truly know our purpose and potential. I am grateful for

all the opportunities and experiences that God has presented me with, and for the strength and wisdom He has granted me to face challenges along the way.

Related Publications

Several contributions in this thesis are already published or accepted for publication in conference and journal papers as well as technical reports. Here we list all these publications.

Accepted Papers

- [1] Marcelo Fonseca Faraj, Alexander van der Grinten, Henning Meyerhenke, Jesper Larsson Träff, and Christian Schulz. High-Quality Hierarchical Process Mapping. In *Proceedings of the 18th Symposium on Experimental Algorithms (SEA)*, volume 160 of LIPIcs, pages 4:1–4:15, 2020.
- [2] Marcelo Fonseca Faraj and Christian Schulz. Buffered Streaming Graph Partitioning. *ACM Journal of Experimental Algorithmics*, Volume 27, pages 1.10:1–1.10:26, 2022.
- [3] Marcelo Fonseca Faraj and Christian Schulz. Recursive Multi-Section on the Fly: Shared-Memory Streaming Algorithms for Hierarchical Graph Partitioning and Process Mapping. In *IEEE International Conference on Cluster Computing (CLUSTER)*, volume 9411, pages 473–483, 2022.
- [4] Ümit V. Çatalyürek, Karen D. Devine, Marcelo Fonseca Faraj, Lars Gottesbüren, Tobias Heuer, Henning Meyerhenke, Peter Sanders, Sebastian Schlag, Christian Schulz, Daniel Seemaier, and Dorothea Wagner. More Recent Advances in (Hyper)Graph Partitioning. *ACM Computing Surveys*, Volume 55, Issue 12, Article No 253, pages 1–38, 2022.
- [5] Adil Chhabra, Marcelo Fonseca Faraj, and Christian Schulz. Local Motif Clustering via (Hyper)Graph Partitioning (Extended Abstract). In *Proceedings of the 15th International Symposium on Combinatorial Search (SoCS)*, pages 261–263, AAAI Press, 2022.
- [6] Adil Chhabra, Marcelo Fonseca Faraj, and Christian Schulz. Local Motif Clustering via (Hyper)Graph Partitioning. In *Proceedings of the Symposium on Algorithm Engineering and Experiments (ALENEX)*, pages 96–109. SIAM, 2023.

- [7] Felix Hausberger, Marcelo Fonseca Faraj, and Christian Schulz. A Distributed Multilevel Memetic Algorithm for Signed Graph Clustering (Short Paper). In *Genetic and Evolutionary Computation Conference Companion (GECCO '23 Companion)*, to appear, 2023.
- [8] Kamal Eyubov, Marcelo Fonseca Faraj, and Christian Schulz. FREIGHT: Fast Streaming Hypergraph Partitioning. In *Proceedings of the 21st Symposium on Experimental Algorithms (SEA)*, to appear, 2023.

Papers under Peer Review

- [1] Felix Hausberger, Marcelo Fonseca Faraj, and Christian Schulz. A Distributed Multilevel Memetic Algorithm for Signed Graph Clustering. 2023.
- [2] Adil Chhabra, Marcelo Fonseca Faraj, and Christian Schulz. Faster Local Motif Clustering via Maximum Flows. 2023.

Technical Reports

- [1] Marcelo Fonseca Faraj, Alexander van der Grinten, Henning Meyerhenke, Jesper Larsson Träff, and Christian Schulz. High-Quality Hierarchical Process Mapping. Technical Report, University of Vienna, Humboldt Universität zu Berlin, and Technical University of Vienna, 2020. (arXiv:2001.07134v2)
- [2] Marcelo Fonseca Faraj and Christian Schulz. Buffered Streaming Graph Partitioning. Technical Report, Heidelberg University, 2021. (arXiv:2102.09384)
- [3] Marcelo Fonseca Faraj and Christian Schulz. Recursive Multi-Section on the Fly: Shared-Memory Streaming Algorithms for Hierarchical Graph Partitioning and Process Mapping. Technical Report, Heidelberg University, 2021. (arXiv:2202.00394)
- [4] Ümit V. Çatalyürek, Karen D. Devine, Marcelo Fonseca Faraj, Lars Gottesbüren, Tobias Heuer, Henning Meyerhenke, Peter Sanders, Sebastian Schlag, Christian Schulz, Daniel Seemaier, and Dorothea Wagner. More Recent Advances in (Hyper)Graph Partitioning. Technical Report, 2022. (arXiv:2205.13202)
- [5] Adil Chhabra, Marcelo Fonseca Faraj, and Christian Schulz. Local Motif Clustering via (Hyper)Graph Partitioning. Technical Report, Heidelberg University, 2022. (arXiv:2205.06176)

- [6] Felix Hausberger, Marcelo Fonseca Faraj, and Christian Schulz. A Distributed Multi-level Memetic Algorithm for Signed Graph Clustering. Technical Report, Heidelberg University, 2022. (arXiv:2208.13618)

- [7] Kamal Eyubov, Marcelo Fonseca Faraj, and Christian Schulz. FREIGHT: Fast Streaming Hypergraph Partitioning. Technical Report, Heidelberg University, 2023. (arXiv:2302.06259)

- [8] Adil Chhabra, Marcelo Fonseca Faraj, and Christian Schulz. Faster Local Motif Clustering via Maximum Flows. Technical Report, Heidelberg University, 2023. (arXiv:2301.07145)

Contents

Abstract	v
Zusammenfassung	vii
Acknowledgements	ix
Related Publications	xiii
Contents	xvii
1 Introduction	1
1.1 Motivation	1
1.2 Main Contributions	3
1.2.1 Streaming Algorithms	4
1.2.2 Local Algorithms	5
1.2.3 Multilevel Algorithms	5
1.3 Outline	6
2 Preliminaries	7
2.1 Graphs and Hypergraphs	7
2.2 Partitions and Clusterings	8
2.3 Process Mapping	9
2.4 Flows	10
2.5 Multilevel Scheme	11
2.6 Evolutionary Algorithms	13
2.7 Computational Models	14
2.8 Instances	14
2.8.1 Graphs	15
2.8.2 Hypergraphs	17
2.8.3 Signed Graphs	17
2.9 Machines	17

2.10	Methodology	19
3	Related Work	21
3.1	(Hyper)Graph Partitioning	21
3.1.1	Streaming Graph Partitioning	21
3.1.2	Restreaming Graph Partitioning	22
3.1.3	Buffered Streaming Graph Partitioning	23
3.1.4	Streaming Hypergraph Partitioning	24
3.2	Local Motif Clustering	25
3.3	Process Mapping	26
3.4	Signed Graph Clustering	29
3.4.1	Contraction-Based	30
3.4.2	Evolutionary	30
3.4.3	Integer Linear Programming	31
4	Streaming Algorithms	33
4.1	Buffered Streaming Graph Partitioning	33
4.1.1	HeiStream	34
4.1.2	Experimental Evaluation	40
4.2	Streaming Process Mapping	52
4.2.1	Online Recursive Multi-Section	52
4.2.2	Experimental Evaluation	58
4.3	Streaming Hypergraph Partitioning	64
4.3.1	FREIGHT	64
4.3.2	Experimental Evaluation	70
4.4	Experimental Comparison	75
4.5	Conclusion	77
5	Local Algorithms	79
5.1	Local Motif Clustering	79
5.1.1	LMCHGP	80
5.1.2	SOCIAL	89
5.1.3	Experimental Evaluation	94
5.2	Conclusion	98
6	Multilevel Algorithms	101
6.1	Multilevel Process Mapping	101
6.1.1	Integrated Mapping	102

6.1.2	Experimental Evaluation	107
6.2	Multilevel Signed Graph Clustering	113
6.2.1	Multilevel Algorithm	114
6.2.2	Distributed Evolutionary Algorithm	116
6.2.3	Experimental Evaluation	120
6.3	Conclusion	124
7	Discussion	125
7.1	Conclusion	125
7.2	Future Work	126
	Bibliography	129
	Curriculum Vitae	149
	List of Publications	151

Chapter 1

Introduction

1.1 Motivation

(Hyper)graphs are incredibly versatile mathematical structures that have proven to be useful for modeling and analyzing a wide range of phenomena. These can include anything from social networks and chemical compounds to computer networks and communication systems. As these systems grow in size and complexity, it becomes increasingly difficult to process their underlying (hyper)graphs on a single computer. This is where (hyper)graph decomposition comes in. The idea behind (hyper)graph decomposition is to break down a large (hyper)graph into smaller subgraphs or partitions, which can be processed separately. By doing this, we can enable more efficient computation on large and complex (hyper)graphs, which has applications in many areas of research.

In fact, (hyper)graph decomposition is becoming increasingly important in a wide range of fields. Its potential applications are vast, including scientific simulations, routing, community detection, computational linear algebra, and more. With the ability to efficiently process large and complex (hyper)graphs, we can gain deeper insights into the structure and behavior of the systems they model. This can lead to breakthroughs in fields ranging from social science to engineering. There are several types of (hyper)graph decomposition problems, each of which has its own constraints, objective functions, challenges, and applications. In this thesis, we specifically explore the following problems: graph partitioning, hypergraph partitioning, local graph clustering, process mapping, and signed graph clustering. We now explain the problems in detail.

The *(hyper)graph partitioning* problem is a very important and well-studied problem. It consists of partitioning the nodes of a given (hyper)graph into a predetermined number of blocks with roughly equal cardinality or weight in order to minimize some metric such as the number of cut (hyper)edges. Specifically in the case of hypergraph partitioning, an additional optimization metric is the *connectivity*, which measures the difference between

the number of blocks that have non-empty intersection with cut hyperedges and the number of hyperedges. A natural application for this problem is the distributed processing of (hyper)graphs, in which multiple processors operate on unique nodes of the (hyper)graph and communicate with one another using message-passing in case of (hyper)edges shared across processors. (Hyper)graph partitioning is NP-hard [1] and there can be no approximation algorithm with a constant ratio for general (hyper)graphs [2]. Thus, heuristics are used in practice. A current trend for partitioning huge (hyper)graphs quickly and using low computational resources are streaming algorithms [3–12].

The *process mapping* problem can be seen as an application and a generalization of the graph partitioning problem. In this problem, we are given a *communication graph* and a *topology* containing processing elements (PEs) alongside with their pair-wise distances. The goal of the problem consists of mapping the nodes of the communication graph onto the topology in such a way that roughly the same (weighted) number of nodes are mapped to each PE and the *total communication cost* is minimized. The total communication cost is usually defined as the sum of weights of the cut edges multiplied by the distance between the PEs containing their respective endpoints. A natural application of process mapping consists of assigning interdependent tasks to high-performance (HPC) systems in order to minimize the total execution time. It can also be useful in other areas, such as manufacturing, where optimizing the flow of materials and resources is critical. All hardness results associated with graph partitioning also apply for process mapping, hence heuristics are used in practice. A special case of topology, known as *hierarchical topology*, has been the subject of widespread research. In this topology, PEs are arranged in a multi-layered hierarchy of modules and sub-modules, and the distance between any two PEs depends exclusively on their nearest shared module within the topology.

The *local graph clustering* problem involves a graph and a seed node, where the objective is to obtain a well-characterized cluster that contains the seed node. Conceptually, a well-characterized cluster is a subgraph that consists of many internal edges and few external edges. More specifically, the quality of a cluster can be quantified by metrics such as *conductance* [13]. Applications of this problem include those that only require analyzing a small, localized portion of a graph rather than the entire graph. This is the case for community-detection on Web [14] and social [15] networks as well as structure-discovery in bioinformatics [16] networks, among other real-world problems. Since minimizing conductance is NP-hard [17], approximative and heuristic approaches are used in practice. Furthermore, the nature and scale of this problem necessitates sub-linear methodologies, i.e., which entail time and memory utilization that is only dependent on the size of the discovered cluster, rather than the entire graph. While traditional approaches to local clustering typically consider the edge distribution when evaluating the quality of a local com-

munity [18–22], novel methods [23–26] have shifted focus to finding local communities based on the distribution of *motifs*, higher-order structures within the graph. Empirical evidence shows that this approach, which can be called *local motif clustering*, is effective at detecting high-quality local communities [23].

The *signed graph clustering* problem involves a signed graph, i.e., a graph in which each edge is associated with a weight that can be either positive or negative. The goal is to partition the nodes of the graph into an unspecified number of *well-characterized* clusters. In the context of signed graphs, a well-characterized cluster is defined as being densely connected by edges with positive weight and sparsely connected by edges with negative weight. Conversely, distinct clusters that are well-characterized exhibit dense inter-connections through edges with negative weight and sparse inter-connections through edges with positive weight. Signed graph clustering necessitates distinct metrics and approaches as opposed to those for traditional, unsigned graph clustering. On the one hand, traditional metrics such as *conductance* [13] and *modularity* [27] are insufficient to address negative edge weights. On the other hand, negative edge weights can make clustering structure more explicit. As a consequence, metrics like edge-cut, which simply measures the sum of edge weights between clusters, can be used for evaluating signed clustering. In signed graphs, the edge-cut is bounded by the sum of negative edge weights. Hence, achieving this edge-cut value through clustering will separate all positive edges within clusters and cut all negative edges, fulfilling the purpose of the problem. In many real-world applications, interactions between two entities can be accurately represented by *signed graphs*, i.e., the sign associated with an edge can indicate when the nature of an interaction between nodes is positive (e.g., attraction, similarity, friendship) or negative (e.g., repulsion, difference, animosity). With that being said, graph clustering has practical applications in areas such as criminology, public health, politics, and analysis of social networks [28]. The problem of finding a clustering of signed graphs with minimum edge-cut is NP-hard [29], hence heuristic algorithms are used in practice.

1.2 Main Contributions

In this section we list our main algorithmic contributions in the context of (hyper)graph decomposition. Our algorithms utilize various algorithmic techniques and data structures to achieve improved performance and better solutions compared to state-of-the-art methods. To make the exposition clear, we split our contributions in three groups of algorithms: streaming algorithms, local algorithms and multilevel algorithms. A summary of each contribution is presented along with a selection of experimental results, which illustrate their superiority in various metrics when compared to the state-of-the-art.

1.2.1 Streaming Algorithms

We develop three (buffered) streaming algorithms for (hyper)graph decomposition problems. More specifically, we propose algorithms for (hyper)graph partitioning and process mapping.

We start by proposing a buffered streaming algorithm for graph partitioning. Our algorithm loads a batch of nodes and then builds a model that represents the loaded subgraph as well as the already present partition structure. This model enables us to apply multilevel algorithms and in turn compute much higher quality solutions of huge graphs on cheap machines than previously possible. To partition the model, we develop a multilevel algorithm that optimizes an objective function that has previously shown to be effective for the streaming setting. Surprisingly, this also removes the dependency on the number of blocks from the running time compared to the previous state-of-the-art. Our algorithm computes considerably better solutions than the state-of-the-art using a very small buffer size. In addition, for large numbers of blocks, our algorithm becomes faster than the state-of-the-art.

Our second streaming algorithm is a shared-memory parallel streaming algorithm for the process mapping problem. It is designed to map a streamed communication graph onto a hierarchical topology by performing recursive multi-sections on the fly. If a hierarchy is not specified as an input, our approach can also be used as a general tool to solve the graph partitioning problem. Our approach is the first streaming algorithm for the process mapping problem. Furthermore, in the context of non-buffered streaming graph partitioning, it has a considerably lower running time complexity in comparison with state-of-the-art. Our experiments indicate that our algorithm is both faster and produces better process mappings than competing streaming tools. In case of graph partitioning, our framework is up to two orders of magnitude faster at the cost of 5% more cut edges compared to a state-of-the-art algorithm.

Our third streaming algorithm solves the hypergraph partitioning problem by extending from a state-of-the-art streaming algorithm for graph partitioning. By using an efficient data structure, we make the overall running of our algorithm linearly dependent on the pin-count of the hypergraph and the memory consumption linearly dependent on the numbers of nets and blocks. The results of our extensive experimentation showcase the promising performance of our algorithm as a highly efficient and effective solution for streaming hypergraph partitioning. Our algorithm demonstrates competitive running time with the Hashing algorithm, with a difference of a maximum factor of four observed on three fourths of the instances. Significantly, our findings highlight the superiority of our algorithm over all existing (buffered) streaming algorithms and even an in-memory algorithm HYPE, with respect to both weighted number of cut hyperedges and connectivity measures.

1.2.2 Local Algorithms

We develop two local algorithms for graph decomposition. In particular, both algorithms are designed to solve the local motif clustering problem. Our algorithms starts by building a (hyper)graph model which represents the motif-distribution around the seed node on the original graph. While the graph model is exact for motifs of size at most three, the hypergraph model works for arbitrary motifs and is designed such that an optimal solution in the (hyper)graph model minimizes the motif conductance in the original network.

In our first algorithm, the (hyper)graph model is then partitioned using a powerful multi-level hypergraph or graph partitioner in order to directly minimize the motif conductance of the corresponding partition in the original graph. Extensive experiments evaluate the trade-offs between the two different models. Moreover, when using the graph model for triangle motifs, our algorithm computes communities that have on average one third of the motif conductance value than communities computed by the state-of-the-art while being faster on average and removing the necessity of a preprocessing motif-enumeration on the whole network.

In our second algorithm, we transform the hypergraph model into a flow model based on the fast and effective algorithm *max-flow quotient-cut improvement* (MQI) [30]. We show that a non-trivial maximum flow exists if and only if a superior solution exists, which is obtained automatically. In our experiments with the triangle motif, our flow-based algorithm produces better communities than the state-of-the-art, while also being up to multiple orders of magnitude faster.

1.2.3 Multilevel Algorithms

We develop two multilevel algorithms for graph decomposition. Multilevel algorithms consist of three main phases: coarsening, construction, and uncoarsening. In coarsening, the graph is recursively contracted into a sequence of smaller graphs which maintain some general structure. In construction, an initial solution is computed on the smallest graph. In uncoarsening, the contractions are recursively undone, and local search methods are used to refine the solution induced by each level of contraction. We propose algorithms for process mapping and signed graph clustering.

We propose and engineer multiple setups of a multilevel algorithm for the process mapping problem. Important ingredients of our algorithm include fast label propagation, more localized local search, initial partitioning, as well as a compressed data structure to compute processor distances without storing a distance matrix. Moreover, our algorithm is able to exploit a given hierarchical structure of the distributed system under consideration. Experiments indicate that our algorithm speeds up the overall mapping process and, due to the integrated multilevel approach, also finds much better solutions in practice. For example,

one configuration of our algorithm yields similar solution quality as the previous state-of-the-art in terms of mapping quality for large numbers of partitions while being a factor 9.3 faster. Compared to the currently fastest iterated multilevel mapping algorithm Scotch, we obtain 16% better solutions while investing slightly more running time.

Our last proposed algorithms are designed to solve the signed graph clustering problem by leveraging some of the most effective techniques from graph partitioning that minimize edge-cut. We engineer all the details of a multilevel algorithm, which encompasses a coarsening-uncoarsening process and efficient local search methods. We also introduce a memetic algorithm that utilizes our multilevel algorithm and further enhances it with natural multilevel recombination and mutation operations. We also parallelize our approach using a scalable coarse-grained island-based strategy which has already shown to be scalable in practice. Experimental results demonstrate that our memetic algorithm outperforms the state-of-the-art with respect to edge-cut, producing significantly better solutions.

1.3 Outline

This dissertation is organized as follows. We begin in Chapter 2 by presenting preliminaries and basic concepts that are used throughout this thesis. We continue by elaborating related work in Chapter 3. Our algorithmic contributions are presented in Chapter 4, Chapter 5, and Chapter 6, which correspond to streaming algorithms, local algorithms, and multilevel algorithms, respectively. Specific conclusions are given in respective chapters dedicated to each family of algorithms, and a general conclusion is provided in Chapter 7.

Chapter 2

Preliminaries

In this chapter, we present the basic concepts used in this dissertation.

2.1 Graphs and Hypergraphs

Let $G = (V = \{0, \dots, n-1\}, E)$ be an *undirected graph* with no multiple or self edges allowed, such that $n = |V|$ and $m = |E|$. Let $c : V \rightarrow \mathbb{R}_{\geq 0}$ be a node-weight function, and let $\omega : E \rightarrow \mathbb{R}_{> 0}$ be an edge-weight function. We generalize c and ω functions to sets, such that $c(V') = \sum_{v \in V'} c(v)$ and $\omega(E') = \sum_{e \in E'} \omega(e)$. Let $N(v) = \{u : \{v, u\} \in E\}$ be the *open neighborhood* of v , and let $N[v] = N(v) \cup \{v\}$ be the *closed neighborhood* of v . We generalize the notations $N(\cdot)$ and $N[\cdot]$ to sets, such that $N(V') = \cup_{v \in V'} N(v)$ and $N[V'] = \cup_{v \in V'} N[v]$. A graph $G' = (V', E')$ is said to be a *subgraph* of $G = (V, E)$ if $V' \subseteq V$ and $E' \subseteq E \cap (V' \times V')$. When $E' = E \cap (V' \times V')$, G' is the subgraph *induced* in G by V' . Let $\overline{V'} = V \setminus V'$ be the *complement* of a set $V' \subseteq V$ of nodes. Let a *motif* μ be a connected graph. *Enumerating* the motifs μ in a graph G consists building the collection M of all occurrences of μ as a subgraph of G . Let $d(v)$ be the *degree* of node v and Δ be the maximum degree of G . Let $d_\omega(v)$ be the *weighted degree* of a node v and Δ_ω be the maximum weighted degree of G . Let $d_\mu(v)$ be the *motif degree* of a node v , i.e., the number of motifs $\mu \in M$ which contain v . We generalize the notations $d(\cdot)$, $d_\omega(\cdot)$, and $d_\mu(\cdot)$ to sets, such that the *volume* of V' is $d(V') = \sum_{v \in V'} d(v)$, the *weighted volume* of V' is $d_\omega(V') = \sum_{v \in V'} d_\omega(v)$, and the *motif volume* of V' is $d_\mu(V') = \sum_{v \in V'} d_\mu(v)$. Let a *spanning forest* of G be an acyclic subgraph of G containing all its nodes. Let the *arboricity* of G be the minimum number of spanning forests of G necessary to cover all its edges.

Let an *undirected signed graph* be a graph $G = (V = \{0, \dots, n-1\}, E)$ which expands the above definition to permit edges with both positive and negative weights. More specifically, let $\omega : E \rightarrow \mathbb{R} \setminus \{0\}$ be a *signed edge-weight function* associated with a signed graph. Let E^- denote the set of edges with negative weight and E^+ denote the set of edges with

positive weight, such that $E^- \cup E^+ = E$, $E^- \cap E^+ = \emptyset$, $m^- = |E^-|$, and $m^+ = |E^+|$. Let $N^-(v) = N(v) \cap E^-$ denote the neighbors of v which are connected to v by an edge with negative weight. Let $N^+(v) = N(v) \cap E^+$ denote the neighbors of v which are connected to v by an edge with positive weight.

Let $H = (V = \{0, \dots, n-1\}, E)$ be an *undirected hypergraph* with no multiple or self hyperedges allowed, with $n = |V|$ nodes and $m = |E|$ hyperedges (or *nets*). A net is defined as a subset of V . The nodes that compose a net are called its *pins*. Let $c : V \rightarrow \mathbb{R}_{\geq 0}$ be a node-weight function, and let $w : E \rightarrow \mathbb{R}_{> 0}$ be a net-weight function. We generalize c and w functions to sets, such that $c(V') = \sum_{v \in V'} c(v)$ and $w(E') = \sum_{e \in E'} w(e)$. A node $v \in V$ is *incident* to a net $e \in E$ if $v \in e$. Let $I(v)$ be the set of incident nets of v , let $d(v) = |I(v)|$ be the *degree* of v , and let $d_w(v) = w(I(v))$ be the *weighted degree* of v . We generalize the notations $d(\cdot)$ and $d_w(\cdot)$ to sets, such that the *volume* of V' is $d(V') = \sum_{v \in V'} d(v)$ and the *weighted volume* of V' is $d_w(V') = \sum_{v \in V'} d_w(v)$. Two nodes are *adjacent* if both are incident to the same net. Let the number of pins $|e|$ in a net e be the *size* of e . Given a cluster $V' \subseteq V$, the *cut* or *cut-net* $cut(V')$ of V' consists of the total weight of the nets crossing the cluster, i.e., $cut(V') = \sum_{e \in E'} w(e)$, in which $E' = \{e \in E : \exists i, j \mid e \cap V' \neq \emptyset, e \cap \bar{V}' \neq \emptyset, i \neq j\}$.

2.2 Partitions and Clusterings

The k -way (*hyper*)graph partitioning problem consists of assigning each node of a (hyper)graph to exactly one of k distinct *blocks* respecting a balancing constraint in order to minimize the weight of the (hyper)edges running between the blocks, i.e., the edge-cut (resp. cut-net). More precisely, it partitions V into k blocks V_1, \dots, V_k (i.e., $V_1 \cup \dots \cup V_k = V$ and $V_i \cap V_j = \emptyset$ for $i \neq j$), which is called a k -partition of the (hyper)graph. The *edge-cut* (resp. *cut-net*) of a k -partition consists of the total weight of the *cut edges* (resp. *cut nets*), i.e., edges (resp. nets) crossing blocks. More formally, let the edge-cut (resp. cut-net) be $\sum_{i < j} \omega(E')$, in which $E' = \{e \in E, \exists \{u, v\} \subseteq e : u \in V_i, v \in V_j, i \neq j\}$ is the *cut-set* (i.e., the set of all cut nets). The *balancing constraint* demands that the sum of node weights in each block does not exceed a threshold associated with some allowed *imbalance* ϵ . More specifically, $\forall i \in \{1, \dots, k\} : c(V_i) \leq L_{\max} = \lceil (1 + \epsilon) \frac{c(V)}{k} \rceil$. For each net e of a hypergraph, $\Lambda(e) := \{V_i \mid V_i \cap e \neq \emptyset\}$ denotes the *connectivity set* of e . The *connectivity* $\lambda(e)$ of a net e is the cardinality of its connectivity set, i.e., $\lambda(e) := |\Lambda(e)|$. The so-called *connectivity metric* $(\lambda-1)$ is computed as $\sum_{e \in E'} (\lambda(e) - 1) \omega(e)$, where E' is the cut-set.

In the *local graph clustering* problem, a graph $G = (V, E)$ and a seed node $u \in V$ are taken as input and the goal is to detect a *well-characterized cluster* (or *community*) $C \subset V$ containing u . A high-quality cluster C usually contains nodes that are densely connected to

one another and sparsely connected to \bar{C} . There are many functions to quantify the quality of a cluster, such as *modularity* [27] and *conductance* [13]. The conductance metric is defined as $\phi(C) = |E'| / \min(d(C), d(\bar{C}))$, where $E' = E \cap (C \times \bar{C})$ is the set of edges shared by a cluster C and its complement. *Local motif graph clustering* is a generalization of local graph clustering where a motif μ is taken as an additional input and the computed cluster optimizes a clustering metric based on μ . In particular, the *motif conductance* $\phi_\mu(C)$ of a cluster C is defined by Benson et al. [31] as a generalization of the conductance in the following way: $\phi_\mu(C) = |M'| / \min(d_\mu(C), d_\mu(\bar{C}))$, where M' are all the motifs μ which contain at least one node in C and one node in \bar{C} . Note that, if the motif under consideration is simply an edge, then $|M'|$ is the edge-cut and $\phi_\mu(C) = \phi(C)$.

Let a *clustering* of a graph $G = (V, E)$ be any partition of V , i.e., a set of *blocks* or *clusters* $V_1, \dots, V_t \subset V$ such that $V_1 \cup \dots \cup V_t = V$ and $V_i \cap V_j = \emptyset$, where $t \in [1, n]$. An abstract view of the clustering is a *quotient graph* Q , in which nodes represent clusters and edges are induced by the connectivity between clusters. We call *neighboring clusters* a pair of clusters that is connected by an edge in the quotient graph. A node $v \in V_i$ that has a neighbor $w \in V_j, i \neq j$, is a boundary node. The *edge-cut* of a clustering consists of the total weight of the edges crossing clusters (also called *cut edges*), i.e., $\sum_{i < j} \omega(E_{ij})$, where $E_{ij} = \{\{u, v\} \in E : u \in V_i, v \in V_j\}$. Note that $\frac{1}{2}d_\omega^-(V)$ is an absolute lower bound for the edge-cut of any clustering. Let a signed graph be *balanced* if there exists a clustering of it with edge-cut equal to $\frac{1}{2}d_\omega^-(V)$, i.e., where all edges with positive weight are inside clusters and all edges with negative weight are cut. If such clustering does not exist, e.g., a triangle with only one edge with negative weight, we say that the signed graph is *unbalanced*. Let the *signed graph clustering* (SGC) consist of obtaining a clustering of an undirected signed graph G in order to minimize the edge-cut.

2.3 Process Mapping

For process mapping applications of hierarchical partitions, assume that we have n processes and a topology containing k PEs. Let $C \in \mathbb{R}^{n \times n}$ denote the communication matrix and let $D \in \mathbb{R}^{k \times k}$ denote the (implicit) topology matrix or distance matrix. In particular, $C_{i,j}$ represents the required amount of communication between processes i and j , while $D_{x,y}$ represents the cost of each communication between PEs x and y . Hence, if processes i and j are respectively assigned to PEs x and y , or vice-versa, the communication cost between i and j will be $C_{i,j}D_{x,y}$. Throughout this thesis, we assume that C and D are symmetric – otherwise one can create equivalent problems with symmetric inputs [32].

In particular, for process mapping applications tackled in this paper, we assume that topologies are organized as homogeneous hierarchies. In this case $S = a_1 : a_2 : \dots : a_\ell$ is a

sequence describing the hierarchy of a supercomputer. The sequence should be interpreted as each processor having a_1 cores, each node a_2 processors, each rack a_3 nodes, and so forth, such that the total number of PEs is $k = \prod_{i=1}^{\ell} a_i$. Without loss of generality, we assume that $a_i \geq 2, \forall i \in \{1, \dots, \ell\}$. Let $D = d_1 : d_2 : \dots : d_{\ell}$ be a sequence describing the distance between PEs within each hierarchy level, meaning that the distance between two cores in the same processor is d_1 , the distance between two cores in the same node but in different processors is d_2 , the distance between two cores in the same rack but in different nodes is d_3 , and so forth. The *process mapping* problem consists of assigning the nodes of a graph to PEs in a communication topology while respecting a balancing constraint in order to minimize the total communication cost. Let $\Pi : \{1, \dots, n\} \mapsto \{1, \dots, k\}$ be the function that maps a node onto its PE. The objective of process mapping is to minimize Equation (2.1).

$$J(C, D, \Pi) := \sum_{i,j} C_{i,j} D_{\Pi(i), \Pi(j)}. \quad (2.1)$$

The exact total communication cost of a given mapping of processes onto PEs depends on a combination of *bandwidth* and *latency*, which depends on a multitude of factors in practice. The total communication cost in latency-based topologies is also known as the *Coco(.)* or *hop-byte* objective function. An alternative objective function to the total communication cost is the bandwidth-based metric *maximum congestion*, which is defined as the maximum number of message exchanges through any link of the topology graph. Another possible objective is the *maximum dilation*, which is defined as the maximum communication cost directly associated with a pair of PEs for a given mapping.

2.4 Flows

Let $G_f = (V_f, E_f)$ be a flow graph. A flow graph has one source node $s \in V_f$, one sink node $t \in V_f$, and a set of remaining nodes $V \setminus \{s, t\}$. All edges $e = (u, v)$ in a flow graph are directed and associated with a nonnegative capacity $cap(u, v)$. An s-t flow is a function $f : V_f \times V_f \rightarrow \mathbb{R}_{>0}$ which satisfies a *capacity* constraint, i. e., $f(u, v) \leq cap(u, v)$, a *symmetry* constraint, i. e., $\forall u, v \in V_f : f(u, v) = -f(v, u)$, and a *flow conservation* constraint, i. e., $\forall u \in V_f \setminus \{s, t\} : \sum_{v \in V_f} f(u, v) = 0$. An edge (u, v) is called *saturated* if $cap(u, v) = f(u, v)$; The total amount of flow moved from s to t is defined as the *value* $|f|$ of f and is computed as follows: $|f| = \sum_{u \in V_f} f(u, t) = \sum_{v \in V_f} f(s, v)$. A given s-t flow f in G_f is *maximum* if, for any s-t flow f' in G_f , $|f'| \leq |f|$. Let $G_r = (V_f, E_r)$ be the *residual graph* associated with a given flow f on G_f , such that $E_r = \{(u, v) \in V_f \times V_f : cap(u, v) - f(u, v) > 0\}$. According to the Max-Flow Min-Cut Theorem [33], the value $|f|$ of a maximum s-t flow f on G_f equals the weight of a minimum s-t cut on G_f , i. e., a

2-way partition of G_f where edge weights equal edge capacities, s and t are in distinct blocks, and the total weight of the cut edges is minimum. To find the sink side of the minimum cut associated with a maximum flow in G_f , a reverse breadth-first search can be performed on G_f starting at the sink node t .

For each node u in a flow graph V_f , let $exc(u) = \sum_{v \in V_f} f(u, v)$ be its *excess* value and $d(u)$ be its potential. A node u is called *active* if $exc(u) > 0$. An edge (u, v) is called *admissible* if $cap(u, v) - f(u, v) > 0$ and $d(u) = d(v) + 1$. The push-relabel [34] algorithm builds a maximum flow by computing a succession of *preflows*, i. e., flows where the flow conservation constraint is substituted by $\forall u \in V_f \setminus \{s, t\} : exc(u) \geq 0$. In the initial preflow, all out-edges of s are saturated, $\forall u \in V_f \setminus \{s\} : d(u) = 0$, and $d(s) = |V_f|$. The initial preflow is evolved via operations *push*, i. e., sending as much flow as possible from an active node through an admissible edge, and *relabel*, i. e., increasing the potential of a node until it becomes active. Preflows induce minimum sink-side cuts, so a maximum flow and a minimum cut are obtained once no node is active.

A common technique to solve flow and cut problems on hypergraphs consists of transforming them in directed graphs and then applying traditional graph-based techniques on them. Among the existing transformations [35, 36], we highlight *clique expansion*, *star expansion*, and *Lawler expansion*. In the *clique expansion*, each net is represented by a clique, i. e., a set of edges connecting each pair of its pins in both directions. In this approach, the weight of each edge is equal to weight of the corresponding net e divided by $|e| - 1$ and parallel edges are substituted by a single edge whose weight is the sum of the weights of the removed edges. In the *star expansion*, each net is represented by an auxiliary artificial node connected to its pins by edges in both directions. In this expansion, the edges have the same weight as the corresponding net. In the *Lawler expansion*, each net e is represented by two auxiliary artificial nodes w_1 and w_2 and a collection of edges. In particular, there is a directed edge (w_1, w_2) which has the same weight as the corresponding net. Additionally, each pin of the corresponding net has an out-edge to w_1 and an in-edge from w_2 , each of them with weight infinity. The three transformation approaches are exemplified in Figure 2.1.

2.5 Multilevel Scheme

In this section, we characterize the *multilevel* scheme, which is a successful heuristic for clustering or partitioning large (hyper)graphs. Before describing the multilevel scheme, we need to define the terms contraction and uncontraction. *Contracting* an edge $e = \{u, v\}$ consists of replacing the nodes u and v by a new node x connected to the former neighbors of u and v and setting $c(x) = c(u) + c(v)$. If replacing edges of the form $\{u, w\}, \{v, w\}$ would generate two parallel edges $\{x, w\}$, a single edge with $\omega(\{x, w\}) = \omega(\{u, w\}) +$




Figure 2.1. Three existing approaches to represent a hypergraph using a directed graph. Nodes and nets of the original hypergraph are respectively represented by black circles and colored areas around them. Auxiliary artificial nodes and edges are respectively represented by circles and arrows, both with the same color as the corresponding net. Bidirectional arrows represent a pair of edges in both directions. Solid edges have finite weight while dashed edges have infinite weight.

$\omega(\{v, w\})$ is inserted. *Contracting* a cluster of nodes $C = \{u_1, \dots, u_\ell\}$ involves replacing them with a new node v whose weight is the sum of the weights of the clustered nodes and is connected to all elements $w \in \bigcup_{i=1}^{\ell} N(u_i)$, $\omega(\{v, w\}) = \sum_{i=1}^{\ell} \omega(\{u_i, w\})$. This ensures the transfer of partition from a coarser to a finer level maintains the edge-cut, as exemplified in Figure 2.3. We define the *contraction* operator as $/$ such that G/V' , with $V' \subseteq V$, is the (hyper)graph obtained by contracting the nodes from V' on G . The *uncontraction* of a node undoes the contraction.

We describe the multilevel scheme within the scope of graph partitioning, although the basic idea is also extensible to other problems, such as process mapping, hypergraph partitioning and graph clustering. A *multilevel approach* consists of three main phases. In the *contraction* (coarsening) phase, successive approximations of an original input graph are created. The contractions quickly reduce the size of the graph and stop as soon as it becomes sufficiently small to be partitioned by an expensive algorithm. In the construction phase, an *initial partitioning* is obtained by partitioning the coarsest graph. Due to the way we define contraction, every partition of the coarsest level implies a corresponding partition of the input graph with equal edge-cut and balance. In the *local improvement* (or *uncoarsening*) phase, we uncontract previously contracted nodes to go back through each level, from the coarsest approximation to the original graph. After each uncoarsening, local improvement algorithms move nodes between blocks in order to improve the objective function or balance. Local search moves nodes between blocks to reduce the objective, as exemplified in Figure 2.4.

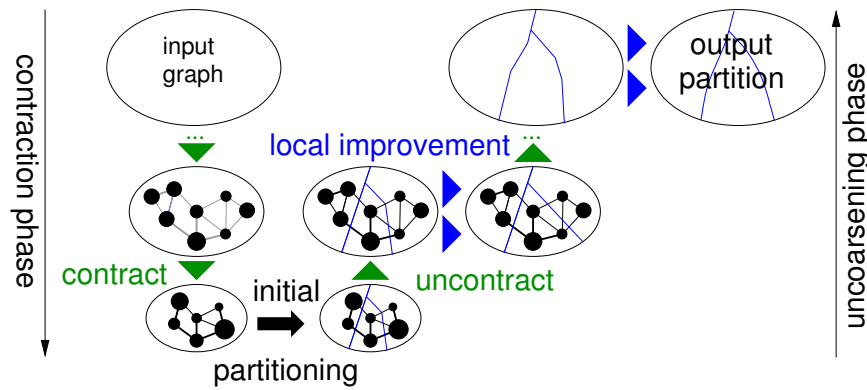


Figure 2.2. Multilevel scheme (adapted from [37]).




Figure 2.3. Contraction of a clustering [38]. Each cluster is represented by a different color on the left hand side graph. Each cluster on the left-hand side is contracted to a single node on the right hand side.

2.6 Evolutionary Algorithms

Evolutionary or *Memetic algorithms* [39] are population-based heuristics that mimic natural evolution to optimize a problem. They use a convenient *notation* to represent the decision variables, called *individuals*, and evaluate their quality through a *fitness* function (simulation or mathematical [40]). A *population* of individuals evolves during the algorithm. *Recombination* [41] (crossover) *exploits* characteristics of previous individuals to create new and better solutions. *Mutation* [42] introduces random variations to *explore* the search space and escape local optima. Evolutionary algorithms do not guarantee optimality, but are effective in exploring and exploiting the solution space.




Figure 2.4. Typical step of a local search algorithm. Cluster assignments are indicated by colors. In this example, a single node is moved to another cluster in order to decrease the overall edge-cut.

2.7 Computational Models

Streaming algorithms usually follow a load-compute-store logic shown in Figure 2.5. The classic streaming model is the *one-pass* model, in which the nodes are loaded one at a time alongside with their (hyper)edges, then some logic is applied to permanently assign them to blocks. This logic can be as simple as a Hashing function or as complex as scoring all blocks based on some objective and then assigning the node to the block with highest score. When assignment decisions of an algorithm for the current node depend on the previous decisions, an algorithm in the model has to store the assignment of the previous loaded nodes and hence needs $\Omega(n)$ space. An extended version of this model is called the *buffered streaming* model. More precisely, a δ -sized *buffer* or *batch* of input nodes with their neighborhood is repeatedly loaded. Partition/block assignment decisions have to be made after the whole buffer is loaded. While we investigate the dependence of our algorithm on this parameter, in practice the parameter will depend on the amount of available memory on a machine. The parameter can be dynamically chosen such that the buffer is “full” if $\Theta(n)$ space has been loaded from the disk. Hence the buffered streaming model asymptotically does not need more space than a one-pass streaming algorithm if this setting is used. This holds true even in the worst case: when a node has degree close to n .

2.8 Instances

Throughout this thesis we present experiments on various kinds of (signed) (hyper)graphs. In this section, we summarize the main properties, the source, and the area of application of the graphs.

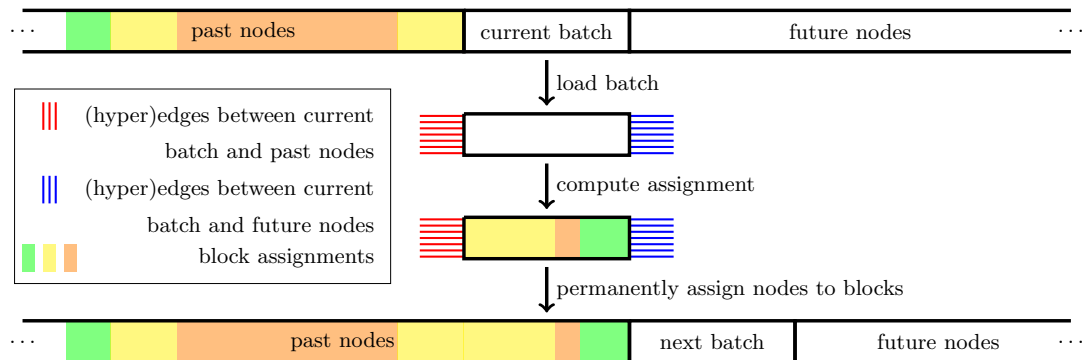


Figure 2.5. Typical layout of streaming algorithm for (hyper)graph partitioning.

2.8.1 Graphs

Experiments with Streaming Algorithms. In our experiments with streaming algorithms, we use graphs from various sources [43–47]. Most of the considered graphs were used for benchmark in previous works on graph partitioning. The graphs wiki-Talk and web-Google, as well as most networks of co-purchasing, roads, social, web, autonomous systems, citations, circuits, similarity, meshes, and miscellaneous are publicly available either in [43] or in [44]. We also use graphs such as eu-2005, in-2004, uk-2002, and uk-2007-05, which are available at the 10th DIMACS Implementation Challenge website [45]. Finally, we include some artificial random graphs. We use the name rggX for *random geometric graph* with 2^X nodes where nodes represent random points in the unit square and edges connect nodes whose Euclidean distance is below $0.55\sqrt{\ln n/n}$. We use the name delX for a graph based on a Delaunay triangulation of 2^X random points in the unit square [46]. We use the name RHGX for random hyperbolic graphs [47, 48] with 10^8 nodes and $X \times 10^9$ edges. Basic properties of the graphs under consideration can be found in Table 2.1. For our experiments, we split the graphs in three disjoint sets. A *tuning* set for the parameter study experiments, a *test* set for the comparisons against the state-of-the-art, and a set of *huge graphs* for special larger scale tests. In any case, when streaming the graphs we use the natural given order of the nodes.

Experiments with Multilevel Process Mapping. The instances used in our experiments with multilevel process mapping come from various sources to test our algorithm. We use the largest six graphs from Chris Walshaw’s benchmark archive [49]. Graphs derived from sparse matrices have been taken from the SuiteSparse Matrix Collection [50]. We also use graphs from the 10th DIMACS Implementation Challenge [45] website. Here, rggX and delX are defined as before. The graphs af_shell9, thermal2, and nlr are from the matrix and

Graph	n	m	Type	Graph	n	m	Type
Tuning Set				Test Set			
coAuthorsCiteseer	227 320	814 134	Citations	Dubcova1	16 129	118 440	Meshes
citationCiteseer	268 495	1 156 647	Citations	hcircuit	105 676	203 734	Circuit
amazon0312	400 727	2 349 869	Co-Purch.	coAuthorsDBLP	299 067	977 676	Citations
amazon0601	403 364	2 443 311	Co-Purch.	Web-NotreDame	325 729	1 090 108	Web
amazon0505	410 236	2 439 437	Co-Purch.	Dblp-2010	326 186	807 700	Citations
roadNet-PA	1 087 562	1 541 514	Roads	ML_Laplace	377 002	13 656 485	Meshes
com-Youtube	1 134 890	2 987 624	Social	coPapersCiteseer	434 102	16 036 720	Citations
soc-lastfm	1 191 805	4 519 330	Social	coPapersDBLP	540 486	15 245 729	Citations
roadNet-TX	1 351 137	1 879 201	Roads	Amazon-2008	735 323	3 523 472	Similarity
in-2004	1 382 908	13 591 473	Web	eu-2005	862 664	16 138 468	Web
G3_circuit	1 585 478	3 037 674	Circuit	web-Google	916 428	4 322 051	Web
soc-pokec	1 632 803	22 301 964	Social	ca-hollywood-2009	1 087 562	1 541 514	Roads
as-Skitter	1 694 616	11 094 209	Aut.Syst.	Flan_1565	1 564 794	57 920 625	Meshes
wiki-topcats	1 791 489	28 511 807	Social	Ljournal-2008	1 957 027	2 760 388	Social
roadNet-CA	1 957 027	2 760 388	Roads	HV15R	2 017 169	162 357 569	Meshes
wiki-Talk	2 388 953	4 656 682	Web	Bump_2911	2 911 419	62 409 240	Meshes
soc-flixster	2 523 386	7 918 801	Social	del21	2 097 152	6 291 408	Artificial
del22	4 194 304	12 582 869	Artificial	rgg21	2 097 152	14 487 995	Artificial
rgg22	4 194 304	30 359 198	Artificial	FullChip	2 987 012	11 817 567	Circuit
del23	8 388 608	25 165 784	Artificial	soc-orkut-dir	3 072 441	117 185 083	Social
rgg23	8 388 608	63 501 393	Artificial	patents	3 750 822	14 970 766	Citations
Huge Graphs				cit-Patents	3 774 768	16 518 947	Citations
uk-2005	39 459 923	783 027 125	Web	soc-LiveJournal1	4 847 571	42 851 237	Social
twitter7	41 652 230	1 202 513 046	Social	circuit5M	5 558 326	26 983 926	Circuit
sk-2005	50 636 154	1 810 063 330	Web	italy-osm	6 686 493	7 013 978	Roads
soc-friendster	65 608 366	1 806 067 135	Social	great-britain-osm	7 733 822	8 156 517	Roads
er-fact1.5s26	67 108 864	907 090 182	Artificial				
RHG1	100 000 000	1 000 913 106	Artificial				
RHG2	100 000 000	1 999 544 833	Artificial				
uk-2007-05	105 896 555	3 301 876 564	Web				

Table 2.1. Graphs for experiments with streaming algorithms.

the numeric section of the DIMACS benchmark set. The graphs eur and deu are large road networks of Europe and Germany taken from [51]. Basic properties of the graphs under consideration can be found in Table 2.2.

Experiments with Local Motif Clustering. In our experiments with local motif clustering, we use graphs from various sources [43–45]. Most of the considered graphs were used for benchmark in previous works in the area. Basic properties of the graphs under consideration can be found in Table 2.3. For our experiments, we split the graphs in two disjoint sets: a *tuning* set for the parameter study experiments and a *test* set for the comparisons against the state-of-the-art. The graphs in the test set are exactly the graphs used in [23].

Graph	n	m	Type	Graph	n	m	Type
Tuning Graphs				Large Walshaw Graphs			
fe_rotor	99 617	662 431	Miscellaneous	598a	110 971	741 934	Meshes
598a	110 971	741 934	Miscellaneous	fe_ocean	143 437	409 593	Miscellaneous
ecology2	999 999	1 997 996	Circuit	144	144 649	1 074 393	Meshes
G3_circuit	1 585 478	3 037 674	Circuit	wave	156 317	1 059 331	Meshes
del22	4 194 304	12 582 869	Artificial	m14b	214 765	1 679 018	Meshes
rgg22	4 194 304	30 359 198	Artificial	auto	448 695	3 314 611	Meshes
UF Graphs				Large Other Graphs			
cop20k_A	99 843	1 262 244	Miscellaneous	af_shell9	504 855	8 542 010	Miscellaneous
2cubes_sphere	101 492	772 886	Miscellaneous	thermal2	1 227 087	3 676 134	Miscellaneous
thermomech_TC	102 158	304 700	Miscellaneous	nlr	4 163 763	12 487 976	Meshes
cfid2	123 440	1 482 229	Miscellaneous	deu	4 378 446	5 483 587	Roads
boneS01	127 224	3 293 964	Miscellaneous	del23	8 388 608	25 165 784	Artificial
Dubcova3	146 689	1 744 980	Miscellaneous	rgg23	8 388 608	63 501 393	Artificial
bmwera_1	148 770	5 247 616	Numerical	del24	16 777 216	50 331 601	Artificial
G2_circuit	150 102	288 286	Circuit	rgg24	16 777 216	132 557 200	Artificial
shipsec5	179 860	4 966 618	Miscellaneous	eur	18 029 721	22 217 686	Roads
cont-300	180 895	448 799	Miscellaneous				

Table 2.2. Graphs for multilevel process mapping experiments.

2.8.2 Hypergraphs

In our experiments, we consider hypergraphs that have been used for benchmark in previous works on hypergraph partitioning. We use the same benchmark as in [52]. This consists of 310 hypergraphs from three benchmark sets: 18 hypergraphs from the ISPD98 Circuit Benchmark Suite [53], 192 hypergraphs based on the University of Florida Sparse Matrix Collection [50], and 100 instances from the international SAT Competition 2014 [54]. The SAT instances were converted into hypergraphs by mapping each boolean variable and its complement to a node and each clause to a net. From the Sparse Matrix Collection, one matrix was selected for each application area that had between 10 000 and 10 000 000 columns. The matrices were converted into hypergraphs using the row-net model, in which each row is treated as a net and each column as a node.

2.8.3 Signed Graphs

For experiments with signed graphs, we use the real-world signed graphs listed in Table 2.4. They were all obtained from the public graph collections SNAP [55] and KONECT [56].

2.9 Machines

We now describe the five machines that are used in the following chapters. With hyper-threading, all used machine are capable of handling a number of threads equal to twice the

Graph	n	m	Triangles	Graph	n	m	Triangles
Tuning Set				Test Set			
citationCiteseer	268 495	1 156 647	847 420	com-amazon	334 863	925 872	667 129
coAuthorsCiteseer	227 320	814 134	2 713 298	com-dblp	317 080	1 049 866	2 224 385
amazon0312	400 727	2 349 869	3 686 467	com-youtube	1 134 890	2 987 624	3 056 386
amazon0505	410 236	2 439 437	3 951 063	com-livejournal	3 997 962	34 681 189	177 820 130
amazon0601	403 364	2 443 311	3 986 507	com-orkut	3 072 441	117 185 083	627 584 181
del22	4 194 304	12 582 869	8 436 672	com-friendster	65 608 366	1 806 067 135	4 173 724 142
del23	8 388 608	25 165 784	16 873 359				
soc-pokec	1 632 803	22 301 964	32 557 458				
rgg22	4 194 304	30 359 198	85 962 754				
rgg23	8 388 608	63 501 393	188 022 664				
in-2004	1 382 908	13 591 473	464 257 245				

Table 2.3. Graphs for local motif clustering experiments.

Graph	n	m	Type
bitcoinalpha	3 783	14 081	Commerce
bitcoinotc	5 881	21 434	Commerce
elec	7 118	100 355	Miscellaneous
chess	7 301	32 650	Miscellaneous
slashdot081106	77 357	466 666	Social
slashdot-zoo	79 116	465 840	Social
slashdot090216	81 871	495 666	Social
slashdot090221	82 144	498 532	Social
wikiconflict	118 100	1 461 058	Social
epinions	131 828	708 507	Social
wikisigned-k2	138 592	712 337	Social

Table 2.4. Signed real-world graphs for signed graph clustering experiments.

number of their cores. **Machine A** has a two six-core Intel Xeon E5-2630 processor running at 2.8 GHz, 64 GB of main memory, and 3 MB of L2-Cache. It runs Ubuntu GNU/Linux 20.04.1 and Linux kernel version 5.4.0-48. **Machine B** has a four-core Intel Xeon E5420 processor running at 2.5 GHz, 16 GB of main memory, and 24 MB of L2-Cache. The machine runs Ubuntu GNU/Linux 20.04.1 and Linux kernel version 5.4.0-65. **Machine C** has a sixteen-core Intel Xeon Silver 4216 processor running at 2.1 GHz, 100 GB of main memory, 16 MB of L2-Cache, and 22 MB of L3-Cache running Ubuntu 20.04.1. The machine can handle 32 threads with hyperthreading. **Machine D** has a sixty-four-core AMD EPYC 7702P processor running at 2.0 GHz, 1 TB of main memory, 32 MB of L2-Cache, and 256 MB of L3-Cache. **Machine E** has four sixteen-core Intel Xeon Haswell-EX E7-8867 processors running at 2.5 GHz, 1 TB of main memory, and 32768 KB of L2-Cache. The machine runs Debian GNU/Linux 10 and Linux kernel version 4.19.67-2.




Figure 2.6. Performance profile comparing running times of three algorithms.

2.10 Methodology

In this section, we describe our experimental methodology. The concepts and details presented here apply throughout the thesis, unless explicitly stated otherwise. Depending on the focus of the experiment, we measure running time, memory consumption and/or solution quality. Depending on the problem under study, solution quality is measured in terms of edge-cut, cut-net, connectivity, communication cost, and/or motif conductance. In general, we perform ten repetitions per algorithm and instance using different random seeds for initialization, and we compute the arithmetic average of the computed objective functions and running time per instance. When further averaging over multiple instances, we use the geometric mean in order to give every instance the same influence on the *final score*.

For a solution generated by an algorithm A , we express its score σ_A (which can measure running time or solution quality) using one or more of the following tools: *improvement* over an algorithm B , computed as $(\frac{\sigma_B}{\sigma_A} - 1) * 100\%$; *ratio*, computed as $(\frac{\sigma_A}{\sigma_{max}})$ with σ_{max} being the maximum score for a given instance or for a given x-axis parameter among all competitors including A ; *relative value* over an algorithm B , computed as $(\frac{\sigma_A}{\sigma_B})$. Bar charts and boxplots are also employed to represent our findings. We use bar charts to visualize the average value of an objective function in relation to an x-axis parameter, where each algorithm is represented by vertical bars of a given color with origin on the x-axis. The bars for every value of this parameter have a common origin and are arranged in terms of their height, allowing all heights to be visible. We use boxplots to give a clear picture of the dataset distribution by displaying the minimum, maximum, median, first and third quartiles, while disregarding outliers.

We also present *performance profiles* which relate the running time (resp. solution quality) of a group of algorithms to the fastest (resp. best) one on a per-instance basis. Their x-axis shows a factor τ while their y-axis shows the percentage of instances for which A has up to τ times the running time (resp. solution quality) of the fastest (resp. best) algo-

rithm. Achieving higher fractions at smaller τ values is considered better. As an example, Figure 2.6 compares the running time of three different algorithms using a performance profile. In this plot, the algorithm SOCIAL is the fastest one for 87% of the instances, while the algorithms LMCHGP and MAPPR are the fastest ones for 12% and 1% of the instances, respectively (see $\tau = 10^0$). Note also that the running time of SOCIAL is within a factor 1.18 of the running times of the fastest competitors for all instances (see y-axis = 1). Furthermore, the running time of LMCHGP and MAPPR are within a factor 10 of the running times of the fastest competitors for 67% and 46% of the instances, respectively (see $\tau = 10^1$).

Chapter 3

Related Work

In this chapter, we give an overview of previous work that has been done on streaming (hyper)graph partitioning, local motif clustering, process mapping, and signed graph clustering.

3.1 (Hyper)Graph Partitioning

There is massive research on (hyper)graph partitioning in the literature. The most prominent tools to partition (hyper)graphs in memory include PaToH [57], Metis [58], hMetis [59], Scotch [60], HYPE [61], KaHIP [62], KaMinPar [63], KaHyPar [52], Mt-KaHyPar [64], and mt-KaHIP [65]. The readers are referred to [66–68] for extensive material and references. Here, we focus on the results specifically related to the scope of this paper. In particular, we provide an exhaustive account of the state-of-the-art solutions available for the (hyper)graph partitioning problem in the context of the (buffered) (re)streaming model.

3.1.1 Streaming Graph Partitioning

Stanton and Kliot [69] propose heuristics to tackle the graph partitioning problem in the streaming model. Among their most prominent heuristic is the one-pass method *linear deterministic greedy* (LDG) which produces solutions with the best overall edge-cut. In this algorithm, node assignments prioritize blocks containing more neighbors and use a penalty multiplier to control imbalance. Particularly, a node v is assigned to the block V_i that maximizes $|V_i \cap (v)|\Phi(i)$ with $\Phi(i)$ being a multiplicative degrading factor defined as $(1 - \frac{|V_i|}{L_{\max}})$. The intuition is that the degrading factor avoids to overload blocks that are already very heavy. In case of ties on the objective function, LDG assigns the node to the block with fewer nodes. Overall, LDG partitions a graph in $O(m + nk)$ time. Moreover, the authors also propose a simple one-pass methods based on *hashing*, which has running time $O(n)$ and produces a poor edge-cut.

Later, Stanton [70] studies the streaming graph partitioning problem from a more theoretical perspective. The author proves that no algorithm can obtain an $o(n)$ -approximation with a random or adversarial stream ordering. Next, two variants of a randomized greedy algorithm are analyzed by using a novel coupling to finite Polya Urn [71] processes, which intuitively explains the performance of the compared algorithms.

Tsourakakis et al. [3] propose Fennel, a simple one-pass partitioning algorithm which adapts the widely-known clustering objective function *modularity* [27]. Roughly speaking, Fennel assigns a node v to a block V_i in order to maximize an expression of type $|V_i \cap (v)| - f(|V_i|)$, where $f(|V_i|)$ is an additive degrading factor. More specifically, the authors defined the Fennel objective function by using $f(|V_i|) = \alpha\gamma|V_i|^{\gamma-1}$, in which γ is a free parameter and $\alpha = m \frac{k^{\gamma-1}}{n^\gamma}$. After a parameter tuning made by the authors, Fennel uses $\gamma = \frac{3}{2}$, which implies $\alpha = \sqrt{k} \frac{m}{n^{3/2}}$. Although the objective function penalizes imbalanced partitions, the authors define the possibility of a hard constraint to enforce balancing. In their experiments, Fennel cuts fewer edges than LDG [69] and, for some instances, it cuts roughly the same number of edges as an offline partitioning algorithm.

Zhang et al. [72] propose AKIN, a streaming graph partitioning algorithm for distributed graph storage systems. AKIN is able to partition graphs where the number of nodes n is not known in advance by allowing the migration of nodes between blocks over time. The assignment decisions are mainly based on the *similarity* between nodes, which is evaluated with the *Jaccard similarity coefficient* [73]. Given the (partial) neighborhoods of two nodes, this coefficient is defined as the ratio of their intersection over their union. Initially, AKIN assumes a *base block* for each node, which is given by a hash function. This base block is taken as a preliminary node assignment as well as a constant-time index for reaching information in the distributed graph storage. More specifically, for each node, AKIN stores in its base block a fixed-length list containing its loaded neighbors with largest degree. This list is used for computing the similarity between nodes. As soon as an edge is loaded, AKIN assigns and migrates it and both of its endpoints to the block which maximizes a similarity-based heuristic. In the experimental evaluation, the version of AKIN keeping up to 100 neighbors per node cuts fewer edges compared to Fennel while maintaining equivalent imbalance and spending 10% more running time.

3.1.2 Restreaming Graph Partitioning

Nishimura and Ugander [74] introduce a restreaming approach to partition the nodes of a graph. Their approach is motivated by scenarios where the same graph is streamed multiple times. In their model, a one-pass partitioning algorithm can pass multiple times through the entire input while the edge-cut is iteratively reduced. The authors propose ReLDG and

ReFennel, which are respective restreaming adaptations of linear deterministic greedy [69] (LDG) and Fennel [3]. On the one hand, ReLDG modifies the objective of LDG to account only for node assignments performed during the current pass when computing block weights. On the other hand, ReFennel uses the same objective as Fennel during restreaming, but its additive balancing degrading factor is increased after each pass in order to enforce balance. Additionally, the authors prove that ReFennel converges after a finite number of restreams even without increasing the degrading factor. Their experiments confirm that their restreaming methods can iteratively reduce edge-cut.

Awadelkarim and Ugander [4] investigate how the order in which nodes are streamed influences one-pass graph partitioning. The authors introduce the notion of *prioritized streaming*, where (re)streamed nodes are statically or dynamically reordered based on some predefined priority. Their approach, which is a prioritized version of ReLDG, uses multiplicative weights of restreaming algorithms and adapts the ordering of the streaming process inspired by balanced label propagation. In their experiments, the authors consider a wide range of stream orderings. The minimum overall edge-cut is obtained using a dynamic node ordering based on their own metric *ambivalence*. This approach is closely followed by a static ordering based on node degree.

3.1.3 Buffered Streaming Graph Partitioning

Patwary et al. [75] propose WStream, a simple streaming graph partitioning algorithm that keeps a sliding window in memory. The authors allow a few hundred nodes in the sliding window in order to obtain more information about a node before it is permanently assigned to a block based on a greedy function. As soon as a node is allocated to a block, one more node is loaded from the input stream into the sliding window, which keeps the window size constant. In their experiments, WStream cuts fewer edges than LDG and more edges than offline multilevel partitioning for most tested graphs.

Jafari et al. [5] perform graph partitioning using a buffered streaming computational model. The authors propose a shared-memory algorithm which repeatedly loads a batch of nodes from the stream input, partitions it using a multilevel scheme, and then permanently assigns the nodes to blocks. Their multilevel scheme is based on a simplified structure where the one-pass algorithm LDG is used for coarsening, computing an initial partition, and refining it. They parallelize LDG in a node-centric way by simply splitting nodes among processors, which yields a parallelization of the three steps of their multilevel scheme. In their experiments, their algorithms cuts fewer edges than LDG while scaling better than offline partitioning algorithms.

3.1.4 Streaming Hypergraph Partitioning

Alistarh et al. [8] propose Min-Max, a one-pass streaming algorithm to assign the nodes of a hypergraph to blocks. For each block, this algorithm keeps track of nets which contain pins in it. This implies a memory consumption of $O(mk)$, which is more than the typical memory consumption of a streaming algorithm for graph partitioning. When a node is loaded, Min-Max allocates it to the block containing the largest intersection with its nets while respecting a hard constraint for load balance. The authors theoretically prove that their algorithm is able to recover a hidden *co-clustering* with high probability, where a co-clustering is defined as a simultaneous clustering of nodes and hyperedges. In the experimental evaluation, Min-Max outperforms five intuitive streaming approaches with respect to load imbalance, while producing solutions up to five times more imbalanced than internal-memory algorithms such as hMetis.

Taşyaran et al. [9] propose improved versions of the algorithm Min-Max [8]. The authors present Min-Max-N2P, a modified version of Min-Max that stores blocks containing each net’s pins instead of storing nets per block, as done in Min-Max. In their experiments, Min-Max-N2P is three orders of magnitude faster than Min-Max while keeping the same cut-net. The authors also introduce three algorithms with reduced memory usage compared to Min-Max: Min-Max-L ℓ , a modification of Min-Max-N2P that employs an upper-bound ℓ to limit memory consumption per net, Min-Max-BF which utilizes Bloom filters for membership queries, and Min-Max-MH that uses hashing functions to replace the connectivity information between blocks and nets. In their experiments, their three algorithms reduce the running time in comparison to Min-Max, especially Min-Max-L ℓ and Min-Max-MH, which are up to four orders of magnitude faster. On the other hand, the three algorithms generate solutions with worse cut-net than Min-Max, especially Min-Max-MH, which increases the cut-net by up to an order of magnitude. Moreover, the authors propose a technique to improve the partitioning decision in the streaming setting by including a buffer to store some nodes and their net sets. This approach operates similarly to Min-Max-N2P, but with the added ability to revisit buffered nodes and adjust their partition assignment based on the connectivity metric. The authors propose three algorithms using this buffered approach: REF that buffers every incoming node but only reassigns those that may improve connectivity, REF_RLX that buffers all nodes and reassigns all nodes in the buffer, and REF_RLX_SV that only buffers nodes with small net sets and reassigns all nodes in the buffer. Their experimental results show that the use of buffered approaches leads to a 5-20% improvement in partitioning quality compared to non-buffered approaches, but with a trade-off of increased runtime.

3.2 Local Motif Clustering

Motif-based clustering has been widely studied in the literature, with works such as [23, 76–79] partitioning all the nodes of a graph into clusters based on motifs. We also address the topic of clustering based on motifs, but our focus is on identifying clusters in the immediate vicinity of a specific seed node, rather than on the entire graph. Several works [21, 80–83] propose local clustering algorithms, but they do not focus on optimizing for motif-based metrics like this thesis. Instead, they use metrics based on edges, like conductance and modularity. In this section, we review previous work on local clustering based on motifs, which is the focus of this thesis.

Rohe and Qin [84] propose a local clustering algorithm based on triangle motifs. Their algorithm starts by initializing a cluster containing only the seed node, and iteratively grows this cluster. Particularly, the algorithm greedily inserts nodes contained in at least a pre-defined number of cut triangles. Huang et al. [85] recover local communities containing a seed node in online and dynamic setups based on higher-order graph structures named Trusses [86]. They define the k -truss of a graph as its largest subgraph whose edges are all contained in at least $(k - 2)$ triangle motifs, hence trusses are a graph structure based on the frequency of triangles. The authors use indexes to search for k -truss communities in time proportional to the size of the recovered community.

Yin et al. [23] propose MAPPR, a local motif clustering algorithm based on the Approximate Personalized PageRank (APPR) method. In a preprocessing phase, MAPPR enumerates the motif of interest in the entire input graph and constructs a weighted graph W , in which edges only exist between nodes that appear in at least one instance of the motif, and their edge weight is equal to the number of occurrences of the motif containing these two endpoints. Afterward, MAPPR uses an adapted version of the APPR method to find local communities in the weighted graph constructed in the preprocessing phase. MAPPR is able to extract local communities from directed input graphs, something that cannot be done using APPR alone.

Zhang et al. [24] propose LCD-Motif, an algorithm that addresses the local motif clustering problem using a modified version of the spectral method. LCD-Motif has two main differences in comparison to the traditional spectral motif clustering method. First, instead of computing singular vectors, the algorithm performs random walks to identify potential members of the searched cluster. They use the span of a few dimensions of vectors, obtained through random walks, as an approximation for the local motif spectra. Second, Instead of using k -means for clustering, LCD-Motif searches for the minimum 0-norm vector within the previously mentioned span, which must contain the seed nodes in its support vector.

Meng et al. [25] propose FuzLhocd, a local motif clustering algorithm that uses fuzzy arithmetic to optimize a modified version of modularity. Given seed node, FuzLhocd starts by detecting probable core nodes of the targeted local community using fuzzy membership. After identifying the probable core nodes of the target local community using fuzzy membership, the algorithm expands these nodes using another fuzzy membership to form a cluster.

Zhou et al. [87] propose HOSPLOC, a local motif clustering algorithm that uses a motif-based random walk to compute a distribution vector, which is then truncated and used in a vector-based partitioning method. The algorithm begins by approximately estimating the distribution vector through a motif-based random walk. To further refine the computation and focus on the local region, HOSPLOC sets all small vector entries to 0. After this preprocessing step, the algorithm applies a vector-based partitioning method [88] on the resulting distribution vector in order to identify a local cluster.

Shang et al. [89] propose HSEI, a local motif clustering algorithm that uses motif and edge information to grow a cluster from a seed node. The algorithm begins by creating an initial cluster consisting of only the seed node. It then adds nodes to the cluster from the seed's neighborhood, selecting them based on their motif degree. The cluster is expanded using a motif-based extension of the modularity function.

3.3 Process Mapping

The algorithms for process mapping can be categorized in two groups. On the one hand, the *single-phase* or *integrated* algorithms combine process mapping with graph partitioning [90, 91], such that the objective of the partitioning – commonly edge-cut – is typically replaced by a function that measures communication cost. On the other hand, the two-phase algorithms decouple partitioning and mapping [32, 92–95]. In a *two-phase* or *decoupled* algorithm, a default graph partitioning algorithm is used to partition a communication graph into k blocks while typically minimizing edge-cut. Afterwards, the quotient graph of the partitioned communication graph is mapped onto PEs in order to minimize the total communication cost J (Equation (2.1)) or other objective function. This step is also known as one-to-one (process) mapping. To the best of our knowledge there is no streaming algorithm specifically designed adapted to optimize for the process mapping problem. We now describe the state-of-the-art for the process mapping problem. The reader is referred to [67, 96] for older works on process mapping.

Jostle and Scotch are integrated algorithms to solve the process mapping problem. Jostle integrates local search into a multilevel scheme to partition the model of computation and communication. In this scheme, it solves the problem on the coarsest level and afterwards performs refinements based on the user-supplied network communication model. Scotch

performs dual recursive bipartitioning to compute a mapping. More precisely, it starts the recursion considering all given processes and PEs. At each recursion level, it bipartitions the communication graph and also the distance graph with a graph bipartitioning algorithm. The first (resp., second) block of the communication graph is then assigned to the first (resp., second) block of the distance graph.

Müller-Merbach [94] propose a greedy construction method to obtain an initial permutation for the one-to-one process mapping problem. The method roughly works as follows: Initially compute the total communication volume for each process and also the sum of distances from each core to all the others. Afterwards, the algorithm proceeds in rounds. In each round, the process with the largest communication volume is assigned to the core with the smallest total distance.

A method to improve an already given solution for one-to-one process mapping was proposed in [93]. The method repeatedly tries to perform swaps in the assignment in a pair-exchange neighborhood $N(\Pi)$ that contains all permutations that can be reached by swapping two elements in Π . Here, swapping two elements means that $\Pi^{-1}(i)$ will be assigned to processor j and $\Pi^{-1}(j)$ will be assigned to processor i after the swap is done. The algorithm then looks at the neighborhood in a cyclic manner. A swap is performed if it reduces the objective. To reduce the runtime, Brandfass et al. [32] introduced a couple of modifications to speed up the algorithm, such as only considering pairs for swapping that can reduce the objective or partitioning the search space into s consecutive blocks and only performing swaps inside those blocks.

Deveci et al. [97] propose a greedy construction algorithm and two refinement algorithms to map processes to PEs based on a topology graph. The construction algorithm iteratively picks the process with highest connectivity to the already mapped processes. Then, this process is assigned to the module which minimizes hop-bytes, which is found by performing a breadth-first search on the topology graph. Both their refinement methods are based on process swapping, but one of them aims at minimizing hop-bytes while the other one aims at minimizing congestion. The authors experimentally compare their algorithms on a torus-based system against the default mapper of the system as well as Scotch [98] and LibTopoMap [99]. Their algorithms induce performance increases of 43% and 23% for a communication-only application and a sparse matrix vector multiplication, respectively, compared to the default mapper of the system, and have the best overall results among the competitors.

Deveci et al. [100] also propose a scheme that exploits geometric partitioning to map processors to PEs. The application data and PEs are partitioned separately using the Multi-Jagged geometric partitioner [101]. Blocks from the data partition are then mapped to corresponding blocks in the PE partition, effectively assigning interdependent data to “nearby”

PEs in the network. The method is appropriate for applications that have geometric coordinates as well as graph data.

Vogelstein et al. [102] solve an application-agnostic problem where there are two graphs with the same number of nodes which are bijectively mapped onto one another in order to minimize the number of induced edge disagreements. The authors call this problem *graph matching problem* and formulate it as a quadratic assignment problem (QAP). They solve it with a non-linear approximation algorithm based on the gradient vector. Their algorithm has complexity $O(k^3)$ and performs better than the previous state-of-the-art [103–105] regarding running time and objective function for over 80% of the QAPLIB benchmark library [106].

Glantz et al. [95] propose two greedy algorithms to bijectively map blocks onto PEs assuming a communication graph. Their most successful algorithm, GreedyAllC, is an adaptation of the greedy mapping algorithm proposed by Müller-Merbach [94]. Glantz et al. [95] modify this algorithm by scaling the distance with the amount of communication to be done. This modification improves the overall mapping quality with respect to the quality measures maximum congestion and maximum dilation.

Glantz et al. [107] propose a local improvement algorithm for one-to-one process mapping in which the hardware topology is a partial cube, i.e., an isometric subgraph of a hypercube. The authors exploit the regularity of these topologies to label PEs as well as processes with bit-strings along convex cuts. These bit-strings permit fast computation of distances between PEs and the implementation of effective hierarchical refinement methods to improve the mapping induced by the labels. Their experimental results show that their algorithm reduces the total communication cost of the mappings produced by state-of-the-art mapping algorithms [95, 98] in a range from 6% to 34% while the mapping time stays within the same order of magnitude.

Schulz and Träff [92] solve the process mapping for a hierarchical topology using a two-phase approach. Their algorithm uses the state-of-the-art partitioner KaHIP [108] for the partitioning phase and then a multi-section for the one-to-one mapping phase. The referred multi-section algorithm recursively partitions the quotient graph of the previously obtained partition throughout the layers of the communication topology in a top-down direction. After this algorithm obtains a one-to-one mapping, a light-weight modification of the swap-based refinement method proposed by Brandfass et al. [32] is executed to further minimize the communication cost. In particular, the authors experimentally show that N_C^{10} , which restricts swapping to processes that have a distance smaller than 10 in the communication graph, is an adequate choice to obtain good solutions with a moderate running time. Their experiments also show that their approach produces mappings with lower total communication cost in comparison with alternative approaches which combine the same KaHIP-based first phase with state-of-the-art algorithms for one-to-one mapping [94, 95, 99]. Kirchbach et al. [109]

further improve the two-pass approach by Schulz and Träff [92]. In particular, the authors specialize the partitioning phase to apply KaHIP multiple times, namely throughout the hierarchical topology. Kirchbach et al experimentally show that their best algorithm is faster while also decreasing the total communication cost in comparison to the algorithm from Schulz and Träff [92].

Predari et al. [110] propose a single-phase distributed algorithm to solve the process mapping for a hierarchical topology. The authors model the system hierarchy as an implicit labeled tree and label the nodes of the communication graph in order to implicitly induce the mapping. The algorithm optimizes the mapping by using an adapted version of parallel label propagation where the labeling scheme is used for quick gain computations. In their experiments, the proposed algorithm has good scalability for up to thousands of PEs while achieving a total communication cost smaller than the distributed algorithms ParMetis [58] and ParHIP [111] and comparable cost to the state-of-the-art sequential mapping algorithm.

Kirchbach et al. [112] address the process mapping problem assuming processes that communicate in a sparse stencil pattern and PEs organized in Cartesian grids. First, the authors prove that the Cartesian mapping is already NP-hard for a two-dimensional Cartesian grid and a one-dimensional stencil. Then three fast construction algorithms which exploit the regularity of the problem are proposed. Their algorithm with best overall results is the *stencil strips* algorithm, which partitions the grid into strips of lengths close to the scaled length of an optimal bounding rectangle of the coordinates of the target stencil. In their experimental evaluation using the MPI_Neighbor_alltoall routine of MPI, their algorithms are up to two orders of magnitude faster than general purpose graph mapping tools such as VieM [113] while resulting in a similar communication performance. Moreover, their algorithms are up to three times faster than other Cartesian grid mapping algorithms while resulting in a much better communication performance.

3.4 Signed Graph Clustering

There is a huge body of research on signed graph clustering. Roughly the same problem is solved under different names and terms in the literature such as *community detection (or mining) in signed networks* [114–116], *correlation clustering* [117, 118], and *clique partitioning problem* [119–122]. Among the most frequent approach for solving the problem, we mention spectral clustering [123–126] and meta-heuristics [115, 127–131]. The readers are referred to the surveys [124, 132, 133] for extensive material and references. Here, we focus on the results specifically related to the scope of this thesis.

3.4.1 Contraction-Based

Keuper et al. [134] propose an algorithm named Greedy Additive Edge Contraction (GAEC). GAEC is an adaptation of the greedy agglomeration algorithm where the criterion to evaluate the strength of interactions between clusters is the summed weight of the edges shared by them. Bailoni et al. [135] propose GASP, a framework for hierarchical agglomerative clustering on weighted signed graphs. The authors prove that their framework is a generalization to many existing clustering algorithms such as GAEC [134] and introduce new algorithms based on yet unexplored special instances of their framework, such as HCC-Sum. GASP is a bottom-up approach where nodes are initially assigned to their own clusters, which are then iteratively merged in a pair-wise fashion. The framework should be configured to use a specific criterion to evaluate the interaction between a pair of clusters. Some of these criteria are the (average) weighted edge-cut between the clusters and the weight of their shared edge with highest absolute weight. The unconstrained variety of GASP starts by merging clusters with the largest positive interaction and stops once the remaining clusters share only negative interactions. The constrained variety of GASP introduces cannot-merge constraints between pairs of clusters and terminates when all the remaining clusters are constrained against their neighbors. This approach greedily selects the pair of clusters with highest absolute interaction, which can be either positive or negative. When it is positive, the clusters are merged, otherwise a constraint is added to prohibit the clusters from being merged until the end of the algorithm. Experimentally, the best algorithms contemplated by GASP with respect to the minimization of edge-cut are GAEC [134] and HCC-Sum. Similarly to GAEC and HCC-Sum, our multilevel algorithm also constructs a clustering based on successive contractions of the graph.

Hua et al. [136] propose FCSG, an algorithm for signed graph clustering which combines a random walk gap (RWG) mechanism with a greedy shrinking method. The RWG mechanism makes a random walk on a version of the graph where negative edges are removed and another random walk on a version of the graph where negative edges are made positive. Using the two random walk graphs that are constructed, the gap between the two graphs is calculated to give information on the natural clustering structure of the signed graph. This information is then used to build a new signed graph whose edge weights more accurately reflect the natural clustering structure of the input signed graph. A clustering is then computed on this new signed graph based on a greedy shrinking approach.

3.4.2 Evolutionary

Che et al. [115] propose MACD-SN, an evolutionary algorithm for signed graph clustering optimizing for *signed modularity*. In their algorithm, an individual represents a clustering

and the *signed modularity* [137] metric is used as fitness function. An individual is denoted by an n -sized string where each node is assigned to a cluster. Individuals for the initial population are built in two steps. First, randomly assign all nodes to clusters. Second, go through nodes in order to assign each node u to its neighboring cluster V_i which minimizes the expression $NID(u, V_i) = |N^- \cap V_i| + |N^+ \cap \overline{V_i}|$. Classical approaches are implemented for selection, recombination, and mutation. Namely, a tournament selection, a randomized two-way crossover recombination, and a uniform random mutation. The authors also implement a mutation operator which finds the cluster whose nodes contain the highest average value of $NID(u, V_i)$ and then simply assign each of its nodes u to the cluster which minimizes $NID(u, V_i)$. Finally, a local search algorithm is executed on the best mutated offspring in each generation. While the operators of MACD-SN are rather simple and traditional, the operators of our evolutionary algorithm are based on a sophisticated multilevel algorithm which we propose in this thesis. However, the approach only scales to networks with a few thousand nodes.

3.4.3 Integer Linear Programming

Many integer linear programming (ILP) formulations were proposed for solving the signed graph clustering problem in order to optimize for edge-cut, as we do here. Grötschel and Wakabayashi [119] propose an ILP formulation consisting of $\Theta(n^2)$ binary variables representing pairs of nodes and $\Theta(n^3)$ transitivity constraints. These constraints ensure that, if two nodes u, v are clustered together, then any other node w will be clustered together with u if, and only if, it is also clustered together with v . Dinh and Thai [120] remove a set of redundant constraints from the ILP by Grötschel and Wakabayashi [119] for a special case of the signed clustering problem and Miyauchi and Sukegawa [121] extend the ILP by Dinh and Thai [120] to signed graph clustering in general. The formulation proposed in [121] removes transitivity constraints associated with two of more edges with negative weight, which reduces the number of constraints to $O(n(n^2 - m^-))$ constraints.

Miyauchi et al. [122] obtain an improved number $O(nm^+)$ of constraints by running the ILP from [121] on a new signed graph where artificial edges are introduced between pairs of unconnected nodes from the input graph. In particular, these artificial edges have a negative weight which is small enough to ensure that every optimal solution in this new graph corresponds to an optimal solution in the input graph. Recently Koshimura et al. [138] proposed a further improvement for the ILP by Miyauchi et al. [122] which removes around 50% of its constraints. Although there have been many improved ILP formulations for signed graph clustering in the last years, this approach does not scale well for large instances (in particular for the size of the instances considered in this paper). In particular, signed graphs with more

than a few thousand nodes become intractable in practice, which is the reason why most instances in the works cited above have no more than a few hundred nodes.

Chapter 4

Streaming Algorithms

In this chapter, we present our contributions in the field of (buffered) streaming (hyper)graph decomposition. Firstly, we introduce the algorithm HeiStream that solves the graph partitioning problem using the buffered streaming model. The algorithm operates by partitioning batches or buffers of nodes with a multilevel algorithm after building a model for each loaded buffer. HeiStream is shown to produce solutions of better quality than previous state-of-the-art algorithms while also being faster when partitioning into large numbers of blocks. Secondly, we propose Online Recursive Multi-Section, a streaming algorithm for the process mapping problem that uses recursive multi-sections on-the-fly. Our approach is the first in the literature designed specifically for process mapping, and its versatility allows it to also solve the graph partitioning problem. Experimental results show that our recursive multisection on-the-fly algorithm produces significantly better solutions in terms of total communication cost while being up to multiple orders of magnitude faster than the state-of-the-art. Thirdly, we introduce FREIGHT, a streaming algorithm for hypergraph partitioning that extends the graph-based algorithm Fennel. Our proposed implementation for FREIGHT is highly efficient, leading to significantly faster computation times and better cut-net and connectivity measures than previous state-of-the-art algorithms. Lastly, we present an experimental comparison of our three algorithms for the graph partitioning problem.

References. This chapter is based on [10] and [11], which is joint work with Christian Schulz, and on [12], which is joint work with Kamal Eyubov and Christian Schulz. The experimental comparison of our three streaming algorithms is new.

4.1 Buffered Streaming Graph Partitioning

In this section, we start to fill the gap currently observed between in-memory and streaming algorithms for graph partitioning algorithms. We propose HeiStream, an algorithm that can

Algorithm 4.1 Structure of HeiStream

```

while  $G$  has not been completely streamed do
    Load batch of nodes
    Build model  $B$ 
    Run multilevel partitioning on model  $B$ 
    Assign nodes of batch to permanent blocks

```

compute significantly better partitions of huge graphs than the currently available streaming algorithms while using a single machine without a lot of memory. We adopt the buffered streaming model which allows a buffer of nodes to be received and stored before making assignment decisions. Our algorithm is carefully engineered to produce partitions of improved quality by using a sophisticated multilevel scheme on a compressed model of the buffer *and* the already assigned nodes. Our multilevel algorithm optimizes for the same objective as the previous state-of-the-art Fennel. However, due to the multilevel scheme used on the compressed model, our local search algorithms have a global view on the optimization problem and hence compute better solutions overall. Lastly, using the multilevel scheme reduces the time complexity from $O(nk + m)$ of Fennel to $O(n + m)$, where k is the number of blocks a graph has to be partitioned in. To this end, experiments indicate that our algorithm can partition huge networks on machines with small memory while computing *better* solutions than the previous state-of-the-art in the streaming setting. At the same time our algorithm is *faster* than the previous state-of-the-art for larger values of blocks k .

4.1.1 HeiStream

4.1.1.1 Overall Structure

We now explain the overall structure of HeiStream. We slide through the streamed graph G by repeating the following successive operations until all the nodes of G are assigned to blocks. First, we load a batch containing δ nodes alongside with their adjacency lists. Second, we build a model B to be partitioned. This model represents the already partitioned nodes as well as the nodes of the current batch. Third, we partition B with a multilevel partitioning algorithm to optimize for the Fennel objective function. And finally, we permanently assign the nodes from the current batch to blocks. Algorithm 4.1 summarizes the general structure of HeiStream and Figure 4.1 shows the detailed structure of HeiStream.

4.1.1.2 Model Construction

We build two different models, which yield a running time-quality trade-off. We start by describing the basic model and then extend this later.

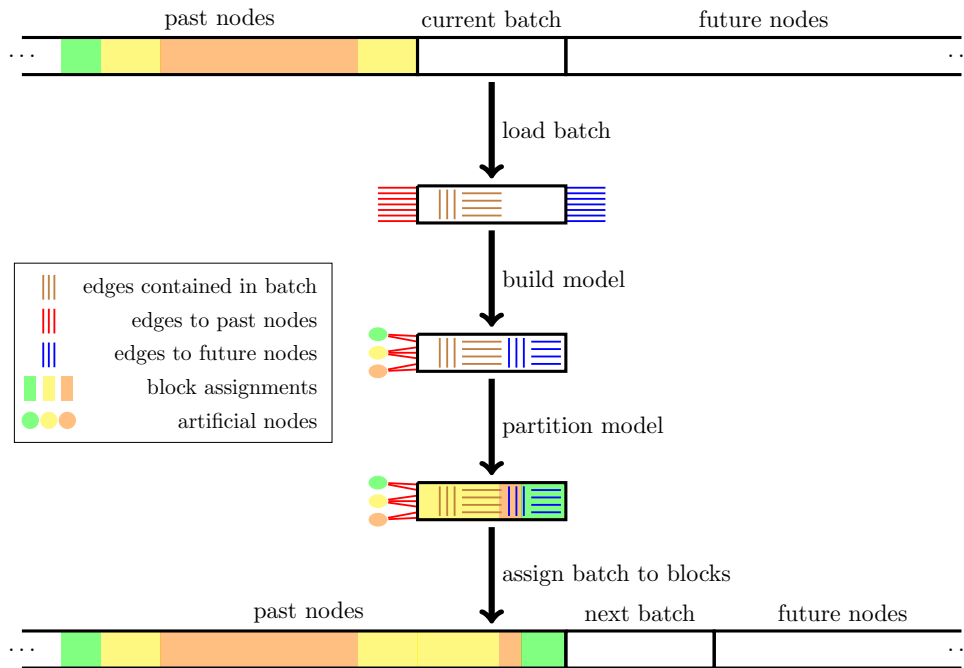


Figure 4.1. Detailed structure of HeiStream. The algorithm starts by loading a batch of nodes alongside with their edges. Next, it builds a meaningful model based on the loaded nodes and edges. This model is then partitioned using a multilevel algorithm. Based on the partition of the model, the nodes of the current batch are permanently assigned to blocks. This whole process is repeated for the next batch until the whole graph is partitioned.

When a batch is loaded, we build the *basic model* B as follows. We initialize B as the subgraph of G induced by the nodes of the current batch. If the current batch is not the first one, we add k artificial nodes to the model. These represent the k preliminary blocks in their current state, i.e., filled with the nodes from the previous batches, which were already assigned. The weight of an artificial node i is set to the weight of block V_i . A node of the current batch is connected to an artificial node i if it has a neighbor from the previous batch that has been assigned to block V_i . If this creates parallel edges, we replace them by a single edge with its weight set to the sum of the weight of the parallel edges. Note that the basic model does ignore edges towards nodes that will be streamed in future batches, i.e., batches that have not been streamed yet.

Our *extended model* incorporates edges towards nodes from future, not yet loaded, batches – if the stream contains such edges. We call edges towards nodes of future batches ghost edges and the corresponding endpoint in the future batch ghost node. Ghost nodes and ghost edges provide partial information about parts of G that have not yet been streamed. Hence, representing them in the model B enhances the partitioning process. Note though that sim-

ply inserting all ghost nodes and edges can overload memory in case there is an excessive number of them. Thus our approach consists of randomly contracting the ghost nodes with one of their neighboring nodes from the current batch. Note that this contraction increments the weight of a node within our model and ensures that if there are more than one node from the current batch connected to the same future node, then there will be edges between those nodes in our model. Also note that the contraction ensures that the number of nodes in all models throughout the batched streaming process is constant. This prevents memory from being overloaded and makes it unnecessary to reallocate memory for B between successive batches. In order to give a lower priority to ghost edges in comparison to the other edges, we divide the weight of each ghost edge by 2 in our model. The construction of the extended model is conceptually illustrated in Figure 4.1.

4.1.1.3 Multilevel Weighted Fennel

Our approach to partition the model B is a multilevel version of the algorithm Fennel. Recall that the multilevel scheme consists of three successive phases: coarsening, initial partitioning, and uncoarsening, as depicted in Figure 2.2.

It should be noted that the artificial nodes present in our model may become very heavy. As an acknowledgement of their representation of permanent block assignments, any contraction or block change of artificial nodes is prohibited. Consequently, these nodes require special handling in our multilevel algorithm. Furthermore, note that the original formulation of Fennel only works for unweighted graphs [3]. However, our model B has nodes and edges that are weighted – as a result of connections to artificial nodes and future nodes that may be contracted into the model, as well as due to the generation of a series of weighted graphs through the multilevel scheme. Hence, the Fennel algorithm requires some adaptation to be utilized in HeiStream. We introduce a generalization of the Fennel algorithm for weighted graphs that can be directly employed in a multilevel algorithm. Within this section, we will present this generalized Fennel objective and provide a detailed explanation of our multilevel algorithm for partitioning the model.

Generalized Fennel. As a necessary feature, our generalization of the Fennel gain function must ensure that the gain of a node at a coarser level equals the sum of the gains of the nodes that it represents at the finer levels. This way the algorithm gets a global view onto the optimization problem on the coarser levels and a very local view on the finer levels. Moreover, it is ensured that on each level of the hierarchy the algorithm works towards optimizing the given objective.

Our generalization of the gain function of Fennel is as follows. Let u be the node that should be assigned to a block. Our generalized Fennel assigns u to a block i that maximizes

Equation (4.1) such that $f(c(V_i)) = \alpha * \gamma * c(V_i)^{\gamma-1}$. Note that this is a direct generalization of the unweighted case. First, if the graph does not have edge weights, then the first part of the equation becomes $|V_i \cap N(u)|$ which is the first part of the Fennel objective. Second, if the graph also does not have node weights, then the second part of the equation is the same as the second part of the equation in the Fennel objective. Moreover, observe that the penalty term $f(c(V_i))$ in our objective is multiplied by $c(u)$. This is done to have the property stated above and formalized in Theorem 4.1. Finally, we keep the original value of the parameters α and γ in order to keep consistency.

$$\sum_{v \in V_i \cap N(u)} \omega(u, v) - c(u)f(c(V_i)) \quad (4.1)$$

Theorem 4.1. *If a set of nodes $S \subseteq V$ is contracted into a node w , the generalized Fennel gain function of w is equal to the sum of the generalized Fennel gain functions of the nodes in S .*

Proof. On the one hand, the generalized Fennel gain of assigning a node w to a block i is computed as:

$$\sum_{v \in V_i \cap N(w)} \omega(w, v) - c(w)f(c(V_i)) \quad (4.2)$$

On the other hand, the sum of the generalized Fennel gains of assigning all nodes in S to a block i consists of:

$$\sum_{u \in S} \sum_{v \in V_i \cap N(u)} \omega(u, v) - \sum_{u \in S} c(u)f(c(V_i)) \quad (4.3)$$

Since the factor $f(c(V_i))$ is identical in (4.2) and (4.3), the two of them are equivalent if equalities (4.4) and (4.5) hold:

$$\sum_{v \in V_i \cap N(w)} \omega(w, v) := \sum_{u \in S} \sum_{v \in V_i \cap N(u)} \omega(u, v) \quad (4.4)$$

$$c(w) := \sum_{u \in S} c(u) \quad (4.5)$$

But equalities (4.4) and (4.5) are trivially true as a property of the contraction process. \square

Multilevel Fennel. In principle, our model B could be partitioned by any multilevel partitioning algorithm which allows fixed nodes. Nevertheless, our preliminary tests using an adaptation of KaHIP [37] generated poor edge-cut. This is the case because internal memory algorithms such as KaHIP and Metis [139] are designed to minimize edge-cut as much as possible while ensuring a low balancing constraint. This leads to two possibilities, both of which are bad for the overall edge-cut within the buffered streaming model: (i) using a large balancing constraint for the first buffers, the internal memory partitioner will tend to assign

all their nodes to a single block, (ii) using a low balancing constraint for each batch, the overall partitioning problem will be aggressively over-restricted. Hence, we design a multilevel algorithm based on the Fennel function, which is a function that successively deals with the trade-off between edge-cut and imbalance associated with graph streaming.

We now explain our multilevel algorithm to partition the model B . Our *coarsening* phase is based on an adapted version of the size-constraint label propagation approach [38]. To be self-contained, we shortly outline the coarsening approach and then show how to modify it to be able to handle artificial nodes. To compute a graph hierarchy, the algorithm computes a size-constrained clustering on each level and contract that to obtain the next level. The clustering is contracted by replacing each of its clusters by a single node (as exemplified in Figure 2.3), and the process is repeated recursively until the graph becomes small enough. This hierarchy is then used by the algorithm. Due to the way we define contraction, it ensures that a partition of a coarse graph corresponds to a partition of all the finer graphs in the hierarchy with the same edge-cut and balance. Note that cluster contraction is an aggressive coarsening strategy. In contrast to matching-based approaches, it enables us to drastically shrink the size of irregular networks. The intuition behind this technique is that a clustering of the graph (one hopes) contains many edges running inside the clusters and only a few edges running between clusters, which is favorable for the edge cut objective.

The algorithm to compute clusters is based on *label propagation* [140] and avoids large clusters by using a *size constraint*, as described in [38]. For a graph with n nodes and m edges, one round of size-constrained label propagation can be implemented to run in $O(n + m)$ time. Initially, each node is in its own cluster/block, i. e., the initial block ID of a node is set to its node ID. The algorithm then works in rounds. In each round, all the nodes of the graph are traversed. When a node v is visited, it is *moved* to the block that has the strongest connection to v , i. e., it is moved to the cluster V_i that maximizes $\omega(\{(v, u) \mid u \in N(v) \cap V_i\})$. Ties are broken randomly. We perform at most L rounds, where L is a tuning parameter.

In HeiStream, we have to ensure that two artificial nodes are not contracted together since each of them should remain in its previously assigned block. We achieve this by ignoring artificial nodes and artificial edges during the label propagation, i. e., artificial nodes cannot not change their label and nodes from the batch can not change their label to become a label of an artificial node. As a consequence, artificial nodes are not contracted during coarsening. Overall, we repeat the process of computing a size-constrained clustering and contracting it, recursively. As soon as the graph is small enough, i. e., it has fewer nodes than an $O(\max(|B|/k, k))$ threshold, it is initially partitioned by an initial partitioning algorithm. More precisely, we use the threshold $\max(\frac{|B|}{2xk}, xk)$, in which x is a tuning parameter. Note that, for large enough buffer sizes, this threshold will be $O(|B|/k)$.

When the coarsening phase ends, we run an *initial partitioning algorithm* to compute an initial k -partition for the coarsest version of B . That means that all nodes other than the artificial nodes, which are already assigned, will be assigned to blocks. To assign the nodes, we run our generalized Fennel algorithm with explicit balancing constraint L_{\max} , i. e., the weight of no block will exceed L_{\max} . To be precise, a node u will be assigned to a block i that maximizes $\sum_{v \in V_i \cap N(u)} \omega(u, v) - c(u)f(c(V_i))$, such that $f(c(V_i)) = \alpha * \gamma * c(V_i)^{\gamma-1}$ and $c(V_i \cup u) \leq L_{\max}$. Note that the algorithm at this point considers all possible blocks $i \in \{1, \dots, k\}$ and hence has complexity proportional to k . However, as the coarsest graph has $O(|B|/k)$ nodes, overall the initial partitioning needs time which is linear in the size of the input model. When initial partitioning is done, we transfer the current solution to the next finer level by assigning each node of the finer level to the block of its coarse representative. At each level of the hierarchy, we apply a *local search* algorithm. Our local search algorithm is the same size-constraint label propagation algorithm we used in the contraction phase but with a different objective function. Namely, when visiting a non-artificial node, we remove it from its current block and then we assign it to the neighboring block which maximizes the generalized Fennel gain function defined in Equation (4.1). Note that, in contrast to the initial partitioning, only blocks of adjacent nodes are considered here. Hence, one round of the algorithm can still be implemented to run in linear time in the size the current level. As in the coarsening phase, artificial nodes cannot be moved between blocks. Differently though, we do not exclude the artificial nodes from the label propagation here. This is the case because the artificial nodes and their edges are used to compute the generalized Fennel gain function of the other nodes. As in the initial partitioning, we use the explicit size constraint L_{\max} of G . As a side note, we also tried to use high-quality offline algorithms as initial partitioning algorithms, however, in preliminary experiments this results in very unbalanced blocks (even with adaptively configured balance constraints) and overall in reduced quality throughout the process. Hence, we did not consider this further.

Assuming geometrically shrinking graphs throughout the hierarchy and assuming that the buffer size δ is larger than the number of blocks k , then the overall running time to partition a batch is linear in the size of the batch. This is due to the fact that the overall running time of coarsening and local search sums up to be linear in the size of the batch, while the overall running time of the initial partitioning depends linearly on the size of the input model. Summing this up over all batches yields overall linear running time $O(n + m)$.

4.1.1.4 Restreaming

We now extend HeiStream to operate in a restreaming setting. During restreaming, the overall structure of the algorithm is roughly the same. Nevertheless we need to implement some adaptations which we explain in this section. The first adaptations concern model construc-

tion. Recall that the nodes from the current batch are already assigned to blocks during the previous pass of the input. We explicitly assign these nodes to their respective blocks in B . Furthermore, ghost nodes and edges are not needed to construct B . This is the case since all nodes from future batches are already known and assigned to blocks, i. e., these nodes will be represented in the artificial nodes. More precisely, we adapt the artificial nodes to represent the nodes from all batches except the current one. Since a partition of the graph is already given, we do not allow the contraction of cut edges during restreaming in the coarsening phase of our multilevel scheme. That means that clusters are only allowed to grow inside blocks. As a consequence, we can directly use the partition computed in the previous pass as initial partitioning for B so we do not need to run an initial partitioning algorithm.

4.1.1.5 Implementation Details

Our implementation of B is based on an adjacency array and consecutive node IDs. We reserve the first δ IDs for the nodes from the current batch, which keep their global order. This means that, when we process the i^{th} batch, nodes IDs can be easily converted from our model B to G and the other way around by respectively summing or subtracting $(i - 1) * \delta$ on their ID. Similarly, we reserve the last k IDs of B for the artificial nodes and keep their relative order for all batches. Note that this configuration separates *mutable* nodes (nodes from current batch) and *immutable* nodes (artificial nodes). This allows us to efficiently control which nodes are allowed to move during coarsening, initial partitioning, and local search. We keep an array of size n store the permanent block assignment of the nodes of G . To improve running time, we use approximate computation of powering in our Fennel function.

4.1.2 Experimental Evaluation

Setup. We performed the implementation of HeiStream and competing algorithms inside the KaHIP framework (using C++) and compiled them using gcc 9.3 with full optimization turned on (-O3 flag). Most of our experiments were run on a single core of Machine A. The only exceptions are the experiments with huge graphs, which were run on a single core of Machine B. When using machine A, we stream the input directly from the internal memory, and when using machine B, that only has 16 GB of main memory, we stream the input from the hard disk.

Baselines. We compare HeiStream against various state-of-the-art algorithms for streaming graph partitioning. Since no official versions of the one-pass streaming and restreaming algorithms are available in public repositories, we implemented them in our framework.

Our implementations of these algorithms reproduce the results presented in the respective papers and are optimized for running time as much as possible. To this end, we implemented Hashing, LDG, Fennel, and ReFennel. Multilevel LDG [5] is also not publicly available. We sent a message to the authors requesting an executable version of their algorithm for our tests but we have not receive any response. Hence, we compare HeiStream against Multilevel LDG based on the results explicitly reported in [5]. We have used two machines: Machine A and Machine B.


Instances. In this section, we use three disjoint sets of graphs. Basic properties of the graphs under consideration can be found in Table 2.1. The *tuning* set is used for the parameter study experiments, the *test* set is used for the comparisons against the state-of-the-art, and the set of *huge graphs* is used for special larger scale tests. In any case, when streaming the graphs we use the natural given order of the nodes. We use $k \in \{2, 3, \dots, 128\}$ for most experiments. We allow an imbalance of 3% for all experiments (and all algorithms). All partitions computed by all algorithms have been balanced.

Methodology. Depending on the focus of the experiment, we measure running time and/or edge-cut. Unless explicitly mentioned otherwise, we average all results of each algorithm grouped by k .


4.1.2.1 Parameter Study

We now present experiments to tune HeiStream and explore its behavior. We do this on the tuning set of our instance set. In our strategy, each experiment focuses on a single parameter of the algorithm while all the other ones are kept invariable. We start with a baseline configuration consisting of the following parameters: 5 rounds in the coarsening label propagation, 1 round in the local search label propagation, and $x = 64$ in the expression of the coarsest model size. After each tuning experiment, we update the baseline to integrate the best found parameter. Unless explicitly mentioned otherwise, we run the experiments of this section over all tuning graphs from Table 2.1 based on the *extended* model construction, i. e., including ghost nodes and edges, for a buffer size of 32 768.


Tuning. We begin by evaluating how the number of label propagation rounds during local search affects running time and solution quality. In particular, we run configurations of HeiStream with 1, 5, and 25 rounds and report results in Figures 4.2a and 4.2b. Observe that the results of the baseline have considerably lower solution quality than the other ones overall. On the other hand, the results of the configurations with 5 and 25 rounds differ slightly to each other. On average, they respectively improve solution quality 3.6% and




(a) Quality improvement plot for label propagation rounds during uncoarsening.



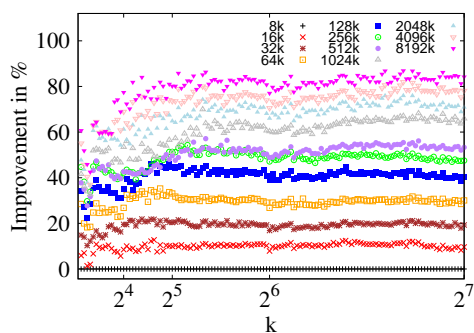
(b) Running time ratio plot for label propagation rounds during uncoarsening.




(c) Quality improvement plot for parameter x from expression $\max(|B|/2xk, xk)$.



(d) Quality improvement plot for restreaming.



(e) Quality improvement plot for buffer size.



(f) Running time ratio plot for buffer size.

Figure 4.2. Results for tuning and exploration experiments. Higher is better for quality improvement plots. Lower is better for running time ratio plots.

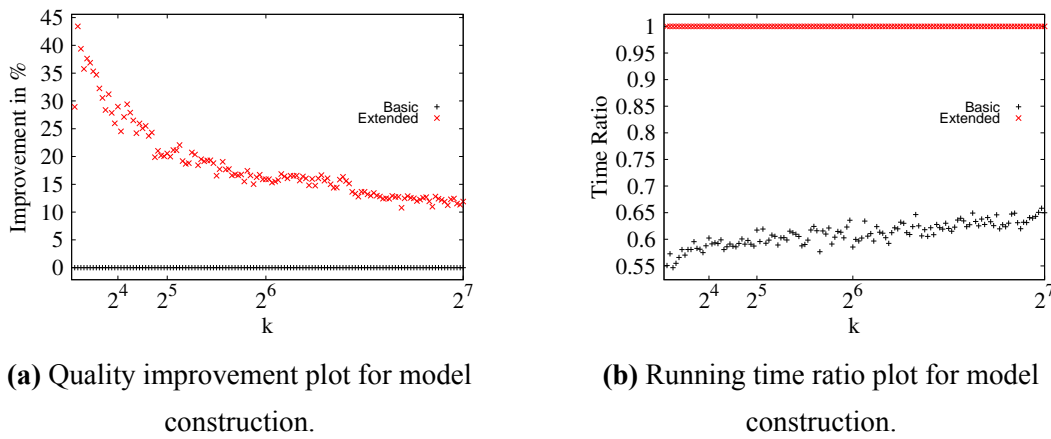


Figure 4.3. Results for tuning and exploration experiments. Higher is better for quality improvement plots. Lower is better for running time ratio plots.

3.7% over the baseline. Regarding running time, they respectively increase 2.2% and 3.4% on average over the baseline. Since the variation of quality for these two configurations is not significant, we decided to integrate the fastest one among them in the algorithm, namely the 5-round configuration.

Next we look at the parameter x associated with the expression $\max(|B|/2xk, xk)$, which determines the size of the coarsest model. We run experiments for $x = 2^i$, with $i \in \{0, \dots, 6\}$, and report results in Figure 4.2c. We omit running time charts for this experiment since the tested configurations present comparable behavior in this regard. Figure 4.2c shows that the baseline presents the overall worst solution quality while $x = 2$ and $x = 4$ present the overall best solution quality. Averaging over all instances, these two latter configurations produce results respectively 10.3% and 11.2% better than the baseline. In light of that, we decide in favor of $x = 4$ to compose HeiStream.

Exploration. We start the exploration of open parameters by investigating how the buffer size affects solution quality and running time. We use as baseline a buffer of 8 192 nodes and successively double its capacity until any graph from the tuning set in Table 2.1 fits in a single buffer. We plot our results in Figures 4.2e and 4.2f. Note that solution quality and running time increase regularly as the buffer size becomes larger. This behavior occurs because larger buffers enable more comprehensive and complex graph structures to be exploited by our multilevel algorithm. As a consequence, there is a trade-off between solution quality and resource consumption. In other words, we can improve partitioning quality at the cost of considerable extra memory and slightly more running time. Otherwise, we can save memory as much as possible and get a faster partitioning process at the cost of lowering solution quality. In practice, it means that HeiStream can be adjusted to produce partitions

as refined as possible with the resources available in a specific system. For the extreme case of a single-node buffer, HeiStream behaves exactly as Fennel, while it behaves as an internal memory partitioning algorithm for the opposite extreme case.

Next, we compare the effect of using the *extended* model, which incorporates ghost nodes, over the *basic* model, which ignores ghost nodes. Figures 4.3a and 4.3b displays the results. The results show that the extended model provides improved quality over the basic model, with an 18.3% improvement on average. This happens because the presence of ghost nodes and edges expands the perspective of the partitioning algorithm to future batches. This has a similar effect to increasing the size of the buffer, but at no considerable extra memory cost. Regarding running time, the results show that the extended model is consistently slower than the basic model for all values of k . Averaging over all instances, the extended model costs 63.9% more running time than the basic model. This increase in running time is explained by the higher number of edges to be processed when ghost nodes are incorporated in the model. As a practical conclusion from the experiment, the extended model can be used for better overall partitions with no significant extra memory but at the cost of extra running time. Otherwise, the basic model can be used for a consistently faster execution at the cost of a lower solution quality.

Finally, we test to what extent solution quality can be improved by restreaming HeiStream multiple times. We investigate this by restreaming each input graph 10 times. We collect results after each pass and plot in Figure 4.2d. The first restream generates a considerable quality jump, with an improvement over the baseline of 24.6% on average. Each following pass has a positive impact on solution quality, which converges to be a 40.9%-improvement on average over the baseline after the last pass. On the other hand, the running time has a roughly linear increase for each pass over the graph. In practice, this adds another degree of freedom to configure HeiStream for the needs of real systems.

4.1.2.2 State-of-the-Art

In this section, we show experiments in which we compare HeiStream against the current state-of-the-art algorithms. Unless mentioned otherwise, these experiments involve all the graphs from the Test Set group in Table 2.1 and focus on two particular configurations of HeiStream, which we refer to as HeiStream(32k) and HeiStream(Int.). The first configuration is based on batches of size 32 768, while the second one has enough batch capacity to operate as an internal memory algorithm – both configurations perform a single pass over the input using the extended model. We also present results for the 2-pass restreaming version of HeiStream(32k), which we refer to as 2-ReHeiStream(32k).

Internal memory algorithms such as Metis [139] and KaHIP [37] are beyond the scope of this thesis since it is common knowledge that internal memory algorithms are better than

Graph	Cut Edges (%)							
	HS(Int.)	MLDG(Int.)	2-ReHS(32k)	2-ReFennel	HS(32k)	Fennel	LDG	Hashing
Dubcova1	13.68	14.26	12.92	29.19	13.68	33.99	33.96	95.62
hcircuit	2.73	17.75	2.04	21.86	2.53	28.97	28.97	90.75
coAuthorsDBLP	15.99	24.82	16.90	24.28	17.80	27.12	27.12	94.80
Web-NotreDame	5.85	11.01	6.15	12.58	9.20	19.52	19.56	95.97
Dblp-2010	11.31	18.52	12.20	22.93	13.42	28.82	28.80	92.49
ML_Laplace	7.93	13.44	7.62	7.82	8.36	7.92	5.77	96.37
coPapersCiteseer	8.23	11.22	8.35	9.63	10.29	12.88	12.27	96.52
coPapersDBLP	14.51	19.29	15.12	16.47	18.33	20.65	20.22	96.39
Amazon-2008	10.09	19.01	12.77	28.92	15.56	37.07	37.07	94.68
eu-2005	11.14	14.57	14.82	25.53	18.64	35.88	31.96	96.44
Web-Google	1.62	9.66	3.96	18.04	9.48	30.64	30.64	96.87
ca-hollywood-2009	35.34	32.51	37.86	41.34	42.36	44.54	45.25	96.62
Flan_1565	7.69	9.36	7.30	10.26	8.12	10.59	6.70	96.61
Ljournal-2008	29.12	34.76	33.58	43.23	38.58	51.43	51.36	96.07
HV15R	14.09	11.48	15.38	15.05	17.48	16.39	15.49	96.84
Bump_2911	8.66	11.73	7.65	8.61	8.23	8.65	8.30	96.19
FullChip	38.20	48.16	42.24	57.39	45.71	61.93	64.23	95.06
patents	15.57	29.12	41.75	52.60	60.56	70.98	70.98	96.88
Cit-Patents	15.75	28.65	39.67	51.62	60.76	72.16	72.16	96.88
Soc-LiveJournal1	29.72	35.69	29.39	34.00	35.62	39.03	45.62	96.66
circuit5M	40.02	34.60	39.20	75.45	41.00	78.42	78.47	96.87
del21	1.38	-	5.41	33.52	8.53	40.21	40.21	93.39
rgg21	1.53	-	1.34	4.09	1.52	5.02	4.88	96.89
soc-orkut-dir	37.86	-	43.19	47.22	54.76	55.85	60.27	96.85
italy-osm	0.13	-	1.19	4.65	1.34	4.80	4.80	78.11
great-britain-osm	0.16	-	1.43	7.18	1.63	7.34	7.34	79.94

Table 4.1. Edge-cut results against competitors for $k = 32$. Internal memory algorithms are on the 2 left columns and streaming algorithms are on the 5 right columns. We refer to setups of HeiStream with specific buffer sizes as HS(Xk), in which each buffer contains $X \times 1024$ nodes. HS(Int.) uses a buffer size of n . We bold the best result for each graph for internal memory approaches and streaming approaches. The results for Multilevel LDG for the five bottom graphs are missing, as the graphs were not part of their benchmark set. Lower is better.

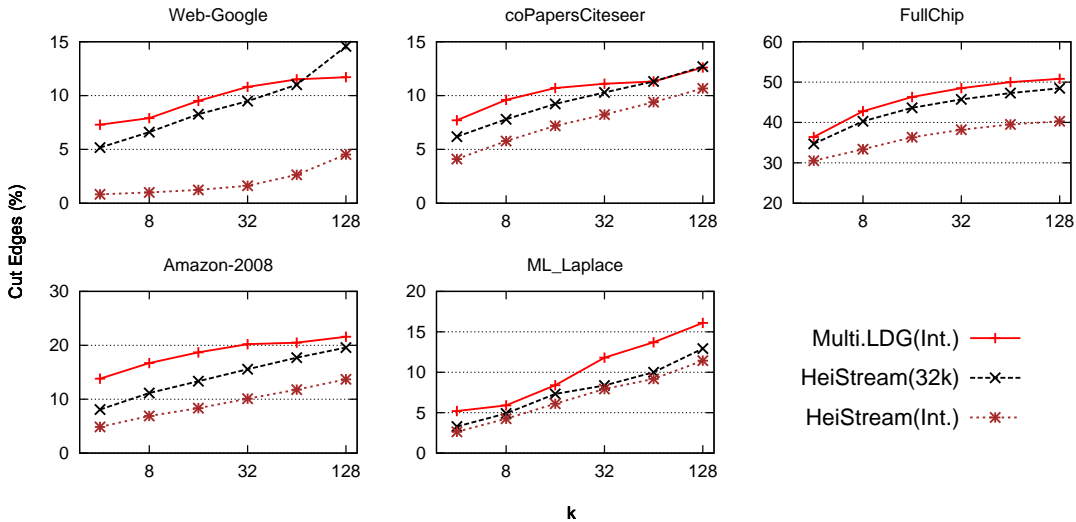



Figure 4.4. Comparison of HeiStream against Multilevel LDG for buffer containing the whole graph.


streaming algorithms regarding partition quality if the instances fits in the memory of a machine. For the sake of reproducibility, we ran Metis and the fast social version of KaHIP over our Test Set group of instances for $k \in \{2, 4, 8, 16, 32, 64, 128\}$. On average, Fennel cuts a factor 7.5 more edges than Metis and the fast social version of KaHIP, while HeiStream(Int.) cuts a factor 2.2 more edges than both. Furthermore, Fennel is respectively a factor 2.2 and a factor 6.0 faster than Metis and the fast social version of KaHIP, while HeiStream(Int.) is 56.3% slower than Metis and 72.9% faster than the fast social version of KaHIP.

Results. We now present a detailed comparison of HeiStream (HS) against the state-of-the-art. In the results, we refer to the internal memory version of Multilevel LDG as MLDG(Int.). Moreover, we refer to the restreaming version of HeiStream and Fennel that passes over the graph 2 times as 2-ReHS and 2-ReFennel respectively. First, we focus on $k = 32$ and later on choose a much wider range for the number of blocks. Table 4.1 shows the percentage of edges cut in the partitions generated by each algorithm for the graphs in the Test Set for $k = 32$. HeiStream(Int.), 2-ReHeiStream(32k), and HeiStream(32k) outperform all the other competitors for the majority of instances. First, all of them outperform Hashing for all graphs and LDG for 23 out of the 26 graphs. Next, they also outperform Fennel in 24 instances. The algorithm 2-ReFennel is outperformed by HeiStream(Int.), 2-ReHeiStream(32k), and HeiStream(32k) in 24, 25, and 17 instances, respectively. Considering only the 21 instances for which there are results reported for Multi.LDG(Int.) in literature,



(a) Quality improvement plot over Fennel.

(b) Relative running time plot over Fennel.



(c) Quality performance profile.

(d) Running time performance profile.

Figure 4.5. Results for comparison against state-of-the-art one-pass (re)streaming algorithms using performance profiles. Higher is better for quality improvement plots.

the algorithms HeiStream(Int.), 2-ReHeiStream(32k), and HeiStream(32k) compute better partitions for 18, 15, and 14 instances respectively.

For a closer comparison against Multi.LDG(Int.), we present Figure 4.4. We plot edge-cut for Multi.LDG(Int.) based on results graphically reported in [5]. In this figure, we show of HeiStream(Int.) and HeiStream(32k) for 5 particular graphs with $k = 4, 8, 16, 32, 64, 128$. For all these instances, HeiStream(Int.) outperforms Multi.LDG(Int.) by a considerable margin. Note that HeiStream(32k) outperforms the *internal memory* version of Multilevel LDG for the majority of instances. We omit additional comparisons against buffered versions of Multilevel LDG, since they provide lower quality than the internal memory version. We ran wider experiments over our whole Test Set for 127 different values of k . Figures 4.5a and 4.5b show a quality improvement plot over Fennel and a relative running time plot, respectively. Figures 4.5c and 4.5d show performance profiles for solution quality and running time, respectively. Observe that HeiStream(Int.) produces solutions with highest quality overall. In particular, it produces partitions with smallest edge-cut for almost 86.9% of the instances and improves solution quality over Fennel 195.0% on average. We now provide some re-




Figure 4.6. Quality performance profile of 2-ReHeiStream(32k) against 2-ReFennel.

sults in which we exclude HeiStream(Int.), since it has access to the whole graph. The best algorithm is HeiStream(32k), which produces the best solution quality for 63.3% of the instances and improves solution quality over Fennel 75.9% on average. It is followed by 2-ReFennel, which is the best tested algorithm from the previous state-of-the-art. In particular, it computes the best partition for 26.8% of the instances and improves on average 19.2% over Fennel. LDG and Fennel come next. Particularly, LDG finds the best partition for 9.8% of the instances and improves on average 2.4% over Fennel. Fennel does not find the best partition for any instance. Finally, Hashing produces the worst solutions, with 72.5% worse quality than Fennel on average. Regarding running time, Hashing is the fastest one for all instances, which is expected since it is the only one with time complexity $O(n)$. The second fastest one is LDG, whose running time is a factor 9.6 higher than Hashing on average. Fennel, HeiStream(32k) and HeiStream(Int.) come next with factors 32.4, 41.4 and 56.2 slower than Hashing on average, respectively. ReFennel is the slowest algorithm of the test, being a factor 64.61 slower than Hashing on average. In a direct comparison, HeiStream(32k) and HeiStream(Int.) are respectively 27.7% and 73.2% slower than Fennel on average. Note that both configurations of HeiStream are faster than Fennel for larger values of k , which is consistent with the fact the running time of HeiStream is $O(n + m)$ while the running time of Fennel is $O(nk + m)$.

Restreaming. Now we directly evaluate the restreaming version of HeiStream. As shown in Table 4.1, 2-ReHeiStream(32k) computes solutions with smaller edge-cut than 2-ReFennel for almost all our Test Set when $k = 32$. Figure 4.6 shows that this is the case for $k \in \{2, \dots, 128\}$. In particular, 2-ReHeiStream(32k) produces the best solution for 98.3% of the instances and improves solution quality over 2-ReFennel by 79.6% on average. The average improvement of 2-ReHeiStream(32k) over 2-ReFennel is comparable to the average improvement of HeiStream(32k) over Fennel. Nevertheless, the percentage of instances for




Figure 4.7. Visualization of partitions generated by the algorithms Hashing, Fennel, HeiStream($1k$) and the fast social version of KaHIP for the graph rgg15, which has 32 768 nodes and 160 240 edges.

which 2-ReHeiStream(32k) outperforms 2-ReFennel is almost 100%, which means that 2-ReHeiStream(32k) almost dominates 2-ReFennel with respect to edge-cut. Note that this is considerably higher than the percentage of instances for which HeiStream(32k) outperforms Fennel (63.3%).

Visualization. As shown, the edge-cut of partitions produced by HeiStream is on average lower than the edge-cut of partitions produced by its competitor streaming algorithms. We shortly look at some visualizations in order to concretely understand why this happens. In Figure 4.7, we show a visual comparison of some partition layouts generated by Hashing, Fennel, HeiStream and the fast social version of KaHIP for the graph rgg15. Since this graph has only 32 768 nodes, we use a buffer size of 1 024 nodes for HeiStream in order to partition the graph over multiple successive batches. There is a leap of partitioning quality from Hashing to Fennel, i. e., well-delimited clusters associated to a same block can be identified

in the partitions generated by Fennel but not in partitions generated by Hashing. Similarly there is a leap of partitioning quality from Fennel to KaHIP, i. e., a block generated by Fennel consists of multiple small clusters that are not mutually connected while a block generated by KaHIP usually consists of a single connected cluster. Note that the partitions produced by HeiStream have intermediary characteristics between the partitions generated by Fennel and KaHIP. More specifically, a block generated by HeiStream consists of fewer and larger clusters than a block generated by Fennel but not as few and as large clusters as those generated by KaHIP. This behavior is a direct consequence of the more or less global view provided by the distinct computational models used by these three algorithms.

4.1.2.3 Huge Graphs

We now switch to the main use case of streaming algorithms: computing high-quality partitions for huge graphs on small machines. The experiments in this section are based on the *huge graphs* listed in Table 2.1 and are run on the relatively small Machine B. Namely, we ran experiments for $k = \{8, 16, 32, 64, 128, 256\}$ and we did not repeat each test multiple times with different seeds as in previous experiments. We also ran Metis and KaHIP on those graphs on this machine, but they fail on all instances since they require more memory than the machine has. For all instances, HeiStream performs a single pass over the input based on the extended model construction. We refer to setups of HeiStream with specific buffer sizes as HeiStream(Xk), in which a buffer contains $X \times 1024$ nodes. In Table 4.2, we report detailed per-instance results with large buffer sizes able to run on Machine B. We exclude from Table 4.2 the IO delay to load the input graph from the disk, since it depends on the disk and is roughly the same independently of the used partitioning algorithm. For completeness, we report this delay (in seconds) for the huge graphs listed in Table 2.1 following their respective order: 131.3, 203.2, 313.2, 294.0, 164.9, 186.1, 340.7, 551.5.

The results show that HeiStream outperforms all the competitor algorithms regarding solution quality for most instances. Notably, HeiStream computes partitions with considerably lower edge-cut in comparison to the one-pass algorithms for 4 out of the tested graphs: uk-2005, sk-2005, uk-2007-05, and RHG1. For the social networks soc-friendster and twitter7, HeiStream is the best for all instances, but the improvement over Fennel and LDG is not so large as in the other instances. One outlier can be seen on the network RHG2. While HeiStream produces fairly small edge-cut values, which are all below 0.7%, Fennel does outperform it and LDG improves solution quality even further on this instance. Furthermore, note that the running time of Fennel increases with increasing k to the point in which it becomes higher than the running time of HeiStream for 5 out of the 8 huge graphs tested.

Graph	k	HeiStream(Xk)			Fennel		LDG		Hashing	
		X	CE(%)	RT(s)	CE(%)	RT(s)	CE(%)	RT(s)	CE(%)	RT(s)
uk-2005	8	1024	4.03	290.23	19.93	37.19	19.97	19.36	73.70	3.28
	16	1024	6.01	300.04	22.76	58.05	22.72	22.58	78.86	3.38
	32	1024	7.65	310.72	25.24	98.78	25.19	29.19	81.39	3.30
	64	1024	8.99	322.14	26.88	183.15	26.81	42.90	82.60	3.28
	128	1024	9.94	346.73	27.89	342.74	27.76	61.87	83.18	3.27
	256	1024	10.68	386.64	28.78	666.22	28.65	109.20	83.46	3.31
twitter7	8	512	41.64	1727.13	45.18	184.17	56.11	180.85	71.66	3.46
	16	512	47.04	1774.92	53.49	213.17	61.73	186.27	76.78	3.57
	32	512	52.59	1884.16	58.15	244.18	66.84	184.90	79.34	3.49
	64	512	57.53	1988.11	62.95	330.46	68.68	197.86	80.62	3.50
	128	512	61.87	2113.34	66.68	504.53	69.94	219.65	81.26	3.93
	256	512	65.47	2357.92	78.32	846.20	71.26	280.99	81.57	3.51
sk-2005	8	1024	3.23	634.79	21.95	55.39	21.13	30.98	81.50	4.17
	16	1024	4.11	648.48	26.36	82.41	25.33	35.44	87.26	4.20
	32	1024	5.32	667.84	29.59	137.42	27.97	43.76	90.11	4.23
	64	1024	7.55	695.60	32.52	238.54	30.18	59.19	91.50	4.21
	128	1024	8.95	733.05	35.87	449.19	32.44	91.36	92.19	4.20
	256	1024	12.02	798.73	40.06	857.76	35.69	150.64	92.55	4.26
soc-friendster	8	1024	27.36	4099.35	30.57	405.68	45.60	381.78	87.53	5.45
	16	1024	34.50	4202.04	45.74	440.11	58.98	361.54	93.77	5.62
	32	1024	39.52	4345.96	54.87	503.90	61.00	408.56	96.89	5.49
	64	1024	46.35	4546.98	59.27	649.34	64.02	422.87	98.45	5.44
	128	1024	52.41	4796.56	60.82	888.14	68.17	475.16	99.22	5.72
	256	1024	57.79	5323.08	64.25	1426.16	71.90	523.66	99.61	5.53
er-fact1.5s26	8	1024	73.27	2216.99	73.44	259.98	73.44	208.81	87.50	5.57
	16	1024	80.18	2292.12	80.40	288.90	80.40	226.01	93.75	5.57
	32	1024	84.36	2400.35	84.63	357.21	84.63	255.60	96.87	5.60
	64	1024	86.99	2534.09	87.31	506.23	87.31	270.58	98.44	5.58
	128	1024	88.72	2725.81	89.10	769.61	89.10	407.59	99.22	5.57
	256	1024	89.99	2913.95	90.45	1329.57	90.45	408.35	99.61	5.65
RHG1	8	1024	0.04	380.04	2.02	86.95	2.02	44.42	91.91	8.37
	16	1024	0.06	391.63	2.12	143.47	2.12	52.11	97.39	8.57
	32	1024	0.09	406.56	2.16	252.15	2.16	65.65	99.19	8.38
	64	1024	0.15	435.71	2.17	450.73	2.17	95.22	99.74	8.33
	128	1024	0.22	482.06	2.18	877.69	2.18	147.22	99.90	8.31
	256	1024	0.34	569.77	2.19	1708.33	2.18	273.76	99.95	8.45
RHG2	8	1024	0.09	621.56	0.05	103.83	0.04	56.23	89.71	8.31
	16	1024	0.13	632.61	0.07	153.29	0.04	60.14	96.15	8.57
	32	1024	0.19	648.68	0.12	262.91	0.05	77.34	98.73	9.85
	64	1024	0.29	674.36	0.18	468.04	0.05	108.39	99.57	8.32
	128	1024	0.44	727.66	0.28	872.68	0.07	157.75	99.85	8.31
	256	1024	0.68	816.60	0.44	1686.18	0.09	278.51	99.92	8.47
uk-2007-05	8	1024	0.54	1024.26	25.23	107.46	25.21	58.19	87.91	8.80
	16	1024	0.60	1045.36	28.02	166.06	28.19	70.72	94.08	8.79
	32	1024	0.70	1058.73	29.40	278.38	29.32	85.42	97.12	8.99
	64	1024	0.92	1099.64	29.94	517.20	29.85	115.18	98.61	9.32
	128	1024	1.31	1163.45	30.73	935.98	30.18	175.01	99.33	8.79
	256	1024	1.95	1280.48	31.65	1808.52	30.70	324.71	99.68	8.92

Table 4.2. Experiments with huge graphs. CE and RT denote cut edges and running time.

Memory Consumption. We now shortly review the amount of memory needed by the streaming algorithms under consideration. First of all note that the memory of HeiStream depends on the size of the buffer that is used. If the buffer only contains one node, then the memory requirements match those of Fennel and LDG. Here, we measure the memory consumption of HeiStream for various buffer sizes and compare it to Fennel and LDG. To do that, we measured the memory consumption of HeiStream(1024k), HeiStream(32k), Fennel and LDG on the three largest graphs (RHG1, RHG2, and uk-2007-05). On average, HeiStream(1024k) consumes respectively 2.5GB, 4.1GB, and 9.9GB of memory to partition these graphs. while HeiStream(32k) consumes 472MB, 521MB, and 3.7GB, and Fennel and LDG use 399MB, 401MB, and 445MB. Note that the increased amount of used memory of HeiStream(1024k) is expected, since we use a fairly large buffer. Given the size of the graphs, we believe those required amounts of memory are more than reasonable.

4.2 Streaming Process Mapping

In this section, we develop Online Recursive Multi-Section, a shared-memory parallel streaming algorithm for process mapping. Our algorithm performs recursive multi-sections on the fly, which allows us to perform multi-sectioning along a hierarchical topology using only one pass over the communication graph. To the best of our knowledge, ours is the first streaming algorithm designed specifically for the process mapping problem. On average, our algorithm computes 41% better process mappings and is 55 times faster than Fennel which ignores the given hierarchy. In scalability tests, our algorithm is only 3 times slower than Hashing when running on 32 threads. If no hierarchy is specified as an input, our approach can also be used as a tool to solve the standard graph partitioning problem. Our approach has a significantly lower running time complexity compared to state-of-the-art non-buffered one-pass partitioning algorithms. Our experiments show that on average, our algorithm is 134.4 times faster than Fennel at the cost of 5% more cut edges.

4.2.1 Online Recursive Multi-Section

4.2.1.1 Overall Scheme

A successful offline algorithm to partition and map the nodes of a communication graph onto PEs is the recursive multi-section [92, 109]. This approach specializes the partitioning process for the case in which the communication cost between two processes (nodes) highly depends on the hierarchy level shared by the PEs (blocks) in which they are allocated. Recall that a hierarchical topology can be represented by a string $S = a_1 : a_2 : \dots : a_\ell$ which in the process mapping application means that each processor has a_1 cores, each node has a_2

Algorithm 4.2 Online Recursive Multi-Section

```

1: for  $u \in V(G)$  do
2:   Load  $N(u)$ 
3:    $X \leftarrow \{V_1^\ell, \dots, V_{a_\ell}^\ell\}, V^* \leftarrow \emptyset$ 
4:   for  $i \in \{\ell, \dots, 1\}$  do
5:     if  $i \neq \ell$  then  $X \leftarrow$  sub-blocks of  $V^*$ 
6:      $V^* \leftarrow \operatorname{argmax}_{W \in X} \{score(W)\}$ 
7:     Assign node to sub-block  $V^*$ 

```

processors, each rack has a_3 nodes, and so forth. The offline recursive multi-section works as follows. First, the whole graph is partitioned in a_ℓ blocks. Then, the subgraph induced by each block of an a_i -partition is recursively partitioned in a_{i-1} sub-blocks until the whole graph is partitioned in $k = \prod_{i=1}^{\ell} a_i$ blocks. This recursive approach has been developed for process mapping and exploits the fact that the communication between PEs is cheaper through lower layers of the communication hierarchy. It creates a hierarchy of partitioning subproblems that directly reflects the hierarchical topology of the system, which yields an improved process mapping in practice [109].

Intuitively, for any given hierarchy S (independent of the application), recursive multi-section can be implemented in the streaming model using ℓ successive passes of any one-pass partitioning algorithm over the input graph, i.e., by restreaming the graph ℓ times and subpartitioning the corresponding blocks in each pass. In our algorithm, which we call Online Recursive Multi-Section, we compress all steps performed during the ℓ passes into a single pass, as follows. After a node is loaded, assign it to one of the a_ℓ blocks $\{V_1^\ell, \dots, V_{a_\ell}^\ell\}$ in layer ℓ . Then, for each layer $i < \ell$, assign the node to one of the a_i sub-blocks of the block chosen in the previous step. After going through all layers, the node is directly assigned to a final block, which makes this approach feasible for online execution. Algorithm 4.2 summarizes the structure of our Online Recursive Multi-Section and Figure 4.8 exemplifies it. Here, score depends on the algorithm logic used, e.g. Fennel or LDG, as well as other parameters that are specific to our multi-section algorithm. We give more details in Section 4.2.1.2. Note that it produces exactly the same result as the version with ℓ passes over the input since it does not violate any data dependency.

Online Recursive Multi-Section can be applied to compute any hierarchical partitioning and is not limited to the process mapping application. However, in case of the process mapping application, as the offline recursive multi-section our streaming algorithm exploits the inherent structure of the problem in two ways: (i) its layout reproduces the hierarchical communication topology; (ii) the top-down order in which the nodes are assigned to sub-blocks of previously assigned blocks reflects the order in which the communication costs

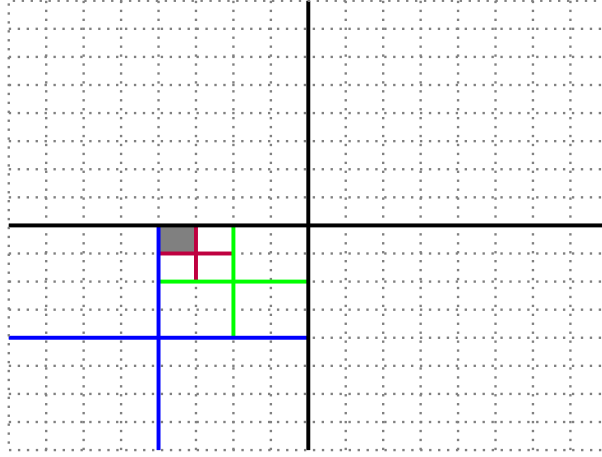


Figure 4.8. Assigning a node with Online Recursive Multi-Section over a grid composed of 256 blocks that are contained within a hierarchy $S = 4 : 4 : 4 : 4$. In the first 4-partition (black), the node is assigned to the lower left block. In the following 4-partitions, it is assigned to the upper right sub-block (blue), its upper left sub-block (green), and finally its upper left sub-block (purple), which is a block from the original partitioning problem.

decrease in a communication topology. In other words, this top-down order reflects the need for primarily avoiding cut edges among modules of higher layers in the communication hierarchy. We give the time and space complexity of our algorithm.

Lemma 4.2. Online Recursive Multi-Section needs $O(k)$ space to store block weights.

Proof. By definition, the multi-section consists of ℓ layers such that a layer $i \in \{1, \dots, \ell\}$ contains exactly $\prod_{r=i}^{\ell} a_r$ blocks whose weight we need to keep track of. As we define $a_r \geq 2 \forall r$, we can write $\prod_{r=i}^{\ell} a_r \leq (1/2^{i-1}) \prod_{r=1}^{\ell} a_r = (1/2^{i-1})k$. Hence, the total number of block weights that we need to keep track of is $\sum_{i=1}^{\ell} \prod_{r=i}^{\ell} a_r \leq \sum_{i=1}^{\ell} (1/2^{i-1})k \leq 2k$. \square

Theorem 4.3. Online Recursive Multi-Section coupled with Fennel or LDG needs $O(n + k)$ memory.

Proof. Due to the hierarchical structure of multi-section, Fennel and LDG may keep track of a single block assignment per node, which is enough to infer all its superblocks. Hence, the space complexity $O(n + k)$ directly follows from Lemma 4.2. \square

Theorem 4.4. Online Recursive Multi-Section coupled with Fennel or LDG has time complexity $O(m\ell + n \sum_{i=1}^{\ell} a_i)$.

Proof. Online Recursive Multi-Section assigns each node u over ℓ layers. Using Fennel or LDG, the running time to assign u in a given layer i is $O(|N(u)| + a_i)$. Accounting for all layers and nodes, this sums up to $O(m\ell + n \sum_{i=1}^{\ell} a_i)$. \square

Corollary 4.5. *Online Recursive Multi-Section coupled with Fennel or LDG has time complexity $O((m+n)\log k)$ if $a_i = b, \forall i \in \{1, \dots, \ell\}$ for any constant $b \geq 2$.*

Proof. Based on the assumption, we derive the claimed bound from Theorem 4.4 by proving $\ell = O(\log k)$ and $\sum_{i=1}^{\ell} a_i = O(\log k)$. The first part trivially holds since $k = \prod_{i=1}^{\ell} a_i = b^{\ell} \Rightarrow \ell = \log_b k$. To prove the second part, notice that $\sum_{i=1}^{\ell} a_i = b\ell$. Since $\ell = \log_b k$, it follows that $k = \prod_{i=1}^{\ell} a_i = b \log_b k = O(\log k)$. \square

4.2.1.2 Partitioning Subproblems

Our Online Recursive Multi-Section algorithm implies a hierarchy of one-pass partitioning subproblems. As self-contained, one-pass partitioning problems, these subproblems can be solved using any one-pass partitioning algorithm in literature. The focus of this section is to analyze the subproblems in relation to the parameters of a hierarchical partitioning problem.

Consider all the partitioning subproblems contained in some layer i of our algorithm. The initial point of consideration is the uniformity of these subproblems, which implies that they are supplied with an induced subgraph containing a similar quantity of nodes and edges, and subsequently partition it among a_i blocks. More specifically, a subproblem in layer i partitions among $k_i = a_i$ blocks and receives as input a graph containing roughly $n_i = n/\prod_{r=i+1}^{\ell} a_r$ nodes and $m_i = m/\prod_{r=i+1}^{\ell} a_r$ edges. Consequently, the size constraint L_i for a block from the subproblems in layer i is determined by the equation $L_i = \lceil (1 + \epsilon)n_i/k_i \rceil \simeq L_{\max} \prod_{r=1}^{i-1} a_r$, which simply denotes the total capacity of all blocks from the original problem that are encompassed within it. The variations observed in the dimensions subproblems at different layers within our algorithm have further implications based on the choice of streaming partitioning algorithm employed to address them, as we show next.

Fennel Mapping. Using Fennel within Online Recursive Multi-Section requires attention to its constant α . Recall that it is defined as $\alpha = \sqrt{km}/n^{3/2}$ for partitioning the whole graph into k blocks with vanilla Fennel. Using this value of α for all partitioning subproblems contained in our algorithm is not a natural choice since we intend to apply Fennel independently for each subproblem. Independently applying the definition of Fennel for each partitioning subproblem contained in our algorithm implies ℓ different parameters $\alpha_i, i \in \{1, \dots, \ell\}$, for all multi-section layers. We derive the value of α_i by applying the Fennel definition $\alpha_i = \sqrt{k_i \frac{m_i}{n_i^{3/2}}}$ and substituting the values of k_i, m_i , and n_i which we have already discussed. It follows that $\alpha_i = \frac{\alpha}{\sqrt{\prod_{r=1}^{i-1} a_r}}$.

LDG Mapping. Combining LDG with Online Recursive Multi-Section is straightforward, since it directly uses the remaining capacity of each block as a multiplicative penalty. Hence,

we can configure LDG for a subproblem in layer i by simply computing this penalty based on the block capacity L_i , whose value we have already discussed.

4.2.1.3 General Partitioning

In this section, we show how to partition a streamed graph into an arbitrary number of blocks using Online Recursive Multi-Section when no explicit hierarchy is given. We do this by creating an artificial hierarchy.

The *recursive bisection* is a successful offline approach to partition graphs into an arbitrary number k of blocks [66]. If k is a power of 2, the algorithm works as a recursive multi-section with $\log_2 k$ layers of 2-way partitioning subproblems. Otherwise, it is irregular and cannot be represented by a string S . In analogy to the offline algorithm, we define an online recursive bisection to partition a graph on the fly when no hierarchy is given. Recall that the whole hierarchy of blocks and sub-blocks has to be kept in memory throughout the execution of Online Recursive Multi-Section, hence the same requirement applies here. We build this hierarchy, which we call *multi-section tree*, as a preliminary step for the streaming partitioning process. In Algorithm 4.3, we define the procedure BuildHierarchy which recursively builds this multi-section tree for any value of k . This procedure receives as input a parent block P for the multi-section tree as well as the endpoints k_L and k_R of the range of blocks to be covered by the multi-section tree. In line 2, it terminates the recursion when P is a leaf of the multi-section tree, which is true when $k_L = k_R$. Otherwise, it creates two sub-blocks for P and inserts them as sons of P in the multi-section tree (line 3). Then, it splits the range $\{k_L, \dots, k_R\}$ in roughly equal parts and performs 2 recursive calls to itself.

We further generalize the recursive bisection to *recursive b-section* for a base b . Given a base $b \geq 2$, a recursive b -section is a recursive multi-section associated with a multi-section tree in which blocks have up to b sub-blocks. Algorithm 4.3 can be adapted to deal with b -section by creating $\min\{b, k_R - k_L + 1\}$ sub-blocks in line 3 and, afterwards, making the same number of recursive calls with proper parameters. We create the multi-section tree by

Algorithm 4.3 Create Blocks for Multi-Section Tree

Input P : Parent block in hierarchy, k_L : Left endpoint of blocks covered by hierarchy. k_R : Right endpoint of blocks covered by hierarchy

- 1: **procedure** BuildHierarchy(P, k_L, k_R)
 - 2: **if** $k_L = k_R$ **then return**
 - 3: $P_L, P_R \leftarrow$ Create sub-blocks for P
 - 4: BuildHierarchy($P_L, k_L, \lfloor \frac{k_L+k_R}{2} \rfloor$)
 - 5: BuildHierarchy($P_R, \lfloor \frac{k_L+k_R}{2} \rfloor + 1, k_R$)
-

calling the command $\text{BuildHierarchy}(\emptyset, 1, k)$ at the cost of $O(k)$. Given a multi-section tree, we solve it by using Algorithm 4.2. Analogously to Theorem 4.3, it is possible to prove that the online recursive b -section respectively stores $O(k)$ blocks and needs $O(n + k)$ memory when coupled with Fennel or LDG. Theorem 4.6 provides a running time bound.

Theorem 4.6. *Online recursive b -section coupled with Fennel or LDG has time complexity $O((m + nb) \log_b k)$.*

Proof. The number of layers in the multi-section tree is up to $\lceil \log_b k \rceil$. In other words, each node should be assigned through up to $1 + \log_b k$ layers. Since all subproblems partition among up to b blocks, then the running time to assign a node u over a layer is $|N(u)| + b$. Accounting for all layers and nodes, this sums up to $(2m + nb)(\log_b k + 1) = O((m + nb) \log_b k)$. \square

Heterogeneous Partitioning. When k is not a power of b , the recursive b -section hierarchy may contain some partitioning subproblems with heterogeneous blocks. We deal with this by computing the size constraint of each block in the multi-section tree individually. For simplicity, we explain how to do this when $b = 2$ (recursive bisection), but this can be easily extended to an arbitrary b . For example when $k = 5$, the two blocks in the first 2-way partitioning subproblem respectively cover 2 and 3 of the blocks from the original 5-way partitioning. Hence, these two blocks shall respectively have capacities $2L_{\max}$ and $3L_{\max}$, where L_{\max} is the size constraint of a block in the original k -way partitioning. Putting it in general terms, each block from the multi-section tree created in line 3 of Algorithm 4.3 covers $t = k_R - k_L + 1$ blocks of the original k -way partitioning. For simplicity, we use t to refer to this number covered by a given block, and we use t_1 and t_2 to refer to the numbers covered by the two blocks of a partitioning subproblem. The size constraint of a block is $t \times L_{\max}$.

When a subproblem has blocks with heterogeneous size constraints, the used partitioning algorithm has to cope with it. We adapt Fennel to address this issue by increasing (decreasing) the constant α used to compute the score of a specific block when its size constraint is lower (higher) than the other block from the same subproblem. Recall that α depends on the numbers of nodes, edges and blocks for a specific subproblem. A subproblem receives as input an induced subgraph with roughly $\frac{t_1+t_2}{k}$ of the nodes and edges from the original k -way partitioning. We redefine the number of blocks as $\frac{t_1+t_2}{t_1}$ for the first block and $\frac{t_1+t_2}{t_2}$ for the second block of a subproblem. This value equals 2 for both blocks when $t_1 = t_2$. Nevertheless, if $t_1 \neq t_2$, this value is larger (smaller) than 2 for the block with smaller (larger) size constraint. Summing up, the value of α for a given block will be \sqrt{t} times smaller than the value α from the original k -way partitioning problem. Consequently,

the imbalance penalty function for Fennel will be more pronounced for blocks with lower capacities, thereby addressing the issue of heterogeneous balancing.. For LDG, a natural adaptation for heterogeneous blocks arises from its very definition, since it directly uses the remaining capacity of each block as a multiplicative penalty.

4.2.1.4 Shared-Memory Parallelization

Since Online Recursive Multi-Section is a node-centric algorithm, we can parallelize it by independently splitting the nodes of the graph among threads. More specifically, it can be achieved with OpenMP by parallelizing the *for* loop in line 1 of Algorithm 4.2. This parallelization requires the nodes from the input graph to be concurrently loaded by distinct threads alongside with their neighborhoods, which is a reasonable assumption in many practical environments. Regarding data consistency, the only source of concern are the block weights, whose values can be concurrently read and incremented by multiple threads. This is important because an inconsistency could compromise the load balance between blocks. We ensure writing consistency by making the incrementation an atomic operation. Potentially, a block can still be overloaded if multiple threads decide to assign a node to it at the same time. Since this is very unlikely, we do not use any synchronization to keep it from happening.

4.2.2 Experimental Evaluation

Setup. We performed our implementations using C++ and compiled them using gcc 9.3 with full optimization turned on (-O3 flag). The majority of our experiments were executed on a single core of Machine C, with the exception of the scalability experiments, which employed up to 32 threads using hyperthreading. Unless otherwise mentioned we stream the input directly from the internal memory to obtain clear running time comparisons. However, note that the streaming algorithm could also be run streaming the graph from hard disk.

Baselines. We identify Fennel and LDG as the state-of-the-art of non-buffered one-pass stream partitioning algorithms which aim at minimizing edge-cut. Since Fennel generates better solutions on average than LDG [3], we focus our experiments on Fennel without loss of generality. We also include Hashing as a competitor, since it is the fastest algorithm for streaming partitioning. To the best of our knowledge, there are no streaming partitioning algorithms specifically designed for process mapping. Since no official versions of Fennel, LDG, and Hashing are available in public repositories, we implemented them in our framework. For comparison purposes, we ran experiments with internal-memory tools, i.e. we compare against the fastest version of the integrated multi-level algorithm proposed in Chap-


ter 6.1 of this dissertation, which we refer to as *IntMap*. In addition, we compare our results against *KaMinPar* [63], a very fast internal-memory algorithm that is orders of magnitudes faster than *mt-Metis* in terms of running time and produces comparable cuts while also enforcing balance (in contrast to *mt-Metis*). In particular, the purpose of running *IntMap* and *KaMinPar* is to provide a reference of streaming algorithms in comparison to internal memory algorithms. We set a timeout of 30 minutes for an algorithm to partition a graph. Only *IntMap* exceeded this time limit for some instances. Hence, we exclude this algorithm from the plots. Our implementations of these algorithms reproduce the results presented in the respective papers and are optimized for running time as much as possible.

Instances. In this section, we use two disjoint sets of graphs: the *tuning* set is used for the parameter study experiments, while the *test* set is used for the comparisons against the state-of-the-art. Basic properties of the graphs under consideration can be found in Table 2.1. In any case, when streaming the graphs we use the natural given order of the nodes. Unless otherwise mentioned, we use the following configurations for process mapping experiments: $D = 1 : 10 : 100$, $S = 4 : 16 : r$, with $r \in \{1, 2, 3, \dots, 128\}$. Hence, $k = 64r$. This is the same configuration used in other studies [92, 109]. Analogously, we use $k = 64s$, $s \in \{1, 2, 3, \dots, 128\}$ for graph partitioning experiments unless mentioned otherwise. We allow a fixed imbalance of 3% for all experiments (and all algorithms) since this is a frequently used value in the partitioning literature. All partitions computed by all algorithms were balanced.


Methodology. We perform two types of experiments: experiments for the process mapping objective (with given hierarchies as specified below) and standard graph partitioning in which we create implicit hierarchies as described above. Depending on the focus of the experiment, we measure running time, edge-cut, and/or the process mapping communication cost defined in Equation (2.1).

4.2.2.1 Parameter Study

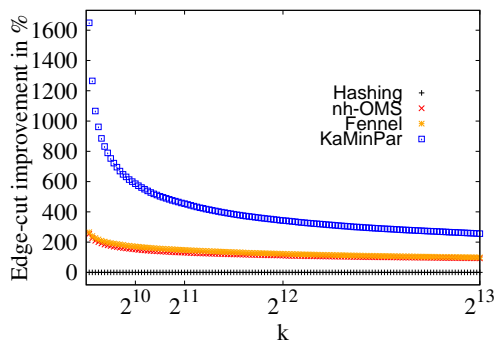
We performed extensive tuning experiments using the graphs from the tuning set in Table 2.1. We briefly summarize the main results here. Online Recursive Multi-Section produces on average 3.89% better mapping and 0.19% better edge-cut when coupled with *Fennel* than when coupled with *LDG*. Hence, we use *Fennel* as our scoring function. Computing adapted values of α for each partitioning subproblem is superior than using the default value of α of the original k -way partitioning. Particularly, it is on average 3.1% faster while producing 9.7% better mapping and cutting roughly the same number of edges. Hence, our algorithm uses *adapted* α values. When no communication hierarchy is given, using the *base* $b = 4$ to



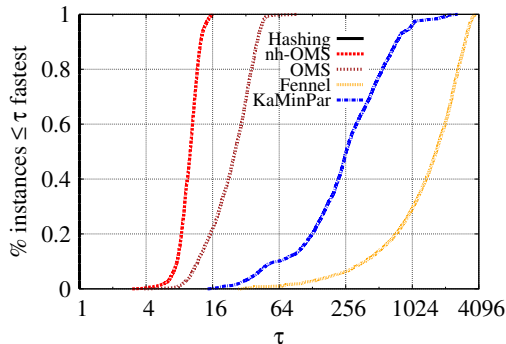
(a) Mapping improvement over Hashing.



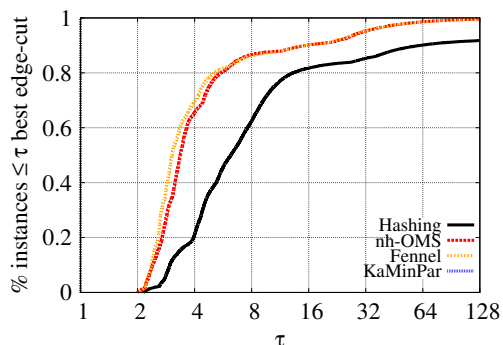
(b) Mapping performance profile.



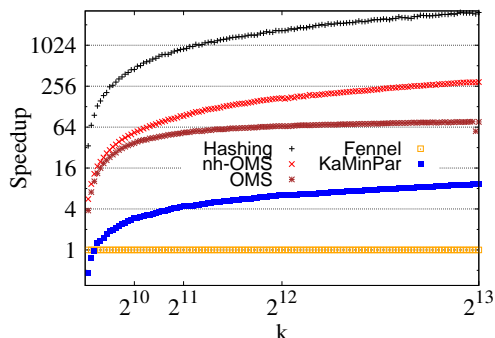
(c) Edge-cut improvement over Hashing.



(d) Running time performance profile.




(e) Edge-cut performance profile.




(f) Total speedup over Fennel.


Figure 4.9. Comparison against the state-of-the-art. Higher is better.




(a) Speedup versus number of used threads for graph soc-orkut-dir.




(b) Running time versus number of used threads for graph soc-orkut-dir.



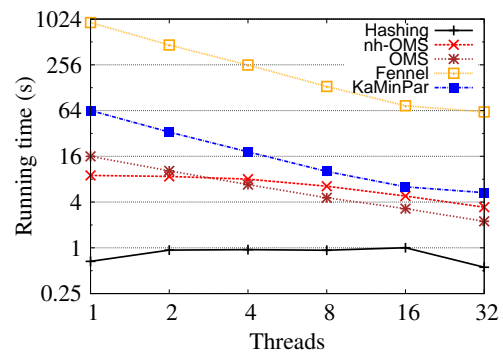
(c) Speedup versus number of used threads for graph HV15R.



(d) Running time versus number of used threads for graph HV15R.



(e) Speedup versus number of used threads for graph soc-LiveJournal1.



(f) Running time versus number of used threads for graph soc-LiveJournal1.

Figure 4.10. Speedup and time comparison for $k = 8192$. Higher is better for speedup. Lower is better for time.

build the multi-section tree is the fastest configuration overall. Using $b = 4$, our algorithm is 16.7% faster and cuts 3.2% fewer edges than using $b = 2$. Hence, our algorithm uses *base* $b = 4$. From now on, we refer to our Online Recursive Multi-Section algorithm as OMS when a communication hierarchy is given and nh-OMS otherwise.

4.2.2.2 State-of-the-Art

In this section, we show experiments in which we compare Online Recursive Multi-Section against the current state-of-the-art.

Solution Quality (Process Mapping). We start by looking at the mapping quality produced by OMS. In Figure 4.9a, we plot the average mapping improvement over Hashing. KaMinPar produces the best mapping overall, with an average improvement of 1117% over Hashing. Among the instances IntMap could solve, it improves on average 7.6% over KaMinPar. IntMap produces the best overall mapping in 67% of the cases it could solve. Note that this is in line with previous works in the area of graph partitioning, i.e. streaming algorithms typically compute worse solutions than internal memory algorithms that always have access to the whole graph. OMS has an average improvement of 257.8% over Hashing, while Fennel improves 153% on average over Hashing. In a direct comparison, OMS produces on average 41% better mappings than Fennel. In Figure 4.9b, we plot the mapping performance profile. In the plot, KaMinPar produces the best overall mapping for all instances. We conclude that OMS produces the best mapping among the streaming competitors.

Solution Quality (Edge-Cut). Next we look at the edge-cut of nh-OMS. In Figure 4.9c, we plot the edge-cut improvement over Hashing. KaMinPar produces the best overall edge-cut, with an average improvement of 3024% over Hashing. IntMap cuts 20% more edges on average than KaMinPar for the instances it solved. Among the streaming algorithms, Fennel and nh-OMS produce improvements of respectively 130.5% and 118.2% on average over Hashing. In a direct comparison, nh-OMS cuts on average 5% more edges than Fennel. In Figure 4.9e, we plot the edge-cut performance profile. KaMinPar produces the smallest edge-cut for all instances. Among the streaming algorithms, Fennel is slightly better than nh-OMS and both are distinctly better than Hashing.

Running Time. We now investigate the running time of OMS and nh-OMS. In Figure 4.9f, we plot the speedup over Fennel. On average, Hashing is 1301 times faster than Fennel, while nh-OMS and OMS are respectively 133 and 55.4 times faster than Fennel. In a direct comparison, Hashing is on average 9.7 times faster than nh-OMS and 23.4 times faster than OMS. KaMinPar comes next with an average speedup of 5.3 over Fennel. In a direct comparison,

Threads	Hashing		nh-OMS		OMS		Fennel		KaMinPar	
	RT	SU	RT	SU	RT	SU	RT	SU	RT	SU
1	0.49	1.0	4.8	1.0	10.7	1.0	673.6	1.0	36.9	1.0
2	0.72	0.7	4.6	1.1	6.4	1.7	346.3	1.9	19.3	1.9
4	0.70	0.7	3.6	1.3	3.9	2.7	184.6	3.6	10.5	3.5
8	0.72	0.7	3.0	1.6	2.6	4.1	96.3	7.0	5.8	6.4
16	0.75	0.7	2.5	1.9	2.3	4.7	54.0	12.5	3.5	10.5
32	0.46	1.1	1.7	2.8	1.3	8.2	44.2	15.2	3.1	11.9

Table 4.3. Average time in seconds (RT) and average speedup (SU) for $k = 8192$.

nh-OMS and OMS are respectively 25.1 and 10.5 times faster than KaMinPar. KaMinPar is on average 2.5 times faster than IntMap for the instances IntMap could solve. In Figure 4.9d, we plot the running time performance profile. Note that the running time of nh-OMS is at most 16 times slower than Hashing for 100% of the experiments, which is in accordance with Theorem 4.6. As the third fastest algorithm, OMS is considerably faster than all the other competitors, including Fennel.

Memory Requirements. We now look at the memory requirements of the different algorithms. We measured this on three graphs of our collection. In this case, we run stream graphs directly from disk for the streaming algorithms. Besides being the fastest algorithm, Hashing needs the least memory overall. For soc-orkut-dir, HV15R, and soc-LiveJournal1, it respectively needs 17MB, 13MB, and 24MB. OMS, nh-OMS, and Fennel have comparable consumption, all of which using 19MB, 14MB, and 25MB for the mentioned graphs, respectively. Finally, KaMinPar respectively uses 4.1GB, 4.1GB, and 1.8GB, while IntMap respectively uses 34GB, 12GB, and 10GB.

4.2.2.3 Scalability

Now we evaluate the scalability of Online Recursive Multi-Section. As in Section 4.2.2.2, we refer to our algorithm as OMS when a process mapping communication hierarchy is given and nh-OMS otherwise. As competitors, we include KaMinPar, Fennel, and Hashing. For a fair comparison against Fennel and Hashing, we implemented them with the same parallelization scheme of our algorithm, i. e., a node-centric parallelization. We do not include IntMap in these experiments since it cannot run in parallel. For these experiments, we selected the 12 graphs from our collection which have at least two million nodes and partitioned them into $k = 8192$ blocks using all algorithms.

In Figure 4.10, we plot speedup and running time versus number of threads for the graphs soc-orkut-dir, HV15R, and soc-LiveJournal1. In Table 4.3, we plot the average running time in seconds and speedup over all graphs as a function of the number of threads. For all graphs,

Hashing presents the worst scalability, with speedups smaller than 1. Although Hashing is theoretically an embarrassingly parallel algorithm, it has two limitations: (i) it is extremely fast, hence the overhead of the parallelization is too large in comparison to the overall running time and (ii) it neither reuses data nor accesses sequential positions in memory, so there are almost no cache hits. On the other hand, Fennel presents the best scalability. Differently than Hashing, it is rather slow but reuses data, e. g., the assignments of previous nodes to blocks, and goes through the array of blocks in order to compute their score. Following Fennel, KaMinPar has the second best scalability which roughly reproduces the behavior reported in [63]. The algorithms OMS and nh-OMS have an intermediary scalability between KaMinPar and Hashing. This is explained by their characteristics, which are intermediary between Fennel and Hashing. Note that OMS is more scalable than nh-OMS. This happens because OMS goes through and scores several blocks in one of the partitioning subproblems contained in the multi-section hierarchy, which favors cache hits, whereas nh-OMS partitioning subproblems have at most 4 blocks. For 32 threads, the average running time of OMS is within a factor 3 of the running time of Hashing.

4.3 Streaming Hypergraph Partitioning

In this section, we propose FREIGHT: a Fast stREamInG Hypergraph parTitioning algorithm that can optimize for cut-net and connectivity. By using an efficient data structure, we make the overall running time of FREIGHT linearly dependent on the pin-count of the hypergraph and the memory consumption linearly dependent on the numbers of nets and blocks. Our proposed algorithm demonstrates remarkable efficiency, with a running time comparable to the Hashing algorithm and a maximum discrepancy of only four in three quarters of the instances. Our study establishes the superiority of FREIGHT over all current (buffered) streaming algorithms and even the in-memory algorithm HYPE, in both cut-net and connectivity measures. This shows the potential of our algorithm as a valuable tool for partitioning hypergraphs in the context of large and constantly changing data processing environments.

4.3.1 FREIGHT

4.3.1.1 Mathematical Definition

In this section, we provide a mathematical definition for FREIGHT by expanding the idea of Fennel to the domain of hypergraphs. Recall that, assuming the nodes of a graph being streamed one-by-one, the Fennel algorithm assigns an incoming node v to a block V_d where d is computed as follows:

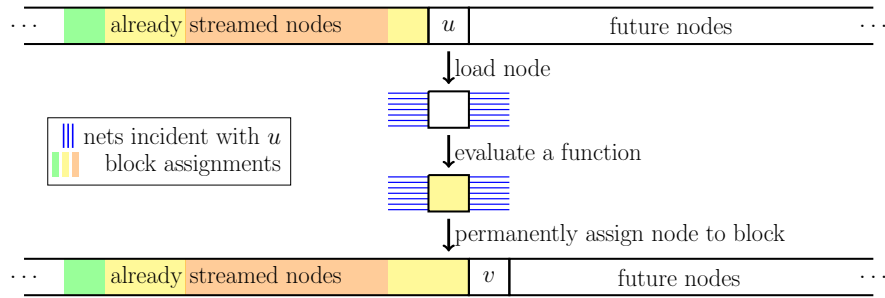


Figure 4.11. Typical layout of streaming algorithm for hypergraph partitioning.

$$d = \underset{i, |V_i| < L_{\max}}{\operatorname{argmax}} \{ |V_i \cap N(v)| - \alpha * \gamma * |V_i|^{\gamma-1} \} \quad (4.6)$$

The term $-\alpha * \gamma * |V_i|^{\gamma-1}$, which penalizes block imbalance in Fennel, is directly used in FREIGHT without modification and with the same meaning. The term $|V_i \cap N(v)|$, which minimizes edge-cut in Fennel, needs to be adapted in FREIGHT to minimize the intended metric, i.e., either cut-net or connectivity. Before explaining how this is adapted, recall that, in contrast to graph partitioning, in hypergraph partitioning the incident nets $I(v)$ of an incoming node v might contain nets that are already cut, i.e., with pins assigned to multiple blocks. The version of FREIGHT designed to optimize for *connectivity* accounts for already cut nets by keeping track of the block d_e to which the most recently streamed pin of each net e has been assigned. More formally, the connectivity version of FREIGHT assigns an incoming node v of a hypergraph to a block V_d with d given by Equation (4.7), where $I_{obj}^i(v) = I_{con}^i(v) = \{e \in I(v) : d_e = i\}$. On the other hand, the version of FREIGHT designed to optimize for *cut-net* ignores already cut nets, since their contribution to the overall cut-net of the hypergraph k -partition is fixed and cannot be changed anymore. More formally, the cut-net version of FREIGHT assigns an incoming node v of a hypergraph to a block V_d with d given by Equation (4.7), where $I_{obj}^i(v) = I_{cut}^i(v) = I_{con}^i(v) \setminus E'$ and E' is the set of already cut nets.

$$d = \underset{i, |V_i| < L_{\max}}{\operatorname{argmax}} \{ |I_{obj}^i(v)| - \alpha * \gamma * |V_i|^{\gamma-1} \} \quad (4.7)$$

Both configurations of FREIGHT interpolate two objectives: favoring blocks with many incident (uncut) nets and penalizing blocks with large cardinality. We briefly highlight that FREIGHT can be adapted for weighted hypergraphs. In particular, when dealing with weighted nets, the term $|I_{obj}^i(v)|$ is substituted by $\omega(I_{obj}^i(v))$. Likewise when dealing with weighted nodes, the term $-\alpha * \gamma * |V_i|^{\gamma-1}$ is substituted by $-c(v) * \alpha * \gamma * c(V_i)^{\gamma-1}$, where the weight $c(v)$ of v is used as a multiplicative factor in the penalty term.

4.3.1.2 Streaming Hypergraphs

In this section, we present and discuss the streaming model used by FREIGHT. Recall in the streaming model for graphs, nodes are loaded one at a time alongside with their adjacency lists. Thus, just streaming the graph (without doing additional computations, implies a time cost $O(m + n)$). In our model, the nodes of a hypergraph are loaded one at a time alongside with their incident nets, as illustrated in Figure 4.11. Our streaming model implies a time cost $O(\sum_{e \in E} |e| + n)$ just to stream the hypergraph, where $O(\sum_{e \in E} |e|)$ is the cost to stream each net e exactly $|e|$ times. FREIGHT uses $O(m + k)$ memory, with $O(m)$ being used to keep track, for each net e , of its cut/uncut status as well as the block d_e to which its most recently streamed pin was assigned. This net-tracking information, which substitutes the need to keep track of node assignments, is necessary for executing FREIGHT. Although FREIGHT consumes more memory than required by graph-based streaming algorithms which often use $O(n + k)$ memory, it is still far better than the $O(mk)$ worst-case memory required by the state-of-the-art algorithms for streaming hypergraph partitioning [8, 9], all of which are also based on a computational model that implies a time cost $O(\sum_{e \in E} |e| + n)$ just to stream the hypergraph.

4.3.1.3 Efficient Implementation

In this section, we describe an efficient implementation for FREIGHT. Recall that, for every node v that is loaded, FREIGHT uses Equation (4.7) to find the block with the highest score among up to k options. A simple method to accomplish this task consists of explicitly evaluating the score for each block and identifying the one with the highest score. This results in a total of $O(nk)$ evaluations, leading to an overall complexity of $O(\sum_{e \in E} |e| + nk)$. We propose an implementation that is significantly more efficient than this approach.

For each loaded node v , our implementation separates the blocks V_i for which $|V_i| < L_{\max}$ into two disjoint sets, S_1 and S_2 . In particular, the set S_1 is composed of the blocks V_i where $|I_{obj}^i(v)| > 0$, while the set S_2 comprises the remaining blocks, i.e., blocks V_i for which $|I_{obj}^i(v)| = 0$. Using the sets provided, we break down Equation (4.7) into Equation (4.8) and Equation (4.9), which are solved separately. The resulting solutions are compared based on their FREIGHT scores to ultimately find the solution for Equation (4.7). The overall process is illustrated in Figure 4.12.

$$d = \operatorname{argmax}_{i \in S_1} \{ |I_{obj}^i(v)| - \alpha * \gamma * |V_i|^{\gamma-1} \} \quad (4.8)$$

$$d = \operatorname{argmax}_{i \in S_2} \{ |I_{obj}^i(v)| - \alpha * \gamma * |V_i|^{\gamma-1} \} = \operatorname{argmin}_{i \in S_2} |V_i| \quad (4.9)$$

Now we explain how we solve Equation (4.8) and Equation (4.9). To solve Equation (4.8), we use the theoretical complexity outlined in Theorem 4.7 and solve it explicitly.

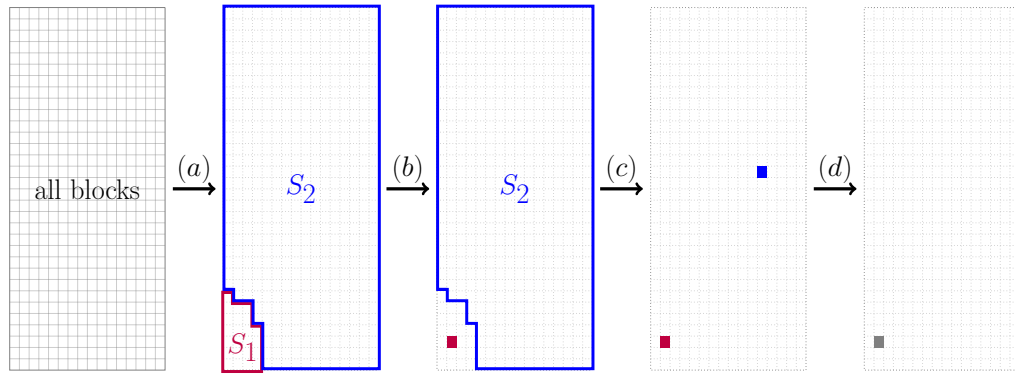


Figure 4.12. Illustration of the process to solve Equation (4.7) for an incoming node u with $k = 512$ blocks. (a) The k blocks are decomposed into S_1 and S_2 , with $|S_1| = O(|I(u)|)$. (b) Equation (4.8) is explicitly solved at cost $O(|I(u)|)$. (c) Equation (4.9) is implicitly solved at cost $O(1)$. (d) Both solutions are then evaluated using their FREIGHT scores to determine the final solution for Equation (4.7).

In contrast, Equation (4.9) is implicitly solved by identifying the block with minimal cardinality. We use an efficient data structure to keep all blocks sorted by cardinality throughout the entire execution, which enables us to solve Equation (4.9) in constant time.

Theorem 4.7. Equation (4.8) can be solved in time $O(|I(v)|)$.

Proof. The terms $|I_{obj}^i(v)|$ in Equation (4.8) can be computed by iterating through the nets of v at a cost of $O(|I(v)|)$ and determining their status as cut, unassigned, or assigned to a block. The calculation of the factors $-\alpha * \gamma * |V_i|^{\gamma-1}$ in Equation (4.8) can be done in time $O(|S_1|) = O(|I(v)|)$, thus completing the proof. \square

Now we explain our data structure to keep the blocks sorted by cardinality during the whole algorithm execution. The data structure is implemented with two arrays A and B , both with k elements, and a list L . The array A stores all k blocks always in ascending order. The array B maps the index i of a block V_i to its position in A . Each element in the list L represents a bucket. Each bucket is associated with a unique block cardinality and contains the leftmost and the rightmost positions ℓ and r of the range of blocks in A which currently have this cardinality. Reciprocally, each block in A has a pointer to the unique bucket in L corresponding to its cardinality. To begin the algorithm, L is set up with a single bucket for cardinality 0 which covers the k positions of A , i.e., its parameters ℓ and r are 1 and k , respectively. The blocks in A are sorted in any order initially, however, as each block starts with a cardinality of 0, they will be ordered by their cardinalities.

When a node is assigned to a block V_d , we update our data structure as detailed in Algorithm 4.4 and exemplified in Figure 4.13. We describe Algorithm 4.4 in detail now. In




Figure 4.13. Illustration of our data structure used to keep the blocks sorted by cardinality throughout the execution of FREIGHT. The array A is represented as a vertical rectangle. Each region of A is covered by a unique bucket, which is represented by a unique color filling the corresponding region in A . The cardinality associated with each bucket is written in the middle of the region of A covered by it. Here we represent the behavior of the data structure when assigning nodes to the block surrounded by a dotted rectangle five times consecutively.

line 1, we find the position p of V_d in A and find the bucket C associated with it. In line 2, we exchange the content of two positions in A : the position where V_d is located and the position identified by the variable r in C , which marks the rightmost block in A covered by C . This variable r is afterwards decremented in line 3 since V_d is now not covered anymore by the bucket C . In lines 4 and 5, we check if the new (increased) cardinality of V_d matches the cardinality of the block located right after it in A . If so, we associate V_d to the same bucket as it and decrement this bucket's leftmost position ℓ in line 6; Otherwise, we push a new bucket to L and match it to V_d adequately in lines 8 and 9. Finally, in line 10, we delete C in case its range $[\ell, r]$ is empty. Figure 4.13 shows our data structure through five consecutive executions of Algorithm 4.4. Theorem 4.8 proves the correctness of our data structure. Theorem 4.9 shows that, using our proposed data structure, we need time $O(1)$ to either solve Equation (4.9) or prove that the solution for Equation (4.8) solves Equation (4.7). Note that our data structure can only handle unweighted vertices. In case of weighted vertices, a bucket queue can be used instead of our data structure, resulting in the same overall complexity and requiring $O(k + L_{\max})$ memory, while our data structure only requires $O(k)$ memory. The overall complexity of FREIGHT, which directly follows from Theorem 4.7 and Theorem 4.9, is expressed in Corollary 4.10.

Algorithm 4.4 Increment cardinality of block V_d in the proposed data structure

```

1:  $p \leftarrow B_d; C \leftarrow A_p.bucket;$ 
2:  $q \leftarrow C.r; c \leftarrow A_q.id; Swap(A_p, A_q); Swap(B_c, B_d);$ 
3:  $C.r \leftarrow C.r - 1;$ 
4:  $C' \leftarrow A_{q+1}.bucket;$ 
5: if  $C.cardinality + 1 = C'.cardinality$  then
6:    $A_q.bucket \leftarrow C'; C'.l \leftarrow C'.l - 1;$ 
7: else
8:    $C'' \leftarrow NewBucket(); A_q.bucket \leftarrow C''; L \leftarrow L \cup \{C''\};$ 
9:    $C''.cardinality \leftarrow C.cardinality + 1; C''.l \leftarrow q; C''.r \leftarrow q;$ 
10: if  $C.r = C.l$  then  $L \leftarrow L \setminus \{C\};$ 

```

Theorem 4.8. *Our proposed data structure keeps the blocks within array A consistently sorted in ascending order of cardinality.*

Proof. We inductively prove two claims at the same time: (a) the variables ℓ and r contained in each bucket from L respectively store the leftmost and the rightmost positions of the unique range of blocks in A which currently have this cardinality; (b) the array A contains the blocks sorted in ascending order of cardinality. Both claims are trivially true at the beginning, since all blocks have cardinality 0 and L is initialized with a single bucket with $\ell = 1$ and $r = k$. Now assuming that (a) and (b) are true at some point, we show that they keep being true after Algorithm 4.4 is executed. Note that line 2 performs the only position exchange in A throughout the whole algorithm. As (a) is assumed, it is the case that V_d swaps positions with the rightmost block in A containing the same cardinality of V_d . Since the cardinality of V_d will be incremented by one and all blocks have integer cardinalities, this concludes the proof of (b). To prove that (a) remains true, note that the only buckets in L that are modified are C (line 3), C' (line 6), and C'' (line 9). Claim (a) remains true for C because V_d , whose cardinality will be incremented, is the only block removed from its range. Claim (a) remains true for C' because line 6 is only executed if the new cardinality of V_d equals the cardinality of C' , whose current range starts right after the new position of V_d in A . Bucket C'' is only created if the new cardinality of V_d is respectively larger and smaller than the cardinalities of C and C' . Since (b) is true, then this condition only happens if there is no block in A with the same cardinality as the new cardinality of V_d . Hence, claim (a) remains true for C'' , which is created covering only the position of V_d in A . \square

Theorem 4.9. *By utilizing our proposed data structure, solving Equation (4.9) or demonstrating that any solution for Equation (4.8) is also a solution for Equation (4.7) can be accomplished in $O(1)$ time.*

Proof. Algorithm 4.4 contains no loops and each command in it has a complexity of $O(1)$, thus the total cost of the algorithm is $O(1)$. Our data structure executes Algorithm 4.4 once for each assigned node, hence it costs $O(1)$ per node. Say we are evaluating an incoming node v . According to Theorem 4.8, the block V_d with minimum cardinality is stored in the first position of the array A , hence it can be accessed in time $O(1)$. In case $V_d \in S_2$, then d is a solution for Equation (4.9). On the other hand, if V_d is in S_1 , the FREIGHT score of V_d will be larger than the FREIGHT score of the solution for Equation (4.9) by at least $|I_{obj}^d(v)| > 0$. In this case, it follows that any solution for Equation (4.8) solves Equation (4.7). \square

Corollary 4.10. *The overall complexity of FREIGHT is $O(\sum_{e \in E} |e| + n)$.*

4.3.2 Experimental Evaluation

Setup. We performed our implementations in C++ and compiled them using gcc 11.2 with full optimization turned on (-O3 flag). Unless mentioned otherwise, all experiments are performed on a single core of Machine C. Unless otherwise mentioned we stream (hyper)graphs directly from the internal memory to obtain clear running time comparisons. However, note that FREIGHT as well as most of the other used algorithms can also be run streaming the hypergraphs from hard disk.

Baselines. We compare FREIGHT against various state-of-the-art algorithms. In this section we will list these algorithms and explain our criteria for algorithm selection. We have implemented Hashing in C++, since it is a simple algorithm. It basically consists of hashing the IDs of incoming nodes into $\{1, \dots, k\}$. The remaining algorithms were obtained either from official repositories or privately from the authors, with the exception of Min-Max, for which there is no official implementation available. Here, we use the Min-Max implementations by Taşyaran et al. [9]. All algorithms were compiled with gcc 11.2.

We run Hashing, Min-Max [8] and its improved versions proposed in [9]: Min-Max-BF, Min-Max-N2P, Min-Max-L ℓ , Min-Max-MH, REF, REF_RLX, and REF_RLX_SV. (see Section 3.1.4 for details on the different Min-Max versions), HYPE [61], and PaToH v3.3 [57]. Hashing is relevant because it is the simplest and fastest streaming algorithm, which gives us a lower bound for partitioning time. Min-Max is a current state-of-the-art for streaming hypergraph partitioning in terms of cut-net and connectivity. The improved and buffered versions of Min-Max proposed in [9] are relevant because some of them are orders of magnitude faster than Min-Max while others produce improved partitions in comparison to it. HYPE and PaToH are in-memory algorithms for hypergraph partitioning, hence they are not suitable for the streaming setting. However, we compare against them because HYPE is among the fastest in-memory algorithms while PaToH is very fast and also computes par-

titions with very good cut-net and connectivity. Note that KaHyPar [52] is the leading tool with respect to solution quality, however it is also much slower than PaToH.


Instances. In our experiments we use 310 hypergraphs which are described in more detail in Section 2.8.2. Prior to each experiment, we converted all hypergraphs to the appropriate streaming formats required by each algorithm. We removed parallel and empty hyperedges and self loops, and assigned unitary weight to all nodes and hyperedges. In all experiments with streaming algorithms, we stream the hypergraphs with the natural given order of the nodes. We use a number of blocks $k \in \{512, 1024, 1536, 2048, 2560\}$ unless mentioned otherwise. We allow a fixed imbalance of 3% for all experiments (and all algorithms) since this is a frequently used value in the partitioning literature. All algorithms always generated balanced partitions, except for HYPE which generated highly unbalanced partitions in around 5% of its experiments.

Methodology. Depending on the focus of the experiment, we measure running time, cut-net, and-or connectivity. We perform 5 repetitions per algorithm and instance using random seeds for non-deterministic algorithms.

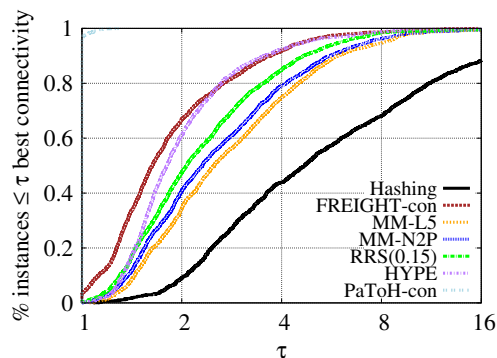
4.3.2.1 State-of-the-Art

In this section, we show experiments in which we compare FREIGHT against the current state-of-the-art of streaming hypergraph partitioning. As already mentioned, we also use two internal-memory algorithms [57, 61] as more general baselines for comparison. We focus our experimental evaluation on the comparison of solution quality and running time. Observe that PaToH and FREIGHT have distinct versions designed to optimize for each quality metric (i.e., connectivity and cut-net). For a meaningful comparison, we only take into account the relevant version when dealing with each quality metric, however, both versions are still considered for running time comparisons. To differentiate between the versions, suffixes -con and -cut are added to represent the connectivity-optimized and cut-net versions respectively.


For clarity, we refrain from discussing state-of-the-art streaming algorithms that are *dominated* by another algorithm. We define a dominated algorithm as one that has worse running time compared to another without offering a superior solution quality in return, or vice-versa. In particular, we leave out Min-Max and Min-Max-BF since they are dominated by Min-Max-N2P, which is referred to as MM-N2P hereafter. Similarly, we omit Min-Max-MH because it is dominated by Hashing. We use a buffer size of 15% for testing the buffered algorithms REF, REF_RLX, and REF_RLX_SV, following the best results outlined in [9]. We omit the first two of them since they are dominated by the latter one,




(a) Connectivity improvement over Hashing.



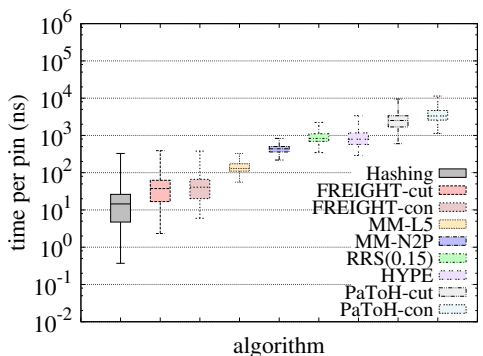
(b) Connectivity performance profiles.



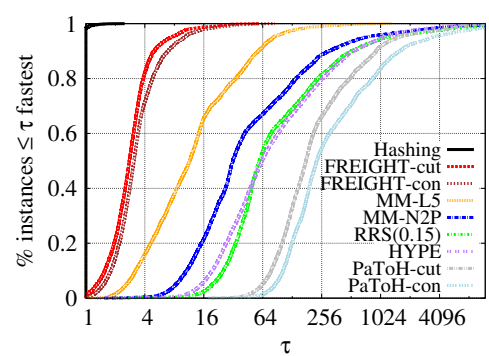
(c) Cut-net improvement over Hashing.



(d) Cut-net performance profiles.



(e) Running time boxplots.



(f) Running time performance profiles.

Figure 4.14. Comparison against the state-of-the-art streaming algorithms for hypergraph partitioning. We show performance profiles, improvement plots over Hashing, and boxplots. Note that PaToH-con, PaToH-cut, and Hashing align almost perfectly with the y-axis in Figures 4.14b, 4.14d, and 4.14f, respectively. Also the curves and bars of MM-N2P and MM-L5 roughly overlap with one another in Figure 4.14d and Figure 4.14c.

which is referred to as RRS(0.15) from now on. Since Min-Max-L ℓ is not dominated by any other algorithm, we exhibit its results with $\ell = 5$, as seen in the best results in [9], and we refer to it as MM-L5 from this point.

Connectivity. We start by looking at the connectivity metric. In Figure 4.14a, we plot the average connectivity improvement over Hashing for each value of k . PaToH-con produces the best connectivity on average, yielding an average improvement of 443% when compared to Hashing. This is in line with previous works in the area of (hyper)graph partitioning, i.e. streaming algorithms typically compute worse solutions than internal memory algorithms, which have access to the whole graph. FREIGHT-con is found to be the second best algorithm in terms of connectivity, outperforming both the internal memory algorithm HYPE and the buffered streaming algorithm RRS(0.15). On average, these three algorithms improve 194%, 171%, and 136% over Hashing, respectively. Finally, MM-N2P and MM-L5 compute solutions which improve 111% and 96% over Hashing on average, respectively. In direct comparison, FREIGHT-con shows average connectivity improvements of 8%, 24%, 39%, and 50% over HYPE, RRS(0.15), MM-N2P, and MM-L5, respectively. Note that each algorithm retains its relative ranking in terms of average connectivity over all values of k .

In Figure 4.14b, we plot connectivity performance profiles across all experiments. Note that PaToH-con produces the best overall connectivity for 96.4% of the instances, while FREIGHT-con produces the best connectivity for 3.1% of the instances and no other algorithm computes the best connectivity for more than 0.35% of the instances. The connectivity produced by FREIGHT-con, HYPE, RRS(0.15), MM-N2P, MM-L5, and Hashing are within a factor 2 of the best found connectivity for 67%, 61%, 47%, 41%, 34%, and 9% of the instances, respectively. In summary, FREIGHT-con produces the best connectivity among (buffered) streaming competitors, outperforming even in-memory algorithm HYPE.

Cut-Net. Next we examine at the cut-net metric. In Figure 4.14c, we plot the cut-net improvement over Hashing. PaToH-cut produces the best overall cut-net, with an average improvement of 100% compared to Hashing. FREIGHT-cut is found to be the second best algorithm with respect to cut-net, superior to internal-memory algorithm HYPE and buffered streaming algorithm RRS(0.15). These three algorithms improve connectivity over Hashing by 37%, 30%, and 17% respectively. Finally, both MM-N2P and MM-L5 improve connectivity by 13% on average over Hashing. In direct comparison, FREIGHT-cut shows average connectivity improvements of 6%, 18%, 22%, and 22% over HYPE, RRS(0.15), MM-N2P, and MM-L5, respectively. Each algorithm preserves its relative ranking in average cut-net across all values of k .

In Figure 4.14d, we plot cut-net performance profiles across all experiments. In the plot, PaToH-cut produces the best overall connectivity for 98.0% of the instances, while FREIGHT-cut and HYPE produce the best cut-net for 6.8% and 5.2% of the instances and all other streaming algorithms (RRS(0.15), MM-N2P, MM-L5, and Hashing) produce the best cut-net for 4.8% of the instances. The cut-net results produced by FREIGHT-cut, HYPE, RRS(0.15), MM-N2P, MM-L5, and Hashing are within a factor 2 of the best found cut-net for 83%, 79%, 69%, 66%, 66%, and 58% of the instances, respectively. This shows that FREIGHT-cut produces the best cut-net among all (buffered) streaming competitors and even beats the in-memory algorithm HYPE.

Running Time. Now we compare the algorithms’ runtime. Boxes and whiskers in Figure 4.14e display the distribution of the running time per pin, measured in nanoseconds, for all instances. Hashing, FREIGHT-cut, and FREIGHT-con are the three fastest algorithms, with median runtimes per pin of 15ns, 38ns, and 41ns, respectively. MM-L5, MM-N2P, HYPE, and RRS(0.15) follow with median runtimes per pin of 130ns, 437ns, 792ns, and 833ns, respectively. Lastly, the algorithms with the highest median runtime per pin are PaToH-cut and PaToH-con, with 2 516ns and 3 333ns respectively. The measured runtime per pin for both HYPE and PaToH align with values reported in prior research [141].

In Figure 4.14f, we show running time performance profiles. Hashing is the fastest algorithm for 98.3% of the instances, while FREIGHT-cut is the fastest one for 1.2% of the instances and no other algorithm is the fastest one for more than 0.4% of the instances. The running time of FREIGHT-cut and FREIGHT-con is within a factor 4 of that of Hashing for 82% and 72% of instances, respectively. In contrast, for only 16% of instances does this occur for MM-L5, and for less than 0.4% of instances for all other algorithms. The close running times of FREIGHT to Hashing are surprising given FREIGHT’s superior solution quality compared to Hashing and all other streaming algorithms and even HYPE.

Further Comparisons. For graph partitioning FREIGHT and Fennel are mathematically equivalent. However, FREIGHT exhibits a lower computational complexity of $O(m + n)$ compared to the standard implementation of Fennel, which has a complexity of $O(m + nk)$ due to evaluating all blocks for each node. To optimize its performance for this use case, we have implemented an optimized version of FREIGHT with a memory consumption of $O(n + k)$, matching that of Fennel. In our experiments, we utilized the same graphs as in [11] and tested with $k \in \{512, 1024, 1536, 2048, 2560\}$. On average, FREIGHT proves to be 109 times faster than the standard implementation of Fennel. Moreover, the performance gap is found to increase as the value of k grow, with FREIGHT reaching up to 261 times faster than Fennel in some instances.

4.4 Experimental Comparison

In this section, we present an unpublished experimental comparison of our three streaming algorithms. We performed experiments for $k = 8, 16, 32, 64, 128, 256$ on the *huge graphs* listed in Table 4.4 using Machine B, which is a relatively modest machine. As in Section 4.1.2.3, we did not repeat each test multiple times with different seeds. As mentioned in Section 4.1.2.3, we also ran Metis and KaHIP on these graphs, but they failed on all instances as they require more memory than the machine has.

Like in Section 4.1.2.3, we refer to setups of HeiStream with specific buffer sizes as HeiStream(Xk), where a buffer contains $X \times 1024$ nodes. Apart from HeiStream, we ran Online Recursive Multi-Section (referred to as nh-OMS) and FREIGHT. In particular, we use here the best parameters of nh-OMS found in the tuning experiments in Section 6.1.2.1. Table 4.4 provides detailed per-instance results with HeiStream using large buffer sizes capable of running on Machine B. We have excluded the IO delay of loading the input graph from the disk in Table 4.4 as it depends on the disk and is roughly the same regardless of the used partitioning algorithm. The delay (in seconds) is reported in Section 4.1.2.3.

The results indicate that HeiStream performs better than its competitors in terms of solution quality for most instances. Notably, it produces partitions with significantly lower edge-cut compared to our other streaming algorithms for four tested graphs, namely uk-2005, sk-2005, uk-2007-05, and RHG1. For social networks soc-friendster and twitter7, HeiStream outperforms all other algorithms, but the improvement over nh-OMS and FREIGHT is not as substantial as in other cases. However, there is an exception on the RHG2 network, where HeiStream produces edge-cut values below 0.7%, but FREIGHT performs better and nh-OMS produces even lower edge-cut values. Overall, FREIGHT and nh-OMS produce comparable quality results, with FREIGHT being slightly better on three graphs and nh-OMS being slightly better on five graphs.

In terms of running time, HeiStream is considerably slower than nh-OMS and FREIGHT. Although the runtime complexity of HeiStream is $O(m+n)$, which is not worse than that of nh-OMS and FREIGHT, its constant factors are larger due to more complex computations. It is also worth noting that the running time of nh-OMS increases with increasing k , while the running time of FREIGHT remains roughly constant for any k . This is expected because the running time of nh-OMS is $O((m+n) \log k)$, while the running time of FREIGHT for graphs is $O(m+n)$. As a result, FREIGHT becomes faster than nh-OMS as k increases.

Graph	k	HeiStream(Xk)			nh-OMS		FREIGHT	
		X	CE(%)	RT(s)	CE(%)	RT(s)	CE(%)	RT(s)
uk-2005	8	1024	4.03	290.23	22.08	26.96	23.87	20.07
	16	1024	6.01	300.04	25.40	27.85	26.79	21.49
	32	1024	7.65	310.72	28.34	34.01	30.15	21.16
	64	1024	8.99	322.14	30.58	36.52	32.10	21.89
	128	1024	9.94	346.73	32.05	41.89	33.31	22.74
	256	1024	10.68	386.64	33.71	44.50	34.35	23.51
twitter7	8	512	41.64	1727.13	60.14	178.59	55.13	173.88
	16	512	47.04	1774.92	72.74	182.74	61.73	184.13
	32	512	52.59	1884.16	81.78	187.33	65.07	179.02
	64	512	57.53	1988.11	87.64	205.12	76.78	183.84
	128	512	61.87	2113.34	92.06	228.06	81.20	185.98
	256	512	65.47	2357.92	94.28	249.24	83.39	188.20
sk-2005	8	1024	3.23	634.79	20.86	40.10	23.89	34.13
	16	1024	4.11	648.48	25.69	43.63	28.87	33.97
	32	1024	5.32	667.84	28.71	49.29	31.90	35.58
	64	1024	7.55	695.60	32.59	53.15	35.20	35.87
	128	1024	8.95	733.05	35.65	61.62	38.59	37.56
	256	1024	12.02	798.73	39.64	64.80	43.15	38.58
soc-friendster	8	1024	27.36	4099.35	37.68	362.80	35.54	383.06
	16	1024	34.50	4202.04	52.70	393.57	50.06	383.60
	32	1024	39.52	4345.96	64.99	426.01	54.70	374.97
	64	1024	46.35	4546.98	76.39	427.77	57.88	388.27
	128	1024	52.41	4796.56	80.88	481.40	60.59	394.40
	256	1024	57.79	5323.08	83.32	477.39	63.32	408.62
er-fact1.5s26	8	1024	73.27	2216.99	74.33	233.06	73.44	231.99
	16	1024	80.18	2292.12	81.59	255.07	80.40	243.52
	32	1024	84.36	2400.35	86.27	293.15	84.63	255.88
	64	1024	86.99	2534.09	89.05	316.62	87.31	253.72
	128	1024	88.72	2725.81	90.82	374.28	89.10	248.61
	256	1024	89.99	2913.95	91.80	393.10	90.45	261.57
RHG1	8	1024	0.04	380.04	1.89	58.12	2.02	44.17
	16	1024	0.06	391.63	1.99	60.36	2.12	44.83
	32	1024	0.09	406.56	2.02	70.22	2.16	46.87
	64	1024	0.15	435.71	2.05	75.38	2.17	45.09
	128	1024	0.22	482.06	2.07	92.34	2.18	47.09
	256	1024	0.34	569.77	2.09	92.16	2.19	46.92
RHG2	8	1024	0.09	621.56	0.04	73.03	0.05	57.02
	16	1024	0.13	632.61	0.04	72.78	0.08	57.87
	32	1024	0.19	648.68	0.04	84.19	0.12	57.34
	64	1024	0.29	674.36	0.05	92.05	0.18	57.78
	128	1024	0.44	727.66	0.06	97.14	0.27	56.64
	256	1024	0.68	816.60	0.08	105.69	0.44	61.23
uk-2007-05	8	1024	0.54	1024.26	23.04	76.25	25.28	62.97
	16	1024	0.60	1045.36	26.69	82.78	28.18	65.81
	32	1024	0.70	1058.73	28.14	94.47	29.46	66.08
	64	1024	0.92	1099.64	29.74	103.21	29.92	66.86
	128	1024	1.31	1163.45	30.80	118.08	30.63	67.78
	256	1024	1.95	1280.48	32.10	125.89	31.54	70.58

Table 4.4. Our algorithms compared. CE and RT denote cut edges and running time.

4.5 Conclusion

In this chapter, we proposed three streaming algorithms for (hyper)graph decomposition: HeiStream, Online Recursive Multi-Section, and FREIGHT. All were discussed in detail and were subjected to extensive experimental evaluation against the state-of-the-art. Lastly, we presented an experimental comparison between them for the graph partitioning problem.

We proposed HeiStream, a buffered streaming graph partitioning algorithm. It combines the buffered streaming model with multilevel graph partitioning techniques and an extension of Fennel to a multilevel algorithm. Compared to the previous state-of-the-art, HeiStream computes significantly better solutions while at the same time being faster in many cases. An important property of HeiStream is that its running time does not depend on the number of blocks, while the previous state-of-the-art streaming partitioning algorithms have running time almost proportional to this number of blocks.

We proposed Online Recursive Multi-Section, a streaming algorithm to compute hierarchical partitionings of graphs. In terms of time complexity our algorithm outperforms previous state-of-the-art one-pass streaming algorithms for graph partitioning while also implicitly optimizing process mapping objectives. To the best of our knowledge, this is the first streaming algorithm for the process mapping problem. We present extensive experimental results in which compare it against the previous state-of-the-art algorithms for streaming non-buffered one-pass graph partitioning, and show that we can speed it up even further with a multi-threaded parallelization. Experiments show that our algorithm is up to two orders of magnitude faster than the previous state-of-the-art while producing solutions with better communication cost and slightly worse edge-cut. Moreover, our algorithm is only three times slower than Hashing when running on 32 threads while computing significantly better results.

We proposed FREIGHT, a streaming algorithm for hypergraph partitioning. Our algorithm leverages an optimized data structure, resulting in linear running time with respect to pin-count and linear memory consumption in relation to the numbers of nets and blocks. The results of our extensive experimentation demonstrate that the running time of FREIGHT is competitive with the Hashing algorithm, with a maximum difference of a factor of four observed in three fourths of the instances. Importantly, our findings indicate that FREIGHT consistently outperforms all existing (buffered) streaming algorithms and even the in-memory algorithm HYPE, with regards to both cut-net and connectivity measures. This underscores the significance of our proposed algorithm as a highly efficient and effective solution for hypergraph partitioning in the context of large-scale and dynamic data processing.

We presented an unpublished experimental comparison of our three streaming algorithms. Although each of them is designed for a specific problem, they are all capable of solving the graph partitioning problem, so we compare them for this problem. In our evaluation HeiStream produces the lowest edge-cut while consuming the most runtime. Our results show that, in terms of graph partitioning, FREIGHT and Online Recursive Multi-Section produce comparable solution quality but FREIGHT becomes considerably faster than Online Recursive Multi-Section for larger numbers of blocks.

In future work, we intend to parallelize HeiStream and extend it to solve other (hyper)graph decomposition problems such as hypergraph partitioning and process mapping. While Online Recursive Multi-Section is already useful for a wide-range of applications that need (hierarchical) partitions very fast, in future work, we plan to parallelize the algorithm in the distributed memory model and want to port it to GPUs. Moreover, we intend to parallelize FREIGHT and try to make it even more effective by adapting its mathematical formulation to more accurately encode the connectivity gains involved.

Chapter 5

Local Algorithms

In this chapter, we propose our contributions in the field of local graph decomposition. Our contributions comprise the detailed design of two algorithms, namely, LMCHGP and SOCIAL, for solving the local motif clustering problem. LMCHGP constructs a (hyper)graph model centered around the seed node in such a way that optimizing for conductivity in the model is equivalent to optimizing for motif conductance in the original graph. The model is then partitioned utilizing high-quality (hyper)graph partitioning techniques available in the literature. On the other hand, SOCIAL builds the same hypergraph model around the seed node but converts it into a flow network that guarantees exact quality improvement. The flow network is then recursively solved using a standard max-flow algorithm until a locally optimal cluster is obtained. In our experiments, both LMCHGP and SOCIAL outperform the state-of-the-art by producing clusters with lower motif conductance value while being significantly faster, up to multiple orders of magnitude.

References. This chapter is based on [142] and [143], which is joint work with Adil Chhabra and Christian Schulz, and on the technical report [144], which is also joint work with Adil Chhabra and Christian Schulz and is currently under submission.

5.1 Local Motif Clustering

Within this section, we propose two novel algorithms to solve the local motif clustering problem using sophisticated combinatorial algorithms.

As a first contribution, we propose LMCHGP (Local Motif Clustering via (Hyper)Graph Partitioning). This is an algorithm based on (hyper)graph partitioning which has two versions: one based on a graph model and the other one based on hypergraph model. To begin, our algorithm constructs a (hyper)graph model that represents the distribution of motifs around the seed node on the original graph. While the graph model accurately captures motifs of

size at most three, the hypergraph model accommodates arbitrary motifs and is designed to minimize motif conductance in the original network. We then partition the (hyper)graph model using a powerful multi-level hypergraph or graph partitioner to directly minimize the motif conductance of the corresponding partition in the original graph.

As a second contribution, we propose an algorithm named SOCIAL (*faSter mOtif Clustering vIa mAximum fLows*). This algorithm optimizes for motif conductance by combining the strongly local hypergraph model proposed in Chapter 5.1 with an adapted version of the fast and effective algorithm *max-flow quotient-cut improvement*(MQI)[30]. Using the same hypergraph model as our previous algorithm, we construct a flow model in which certain cuts correspond one-to-one with subsets of the initial cluster that include the seed node and have lower motif conductance than that of the whole cluster. We then utilize a push-relabel algorithm to find such a cut and recursively repeat the process, or prove that the current cluster is optimal among all its sub-clusters containing the seed node.

We extensively evaluate our proposed algorithms LMCHGP and SOCIAL against the state-of-the-art. In experiments involving triangle motifs, LMCHGP and SOCIAL compute clusters with a motif conductance value lower than the state-of-the-art, while also being up to multiple orders of magnitude faster.

5.1.1 LMCHGP

5.1.1.1 Overall Strategy

Given a graph $G = (V, E)$, a seed node u , and a motif μ , our strategy for local clustering is based on four consecutive phases. First, we select a set $S \subseteq V$ containing u and close-by nodes. From now on, we refer to this set S as a *ball around u* . Second, we enumerate the collection M of occurrences of the motif μ which contain at least one node in S . Next, we build a graph or a hypergraph model H_μ depending on the configuration of the algorithm. In particular, we design H_μ in such a way that the motif-conductance metric in G can be computed directly in H_μ . Then, we partition this model into two blocks using a high-quality (hyper)graph partitioning algorithm. The obtained partition of H_μ is directly translated back to G as a local cluster around the seed node. Figure 5.1 provides a comprehensive illustration of the consecutive phases of our algorithm. Note that (hyper)graph partitioning algorithms do not optimize for traditional clustering objectives such as conductance. Instead, they aim at minimizing the edge-cut (resp. cut-net) value while respecting a hard balancing constraint. To improve for the correct objective, we repeat the partitioning phase β times with different imbalance constraints and pick the clustering with best motif conductance. Especially for the graph-based version of our model H_μ , we subsequently run a special label propagation for each of these β iterations in order to increase the chances of reaching a local minimum

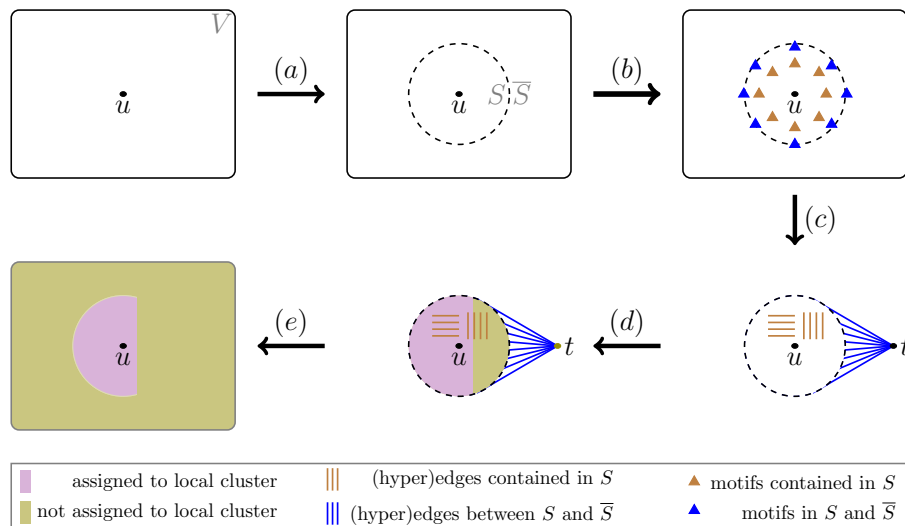


Figure 5.1. Illustration of the phases of our algorithm. (a) Given a seed node u and a graph G , a ball S around u is selected. (b) Motif occurrences of μ with at least a node in S are enumerated. (c) The (hyper)graph model H_μ is built by converting motifs into (hyper)edges and contracting \bar{S} into a node r . (d) The model is partitioned into two blocks using a multi-level (hyper)graph partitioner. (e) The partition of H_μ is converted in a local cluster around the seed node in G .

motif conductance. Moreover, the first three phases of our strategy are repeated α times with different balls around the seed node in order to better explore the vicinity of the seed node in the original graph. Our overall strategy including the mentioned repetitions is outlined in Algorithm 5.1.

5.1.1.2 Ball around Seed Node

Our approach to select S is a fixed-depth breadth-first search (BFS) rooted on u . More specifically, we compute the first ℓ layers of the BFS tree rooted on u , then we include all its nodes in S . For each of the α repetitions of our overall algorithm, we use different amounts ℓ of layers for a better algorithm exploration. Two exceptional cases are handled by our algorithm, namely a ball S that is either too small or disconnected from \bar{S} . We avoid the first exceptional case by ensuring that S contains 100 or more nodes in at least one repetition of our overall algorithm. More specifically, in case this condition is not automatically met, then we accomplish it in the last repetition by growing additional layers in our partial BFS tree while it contains fewer than 100 nodes. The number 100 is based on the findings of Leskovec et al. [19], which show that most well characterized communities from real-world graphs have a relatively small size, in the order of magnitude of 100 nodes. If the second

Algorithm 5.1 Local Motif Clustering via (Hyper)Graph Partitioning**Input** graph $G = (V, E)$; seed node $u \in V$; motif μ **Output** cluster $C^* \subseteq V$

```

1:  $C^* \leftarrow \emptyset$ 
2: for  $i = 1, \dots, \alpha$  do
3:   Select ball  $S$  around  $u$ 
4:    $M \leftarrow$  Enumerate motifs in  $S$ 
5:   Build (hyper)graph model  $H_\mu$  based on  $S$  and  $M$ 
6:   for  $j = 1, \dots, \beta$  do
7:     Partition model  $H_\mu$  into  $(C, \bar{C})$ , where  $u \in C$ 
8:     if  $C^* = \emptyset \vee \phi_\mu(C) < \phi_\mu(C^*)$  then
9:        $C^* \leftarrow C$ 
10: Convert  $C^*$  into a local motif cluster in  $G$ 

```

exceptional case happens, it means that the whole BFS tree rooted on the seed node has at most ℓ layers. In this case, we simply stop the algorithm and return the entire ball S , which corresponds to an optimal community with motif conductance 0 provided that there is at least one motif in S .

The approach described above makes sure that there is a reasonable chance that a well characterized community containing u is contained in S , since it has at least 100 nodes [19] which are all very close to u . This likelihood is further increased due to the multiple repetitions of our overall algorithm using balls S of different sizes. Our BFS approach to select S can be executed in time linear on the subgraph induced in G by the closed neighborhood $N[S]$ of S . After selecting S , the further phases of our algorithm do not deal with the whole graph, but exclusively with S , its edges, and its motif occurrences. As a consequence, our algorithm operates on a much smaller problem dimension than the size of input graph G , hence its running time corresponds to the same smaller problem dimension. The number α of repetitions as well as the amount ℓ of layers used in each repetition are tuning parameters.

5.1.1.3 Motif Enumeration

We now describe and discuss the motif-enumeration phase of our algorithm. We optimally solve it for the triangle motif in time roughly linear on the size of the subgraph induced in G by the closed neighborhood $N[S]$ of S . Moreover, we show that there are good heuristics approaches to enumerate higher-order motifs efficiently. The general problem of finding out if a given motif is a subgraph of some graph is NP-hard [145], hence the enumeration of all such motifs is also NP-hard. Nevertheless, some simpler motifs can be enumerated in polynomial time, which is the case for the triangle motif [146]. Triangles, which can

be defined as cycles or cliques of length three, have a wide variety of relevant applications on network analysis and clustering [147–149]. Without loss of generality, we specifically focus on the triangle motif within our algorithm. Nevertheless, note that many other small motifs can also be polynomially enumerated, such as small (directed and undirected) paths and cycles. Moreover, our overall algorithm can also be adapted for more arbitrary motifs if we relax the optimality of the enumeration, which can be done using efficient heuristics such as the one proposed by Kimmig et al. [150]. A simple and exact algorithm for triangle enumeration was proposed by Chiba and Nishizeki [151]. Roughly speaking, this algorithm works by intersecting the neighborhoods of adjacent nodes. For each node v , the algorithm starts by marking its neighbors with degree smaller than or equal to its own degree. For each of these specific neighbors of v , it then scans its neighborhood and enumerates new triangles as soon as marked nodes are found. The running time of this algorithm is $O(ma) = O(m^{\frac{3}{2}})$, where a is the arboricity of the graph. For the motif-enumeration phase of our algorithm, we apply the algorithm of Chiba and Nishizeki [151] only on the subgraph induced in G by $N[S]$. This is enough to find all triangles containing at least one node in S , as exemplified by transformation (a) in figures 5.2 and 5.3. Assuming a constant-bounded arboricity, the overall cost of our motif-enumeration phase for triangles is $O(|N[S] \times N[S] \cap E|)$.

5.1.1.4 Hypergraph Model

In this section, we conceptually describe the hypergraph version of our model H_μ and explain how to build it. We show that it can be constructed in time linear on the amount of nodes in S and motifs in M . We also discuss advantages and limitations of our hypergraph model approach.

Our hypergraph model H_μ is built in two conceptual operations. First, define a hypergraph containing V as nodes and a set E of nets such that, for each motif in M , E has a net with pins equal to the endpoints of this motif. Then, we contract together all nodes in \bar{S} into a single node r and substitute parallel nets by a single net whose weight is equal to the summed weights of the removed parallel nets. More formally, we define the hypergraph version of our model as $H_\mu = (S \cup \{r\}, E)$ where the set E of nets contains one net e associated with each motif occurrence $G' = (V', E') \in M$ such that $e = V'$ if $V' \subseteq S$, and $e = V' \cap S \cup \{r\}$ otherwise. In the former case the net has weight 1, in the latter case the net has weight equal to the amount of motif occurrences in M represented by it. In a typical hypergraph contraction, the weight of r would be $c(r) = c(\bar{S})$. Nevertheless, we opt to make the nodes of H_μ unweighted, which is convenient for the purpose of our overall algorithm as will become clear later. In practice, the hypergraph version of H_μ can be built by instantiating the nodes in $S \cup \{r\}$ and the nets in E . Assuming that the number of nodes in μ is a constant, our model is built in time $O(|S| + |M|)$ and uses memory $O(|S| + |M|)$.




Figure 5.2. Example of motif-enumeration and model-construction phases of our algorithm for triangle motif and the hypergraph model. In the left, the nodes of G are split into sets S and \bar{S} . In the center, motif occurrences containing nodes in S are enumerated. In the right, H_μ is built by converting motifs into nets and contracting \bar{S} into a node r .

The construction of H_μ is illustrated in transformation (c) of Figure 5.1 and demonstrated for a particular example in transformation (b) of Figure 5.2.

Observing the relationship between G and H_μ , we can distinguish three groups of components. The first group comprises nodes in S and motifs with all endpoints in S , all of which are represented in H_μ without any contraction as nodes and nets. The second group consists of nodes in \bar{S} and motifs with all endpoints in \bar{S} , which are compactly represented in H_μ as the contracted node r . The third group comprises motifs with nodes in both S and \bar{S} , all of which are abstractly represented in H_μ as nets containing individual pins in S as well as the pin r . Summing up, our hypergraph model is a concise representation of the whole graph G where relevant information for local motif clustering is emphasized in two perspectives: Edges are omitted while motifs are made explicit and global information is abstracted while local information is preserved in detail. Theorem 5.1 shows that the cut-net of a partition of H_μ directly corresponds to the motif-cut of an equivalent partition in G if our motif enumeration step is exact. On top of that, Theorem 5.2 shows that the motif conductance of this equivalent partition of G can be directly computed from H_μ assuming $d_\mu(S) \leq d_\mu(\bar{S})$. Assuming $d_\mu(S) \leq d_\mu(\bar{S})$ is fair since S is ideally much smaller than \bar{S} . Enumerating the motifs in \bar{S} is not reasonable for a local clustering algorithm, but we did verify that our assumption holds during all our experiments.

Theorem 5.1. *Every k -way partition P of the hypergraph model H_μ corresponds to a distinct k -way partition P' of G , such that the cut-net of P is equivalent to the motif-cut of P' , assuming an exact motif enumeration step.*

Proof. For simplicity, we prove the claim assuming that parallel nets are not substituted by a single net whose weight is equal to their summed weights. This proof directly extends to our model since the contribution of a contracted cut net to the overall cut-net equals the contribution of the parallel nets represented by it. Due to the design of our hypergraph model H_μ , there is a direct correspondence between its nodes and the nodes of G . Hence, any partition P of H_μ corresponds to a partition P' of G where corresponding nodes are simply assigned to the same blocks. Since \bar{S} is represented by the single node r in H_μ , no motif occurrence totally contained in \bar{S} can be cut in P' . All the remaining motif occurrences in G can be potentially cut in P' , but these motif occurrences are bijectively associated with the nets of H_μ with a direct correspondence between motif endpoints in G and net pins in H_μ . As a consequence, a motif occurrence of G is cut in P' if, and only if, the corresponding net in H_μ is cut in P . \square

Theorem 5.2. *Given a 2-way partition $P = (C, \bar{C})$ of our hypergraph model H_μ with $r \in \bar{C}$, the motif conductance $\phi_\mu(C')$ of the corresponding 2-way partition $P' = (C', \bar{C}')$ of G is the ratio of the cut-net of P to $d_w(C)$, assuming an exact motif enumeration step and $d_\mu(S) \leq d_\mu(\bar{S})$.*

Proof. From Theorem 5.1, the motif-cut of P' can be substituted by the cut-net of P in the numerator of the definition of $\phi_\mu(C')$. To complete the proof, it suffices to show that the denominator of $\phi_\mu(C')$, namely $\min(d_\mu(C'), d_\mu(\bar{C}'))$, is equal to $d_w(C)$. Due to the design of H_μ , the values of $d_\mu(C')$ and $d_w(C)$ are identical. Our assumption $r \in \bar{C}$ leads to $\bar{S} \subseteq \bar{C}'$ and $C' \subseteq S$, which respectively imply $d_\mu(\bar{S}) \leq d_\mu(\bar{C}')$ and $d_\mu(C') \leq d_\mu(S)$. Since $d_\mu(S) \leq d_\mu(\bar{S})$, hence $d_w(C) = d_\mu(C') \leq d_\mu(S) \leq d_\mu(\bar{S}) \leq d_\mu(\bar{C}')$. \square

5.1.1.5 Graph Model

In this section, we describe the graph version of our model H_μ and explain how to build it. Similarly to the hypergraph version of this model, our graph model can be built in time linear on the amount of nodes in S and motifs in M . We discuss advantages and limitations of the graph model in comparison to the hypergraph model.

The graph version of our model H_μ is built in two conceptual operations. The first one consists of obtaining the weighted graph W proposed by Benson et al. [31] and used in the state-of-the-art algorithm MAPPR [23]. The graph W contains V as nodes and a set of edges such that two conditions are met: (i) there is an edge between a pair of nodes if, and only if, both nodes belong at the same time to at least a motif in G ; (ii) the weight of an edge is equal to the number of motif occurrences containing both its endpoints. The second operation consists of contracting all nodes in \bar{S} into a single node r and substitute parallel edges by a single edge whose weight is equal to the summed weights of the removed edges. The




Figure 5.3. Example of motif-enumeration and model-construction phases of our algorithm for triangle motif and the graph model. In the left, G is shown with its nodes split into the sets S and \bar{S} . In the center, the motif occurrences containing at least a node in S are enumerated and the weight of an edge equals the number triangles it touches. In the right, the model H_μ is built by contracting \bar{S} into a single node r . The weight of an edge is represented by its thickness.

construction of the graph version of H_μ is illustrated in transformation (c) of Figure 5.1 and demonstrated for a particular example in transformation (b) of Figure 5.3. More formally, we define the model as $H_\mu = (S \cup \{r\}, E_\mu)$ where E_μ contains an edge e for each pair of nodes sharing a motif $G' = (V', E')$ provided that at least one of its endpoints is contained in S . The weight of r is set to $c(r) = c(\bar{S})$. Similarly to our approach with the hypergraph version of the model, we opt to make the nodes of the graph model unweighted in our experiments. In practice, the graph version of H_μ can be built by instantiating the nodes in $S \cup \{r\}$ and directly computing the edges in E_μ and their weights. Assuming that the number of nodes in μ is a constant, our graph model is built in time $O(|S| + |M|)$ and uses memory $O(|S| + |E_\mu|)$. Especially for the triangle motif, the memory requirement of the graph model is $O(|S| + |N[S] \times N[S] \cap E|)$, which is linear on n and m in the worst case.

We reproduce here Theorem 5.3 by Yin et al. [23], which shows that conductance in the weighted graph W is equivalent to motif conductance in G as long as the motif has at most 3 nodes. Based on this result, Theorem 5.4 shows that we can compute motif conductance directly from our graph model H_μ if $d_\mu(S) \leq d_\mu(\bar{S})$. As we mentioned, assuming $d_\mu(S) \leq d_\mu(\bar{S})$ is fair since S is ideally much smaller than \bar{S} . Recall that the hypergraph version of our model H_μ is flexible enough to represent any motif as a net such that Theorems 5.1 and 5.2 continue valid. Although the graph version of our model can technically represent any motif, Theorem 5.4 is only valid for motifs with at most three nodes, while other models only allow a heuristic computation of the motif conductance [31]. Nevertheless, the biggest drawback

of the hypergraph-based approach is the need for storing up to $|M|$ nets, which costs $O(n^3)$ in the worst case. In contrast, the memory needed to store our graph model is $O(n^2)$ in the worst case and $O(n + m)$ specifically for the triangle motif.

Theorem 5.3 (Theorem 4.1 by Yin et al. [23]). *Given a 2-way partition $P'' = (C'', \overline{C}'')$ of the weighted graph W , the motif conductance $\phi_\mu(C')$ of the corresponding 2-way partition $P' = (C', \overline{C}')$ in G is equal to the conductance $\phi(C'')$ of C'' in W , assuming a motif with at most three nodes.*

Theorem 5.4. *Given a 2-way partition $P = (C, \overline{C})$ of the graph version of model H_μ with $r \in \overline{C}$, the motif conductance $\phi_\mu(C')$ of the corresponding 2-way partition $P' = (C', \overline{C}')$ of G is the ratio of the edge-cut of P to $d_\omega(C)$, assuming an exact motif enumeration step, $d_\mu(S) \leq d_\mu(\overline{S})$, and a motif with at most three nodes.*

Proof. From Theorem 5.3, the motif conductance of any 2-way partition of G is equal to the conductance of the equivalent partition in W assuming a motif with at most 3 nodes. Since we assume $d_\mu(S) \leq d_\mu(\overline{S})$, hence the conductance in W of any community $C'' \in S$ is equal to its edge-cut divided by the volume of C'' . From the construction of our graph model H_μ , the assumed community C has an equivalent community C'' in W with same edge-cut and same volume, which completes the proof. \square

5.1.1.6 Partitioning

In this section, we describe the (hyper)graph partitioning phase of our local motif clustering algorithm. We present the used (hyper)graph partitioning algorithms and discuss how we enforce feasibility of the found solution and maximize its quality. Moreover, we provide remarks about the running time of our partitioning phase.

The partitioning phase of our algorithm consists of a 2-way partitioning of H_μ . When using the hypergraph model, the partition is computed by the multi-level hypergraph partitioner KaHyPar [52]. When using the graph model, the partition is computed by the multi-level graph partitioner KaHIP [108]. These partitioners contain sophisticated algorithms to produce low-cut partitions of (hyper)graphs efficiently. As already shown, any 2-way partition of H_μ automatically corresponds to a community in G . Nevertheless, our aim is to obtain a *consistent* partition of H_μ , which we define as a partition where the seed node u and the contracted node r are in different blocks. This consistent partition ultimately corresponds to a local community in G which contains the seed node u and is completely contained in the ball S . This *consistency* criterion is important since the nodes in \overline{S} have not been explored by our algorithm and are farther from u than the nodes in S . Note that the 2-way partition illustrated in Figure 5.1 generates a consistent local community according to our definition of it. While KaHyPar allows partitioning hypergraphs with fixed nodes, KaHIP does not

offer such functionality. Nevertheless, we ensure block feasibility for both versions of our algorithm by simply assigning the seed node to the block that does not contain r after the partition is computed (before computing the motif conductance). Although simplistic, this approach has not affected solution quality considerably for the hypergraph-based version of our algorithm in preliminary comparisons against a fixed nodes-based approach.

Besides corresponding to a consistent local clustering on the original graph, our solution should have as low a motif conductance as possible. However, KaHyPar and KaHIP are randomized algorithms and do not directly optimize for this objective, but rather minimize the cut value while enforcing a hard balancing constraint. To improve our results, we explore different combinations of edge-cut (resp. cut-net) and imbalance by repeating the partitioning procedure β times with random balancing constraints for each built (hyper)graph model, where β is a tuning parameter. For each obtained partition, our algorithm computes the motif conductance of the corresponding local cluster as shown in Theorem 5.4 (resp. Theorem 5.2) and keeps the partition with the best motif conductance.

KaHyPar and KaHIP run in time close to linear in practice. Nevertheless, the partitioning phase can be the dominating operation of our overall local clustering algorithm. This is the case because KaHyPar and KaHIP use sophisticated algorithms and data structures in order to minimize the cut value, which increases constant factors in the algorithm running time complexity. KaHyPar can be especially much slower than KaHIP since the number of nets in the hypergraph model can be considerably larger than the number of edges in the graph model. We can make the hypergraph version of our algorithm faster by using Mt-KaHyPar [152] instead of KaHyPar. Mt-KaHyPar obtains significant speedups by using allowing shared-memory parallel execution. In both cases, running time can be improved further by using parallelized tools such as Mt-KaHyPar [152] or KaMinPar [153]. However, parallelization is not the focus of this thesis.

Local Search. We implement a local search inspired by *label propagation* [140] for the graph model-based version of our algorithm. This local search is designed to optimize for the correct objective, i.e., motif conductance. We apply it directly on each 2-way partition generated by KaHIP in order to increase the chances of reaching a local minimum motif conductance. The local search algorithm works in rounds. In each round, it visits all nodes of H_μ in a random order, starting with the labels being the current assignment of nodes to blocks. When a node v is visited, it is moved to the opposite block if this movement causes a decrease in the motif conductance of the clustering. Movements of nodes with zero-gain can occasionally occur with 50% probability. We ensure that the seed node u and the contracted node r continue in opposite blocks by simply skipping them. We stop the local search when a local optimum is reached or after at most ℓ rounds, where ℓ is a tuning parameter.

5.1.2 SOCIAL

5.1.2.1 Overall Strategy

Given a graph $G = (V, E)$, a seed node u , and a motif μ , our strategy for local clustering is based on the following phases. First, we select a set $S \subseteq V$ containing u and close-by nodes. As in Section 5.1.1.1, we refer to this set S as a *ball around u* . Second, we enumerate the collection M of occurrences of the motif μ which contain at least one node in S . Third, we build a hypergraph model H_μ in such a way that the motif-conductance of any cluster $C \subseteq S$ in G can be computed directly in H_μ . Fourth, we set $C_0 = S$ as our initial cluster and use it to build our MQI-based [30] flow model G_f from the hypergraph model H_μ . Fifth, we use G_f to either find a new cluster $C \subset C_0$ containing u with strictly smaller motif conductance than C_0 or prove that such cluster does not exist. While $C \subset C_0$ is found, we take it as our new initial cluster, rebuild G_f , and repeat the previous phase. When eventually no such strict sub-set is found, the best obtained cluster is directly translated back to G as a local cluster around the seed node. Figure 5.4 provides a comprehensive illustration of the consecutive phases of SOCIAL. Note that there is no guarantee of finding the best overall cluster including u strictly contained in S . Instead, we find a succession of clusters with strictly decreasing cardinality and motif conductance until a local optimum is reached. To better explore the vicinity of u in G and overcome the fact we only find clusters inside S , we repeat the overall strategy α times with distinct balls S . Our overall algorithm including the mentioned repetitions is outlined in Algorithm 5.2.

5.1.2.2 Hypergraph Model

The hypergraph model employed in SOCIAL is constructed in a manner that is identical to the hypergraph model utilized in LMCHGP. Specifically, the construction process entails the creation of a ball S centered around the seed node, as detailed in Section 5.1.1.2. Subsequently, the collection M of occurrences of the motif μ that are incident to this ball are enumerated according to the procedure outlined in Section 5.1.1.3. The hypergraph model H_μ is then constructed, as illustrated in Figure 5.2 and described in Section 5.1.1.4, to ensure the validity of Theorem 5.2.

5.1.2.3 Flow Model

In this section, we describe the process of constructing our MQI-based flow model G_f using the hypergraph model H_μ and an initial cluster $C_0 \subseteq S$ which contains the seed node u . There are three possible implementations of G_f based on the three already explained techniques to represent hypergraphs using graphs, namely *clique expansion*, *star expansion*, and *Lawler expansion* (see Figure 2.1). We show a bijective correspondence between certain s-t

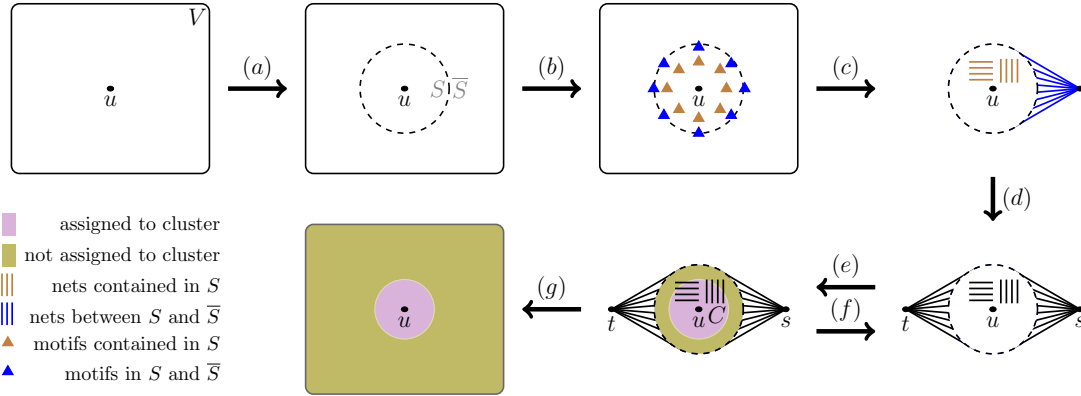


Figure 5.4. Illustration of the phases of SOCIAL. (a) Given a seed node u and a graph G , a ball S around u is selected. (b) Motif occurrences of μ with at least a node in S are enumerated. (c) The hypergraph model H_μ is built by converting motifs into nets and contracting \bar{S} into a single node. The ball S is taken as the initial cluster C_0 . (d) The flow model G_f is built based on C_0 in H_μ . (e) A cluster $C \subseteq C_0$ containing u is found using maximum flows. (f) While $C \subset C_0$, the model G_f is rebuilt based on C , which is taken as the initial cluster C_0 . (g) When eventually $C = S$, C is converted in a local cluster around the seed node in G .

cuts in G_f and clusters $C \subseteq C_0$ in G that include the seed node u and have motif conductance less than that of C_0 .

We start by converting our hypergraph model H_μ in a directed graph using the chosen net expansion technique. Second, we find a corresponding cluster C'_0 for C_0 in the created graph. For the *clique expansion*, $C'_0 = C_0$ since this transformation does not create artificial nodes. For the *star expansion*, C'_0 consists of C_0 and also the auxiliary artificial nodes connected to at least one node in C_0 . For the *Lawler expansion*, C'_0 consists of C_0 , the auxiliary artificial nodes w_1 having in-edges only from nodes in C_0 , and the auxiliary artificial nodes w_2 having out-edges to at least one node in C_0 . Third, we contract \bar{C}'_0 to a single *source* node s and then remove all its in-edges. Fourth, we multiply the weight of all the remaining edges by $d_w(C_0)$, i.e., the weighted volume of C_0 in H_μ . Fifth, we introduce a sink node t and include in-edges to it from each of the nodes $v \in C_0 \setminus \{u\}$, such that the weight of (v, t) is set to $cut(C_0)d_w(v)$, i.e., the cut-net of C_0 in H_μ multiplied by the weighted degree of v in H_μ . Finally, we include an edge (u, t) from the seed node to the sink and set its weight to infinity. Our flow network model G_f is concluded by setting edge capacities to match edge weights. Figure 5.5 shows the three possible configurations of our flow model G_f for a given hypergraph model H_f and an initial cluster C_0 .

Algorithm 5.2 Local Motif Clustering via Maximum Flows**Input** graph $G = (V, E)$; seed node $u \in V$; motif μ **Output** cluster $C^* \subseteq V$

- 1: $C^* \leftarrow \emptyset$
- 2: **for** $i = 1, \dots, \alpha$ **do**
- 3: Select ball S around u
- 4: $M \leftarrow$ Enumerate motifs in S
- 5: Build hypergraph model H_μ based on S and M
- 6: $C \leftarrow S$
- 7: **do**
- 8: $C_0 \leftarrow C$
- 9: Build flow model G_f based on C_0 in H_μ
- 10: Solve G_f to obtain cluster $C \subseteq C_0$ including u
- 11: **while** $C \subset C_0$
- 12: **if** $C^* = \emptyset \vee \phi_\mu(C) < \phi_\mu(C^*)$ **then**
- 13: $C^* \leftarrow C$
- 14: Convert C^* into a local motif cluster in G

We now analyze the theoretical guarantees provided by the defined flow model G_f . Theorem 5.5 shows that there is a set $C \subset C_0$ in G including the seed node u with motif conductance smaller than that of C_0 if, and only if, the value of the maximum flow on G_f is less than $\text{cut}(C_0)d_w(C_0)$, which is the weight of the trivial cut $(\{s\}, V(G_f) \setminus \{s\})$. In the proof, we show how such improved cluster C can be directly obtained from a maximum flow on G_f . Assumptions (a) and (b) in Theorem 5.5 are the same used in Theorem 5.2, which were previously shown to be reasonable in practice. Note that the claim is only valid for motifs with three nodes for clique and star expansion models, while it is valid in general for the Lawler expansion model.

Theorem 5.5. *There is a set $C \subset C_0$ in G including the seed node u with motif conductance smaller than that of C_0 if, and only if, the maximum flow on G_f is less than $\text{cut}(C_0)d_w(C_0)$ under the following assumptions:*

1. *the motif enumeration phase is exact;*
2. *$d_\mu(S) \leq d_\mu(\bar{S})$ in G ;*
3. *in case G_f is based on clique expansion or star expansion, the motif μ has three nodes;*




Figure 5.5. Flow model G_f given a hypergraph model H_μ and an initial cluster C_0 . Nodes and nets of H_μ are respectively represented by black circles and brown areas around them. The seed node u is circled in white and the initial cluster C_0 is surrounded by a dotted ellipse. Auxiliary artificial nodes and edges used in each net-expansion are respectively represented by brown circles and arrows. Bidirectional arrows represent pairs of edges in both directions. The seed node s , the sink node t , and the in-edges of t are respectively represented by a green circle, a blue circle, and blue arrows. Solid and dashed arrows respectively represent edges with finite and infinite weight.

Proof. We start with the backward direction, i.e., if the maximum flow on G_f is less than $cut(C_0)d_w(C_0)$, then there is a sub-set of C_0 in G including the seed node u with motif conductance smaller than that of C_0 . According to the Max-Flow Min-Cut Theorem [33], the weight of the maximum s-t flow on a network equals the weight of its minimum s-t cut, hence it follows that there is an s-t cut (B_1, B_2) of G_f with weight smaller than $cut(C_0)d_w(C_0)$. Without loss of generality, let $s \in B_1, t \in B_2, C = C_0 \cap B_2$, and hence, by definition, $C_0 \setminus C = C_0 \cap B_1$. Necessarily $u \in C$, otherwise the edge (u, t) , which has infinite weight, would be cut. There are two kinds of edges from B_1 to B_2 . First, there are the edges (x, t) , with $x \in C_0 \setminus C$. By design, the total weight of these edges is given by $cut(C_0)d_w(C_0 \setminus C)$. The second kind consists of edges from B_1 to $B_2 \setminus \{t\}$. These edges vary based on the net-expansion technique used, but their total weight is $cut(C)d_w(C_0)$ by design under the given assumptions.

Now we show that the total weight of the edges from B_1 to $B_2 \setminus \{t\}$ is $cut(C)d_w(C_0)$ for the three net-expansion techniques. In the clique expansion under assumption (c), each cut net e of C in H_μ corresponds directly to two cut in-edges of B_2 in G_f . The weight of each of these edges is, by design, set to $w(e)d_w(C_0)/2$, so they add up to the specified total weight. In the star expansion under assumption (c), for each cut net e of C in H_μ there is a single cut in-edge of B_2 in G_f that connects an auxiliary artificial node and a node from C_0 . The weight of this cut edge is set by design to $w(e)d_w(C_0)$, so the total weight of these edges is as claimed. In the Lawler expansion, for each cut net e of C in H_μ there is

exactly one cut in-edge of B_2 in G_f that connects two auxiliary artificial nodes, namely $w_1 \in B_1$ and $w_2 \in B_2$. If this were not the case, there would be an edge from B_1 to B_2 with infinite weight. The weight of this single cut edge is set by design to $w(e)d_w(C_0)$, so the sum of the weights of the cut edges is as stated. By adding up the weights of the two kinds of edges that cross the cut (B_1, B_2) and verifying that their total weight is less than $cut(C_0)d_w(C_0)$, we derive Equation (5.1), which can further be simplified to Equation (5.2). We conclude the proof of the backward direction by applying Theorem 5.2, Equation (5.2), and the assumptions (a) and (b).

$$cut(C_0)d_w(C_0 \setminus C) + cut(C)d_w(C_0) < cut(C_0)d_w(C_0) \quad (5.1)$$

$$\frac{cut(C)}{d_w(C)} < \frac{cut(C_0)}{d_w(C_0)} \quad (5.2)$$

Now we prove the forward direction, i.e., given a set $C \subset C_0$ including the seed node u in G with motif conductance smaller than that of C_0 , then the maximum flow on G_f is less than $cut(C_0)d_w(C_0)$. With the assumptions (a) and (b) in place, Equation (5.2) holds true for the selected set C . Since Equation (5.2) can be rewritten as Equation (5.1). To complete the proof, we will show that there exists an s-t cut (B_1, B_2) of G_f such that $s \in B_1, t \in B_2, B_2 \cap C_0 = C$, and the total weight of the in-edges of B_2 is $cut(C_0)d_w(C_0 \setminus C) + cut(C)d_w(C_0)$. Using the Max-Flow Min-Cut Theorem [33] again, it follows that if there is an s-t cut with a weight of less than $cut(C_0)d_w(C_0)$, the maximum flow value on G_f must also be less than $cut(C_0)d_w(C_0)$. Since $B_2 \cap C_0 = C$, it follows that $B_1 \cap C_0 = C_0 \setminus C$, hence there are $|C_0 \setminus C|$ cut edges of the form (x, t) . By design, the total weight of these edges is $cut(C_0)d_w(C_0 \setminus C)$.

Now we show that the weights of the remaining cut edges, i.e., edges from B_1 to $B_2 \setminus \{t\}$ add up to $cut(C)d_w(C_0)$ for each net-expansion technique. In the clique expansion, we forcibly have $B_2 = C \cup \{t\}$. Under assumption (c), each cut net e of C in H_μ corresponds to two cut in-edges of B_2 in G_f , both with weight set to $w(e)d_w(C_0)/2$ by design, hence they add up to the specified total weight. In the star expansion, under assumption (c), we make $B_2 = C \cup A \cup \{t\}$, where A is the set of artificial nodes a with $|N(a) \cap C| \geq |N(a)|/2$. For each cut net e of C in H_μ there is a single cut in-edge of B_2 in G_f that connects an auxiliary artificial node and a node from C_0 . The weight of this cut edge is set by design to $w(e)d_w(C_0)$, so the total weight of these cut edges is as stated. In the Lawler expansion, we make $B_2 = C \cup A \cup \{t\}$, where A consists of the auxiliary artificial nodes w_1 having in-edges only from nodes in C , and the auxiliary artificial nodes w_2 having out-edges to at least one node in C . Hence, for each cut net e of C in H_μ there is exactly one cut in-edge of B_2 in G_f that connects two auxiliary artificial nodes, namely $w_1 \in B_1$ and $w_2 \in B_2$. The

weight of this single cut edge is set by design to $w(e)d_w(C_0)$, so the sum of the weights of the cut edges is as expected. \square

SOCIAL utilizes a push-relabel approach to iteratively search for a maximum s-t flow in the model G_f . If the found maximum flow is strictly smaller than $cut(C_0)d_w(C_0)$, then we can directly find a minimum cut with the same weight as it and, consequently, a cluster $C \subset C_0$ containing the seed node u that has a strictly smaller motif conductance value $\phi_\mu(C)$ than that of C_0 in G . If such a cut is found, the algorithm repeats the process recursively setting the identified sub-cluster C as the new initial cluster, i.e., it constructs a new flow model based on H_μ and the initial cluster and uses the push-relabel algorithm to continue searching for sub-clusters with even lower motif conductance values. If, on the other hand, the maximum flow is not strictly smaller than $cut(C_0)d_w(C_0)$, it means that the current cluster C_0 is optimal among all of its sub-clusters containing the seed node u , and the algorithm terminates for the given ball S .

5.1.3 Experimental Evaluation

Setup. We implemented LMCHGP and SOCIAL in C++ within the KaHIP framework. In particular, LMCHGP was implemented using the public libraries for KaHyPar [52] and KaHIP [37]. We use the fastest configuration of these libraries throughout our experiments. Note that using stronger configurations would likely yield better solutions at the cost of higher running time. We compiled all algorithms using gcc 11.2 with full optimization turned on (-O3 flag). Unless mentioned otherwise, all experiments are performed on a single core of Machine D. For the reported experiments, we use a time limit of one hour for our overall algorithm. This time limit is checked between repetitions of each of our algorithms, hence it can be violated in case a particular partitioning procedure takes too long.

Baselines. We experimentally compare our algorithms against the state-of-the-art competitors, namely MAPPR [23]. We also ran preliminary experiments with HOSPLOC [87]. However, the algorithm is very slow even for small graphs and not scalable as their algorithm works using an adjacency matrix and hence needs $\Omega(n^2)$ space and time. Moreover, experiments done in their paper are on graphs that are multiple orders of magnitude smaller than the graphs used in our evaluation. Hence, we are not able to run the algorithm on the scale of the instances used in this thesis. We were not able to explicitly compare against LCD-Motif [24] since their code is not available (neither public, nor privately¹) and the data presented in the respective paper does not warrant explicit comparisons (e.g. seed nodes are typically not presented in the papers). However, we try to make implicit comparisons in Section 5.1.3.3.

¹Personal communication with the authors

Graph	SOCIAL			LMCHGP			MAPPR		
	ϕ_μ	$ C $	t(s)	ϕ_μ	$ C $	t(s)	ϕ_μ	$ C $	t(s)
com-amazon	0.031	76	<0.01	0.037	64	0.22	0.153	58	2.68
com-dblp	0.090	58	0.02	0.115	56	0.38	0.289	35	3.04
com-youtube	0.125	1832	4.52	0.172	1443	7.93	0.910	2	10.44
com-livejournal	0.158	494	3.33	0.244	387	8.17	0.507	61	173.80
com-orkut	0.273	1041	256.21	0.150	13168	496.94	0.407	511	923.26
com-friendster	0.388	2060	1194.50	0.368	10610	1339.99	0.741	121	16565.99
Overall	0.178	453	2.33	0.181	823	12.67	0.500	50	79.34

Table 5.1. Average comparison against state-of-the-art.

We compare our results against the globally best cluster computed for each seed node by MAPPR using its standard parameters ($\alpha = 0.98$, $\epsilon = 10^{-4}$).

Instances. Basic properties of the graphs used in our experiments can be found in Table 2.3. For our experiments, we split the graphs in two disjoint sets: a *tuning* set for the parameter study experiments and a *test* set for the comparisons against the state-of-the-art. The graphs in the test set are exactly the graphs used in the MAPPR paper [23]. Prior to our experiments, we removed parallel edges, self-loops, and directions of edges and assigning unitary weight to all nodes and edges. For each graph, we pick 50 random seed nodes and use all of them as input for each algorithm.

Methodology. For the sake of simplicity, all our experiments are based on the triangle motif, i.e., the undirected clique of size three. However, in general our method is also applicable to larger motifs. Since this motif has three nodes, Theorems 5.2, 5.4, and 5.5 are valid in general. In other words, this setup ensures that the graph model and the hypergraph model of LMCHGP are exact and guarantees that SOCIAL yields exact improvements for all net expansion techniques. We ensure the integrity of our results by using the same motif-conductance evaluator function for all tested algorithms.

We measure running time, motif-conductance, and/or size of the computed cluster. When averaging motif conductance over multiple instances, the final score is computed via arithmetic mean. This is a necessary averaging strategy since motif conductance can be zero, which makes the geometric mean infeasible to compute. When averaging running time or cluster size over multiple instances, we still use the geometric mean in order to give every instance the same influence on the *final score*.

5.1.3.1 Parameter Study

We tune the parameters of LMCHGP using the graphs in the Tuning Set of Table 2.3, which are disjoint from the graphs used for the evaluation against state-of-the-art. In a comparison of the hypergraph-based version of our algorithm against its graph-based version, each approach produces the best motif conductance for around 50% of the instances. Nevertheless, the hypergraph-based version is 23 times slower on average and uses up to 68.7 times more memory since the hypergraph model stores a large number of nets. Hence, we exclusively use the graph-based version of our algorithm for the remaining experiments. Next, we compare the effect of using different values for the parameters β and α . The results show the expected regular relationship between solution quality, running time, and these parameter: the larger α (resp. β), the smaller the motif conductance and the larger the running time. Finally, we checked that the impact of including the label propagation local search in LMCHGP is, on average, a 13% decrease in the motif conductance at the cost of only 1.5% more running time. Summing up, we use the following parameters for LMCHGP: graph model, label propagation, $\alpha = 3$, $\ell \in \{1, 2, 3\}$, and $\beta = 80$. In view of that fact that the parameters α and ℓ utilized in SOCIAL carry the same significance as their counterparts in LMCHGP, we directly incorporate the finely-tuned values of these parameters from the latter in SOCIAL. In addition, the clique expansion method proves to be the most efficient for our three-node motif as it requires no auxiliary artificial nodes and involves the fewest number of auxiliary artificial edges compared to all other net expansion techniques. In summary, we use the following parameters for SOCIAL: clique expansion, $\alpha = 3$, and $\ell \in \{1, 2, 3\}$.

5.1.3.2 State-of-the-Art

In this section we present and discuss our experiments against the state-of-the-art. All the reported results are based on the Tet Set of Table 2.3. In the performance profile plots shown in Figures 5.6a and 5.6b, we compare MAPPR [23], LMCHGP, and SOCIAL. In Table 5.1, we show average results for each graph.

As shown in Figure 5.6a, SOCIAL obtains the best or equal motif conductance value for 62% of the instances, while LMCHGP and MAPPR respectively obtain the best or equal motif conductance for 49% and 19% of the instances. This result can be explained with two observations. First, SOCIAL and LMCHGP explore the solution space considerably better than MAPPR, since both build the (hyper)graph model multiple times, while MAPPR simply uses the APPR algorithm. Second, SOCIAL is based on a flow approach which directly optimizes for motif conductance, whereas LMCHGP is based on a graph partitioning algorithm which is repeated multiple times to compensate for its design to minimize the number of cut motifs rather than motif conductance. In Table 5.1, SOCIAL outperforms LMCHGP for 4 of

the 6 graph and overall, and outperforms MAPPR for all graphs and overall. On the other hand, LMCHGP outperforms SOCIAL for 2 of the 6 graphs, and outperforms MAPPR with respect to motif conductance and running time for all graphs and overall. Overall, SOCIAL computes clusters with motif conductance 0.178 while LMCHGP and MAPPR compute clusters with motif conductance 0.181 and 0.500, respectively.




Figure 5.6. Performance profiles comparing our algorithms against the state-of-the-art in terms of solution quality and running time.

As exhibited in Figure 5.6b, SOCIAL is the fastest one for 87% of the instances, while LMCHGP and MAPP are the fastest ones for 12% and 1% of the instances, respectively. Furthermore, the running time of SOCIAL is within a factor 1.18 of the running times of the fastest competitors for all instances. SOCIAL is respectively up to 237 and 144 063 times faster than LMCHGP and MAPP while being a factor 5.4 and 34.1 faster than them on average. The reason for MAPP being considerably slower than the other algorithms is that it must enumerate motifs across the entire graph, while SOCIAL and LMCHGP only require enumeration of motifs in a ball around the seed node. The reduced but still substantial difference in running time between SOCIAL and LMCHGP is a result of LMCHGP's repeated partitioning of each ball around the seed node, while SOCIAL employs a flow model to greedily improve the motif conductance metric until a local optimum cluster is obtained. Table 5.1 demonstrates that, on average, SOCIAL has a lower running time than LMCHGP, and LMCHGP has a lower running time than MAPP, for each graph as well as overall.

For a more intuitive analysis of the quality of our results, Figure 5.7 plots motif conductance versus cluster size for all communities computed by the three algorithms. Observe that the communities found by SOCIAL are densely localized in the lower left area of the chart, which is the region with smaller motif conductance and smaller cluster size. On the




Figure 5.7. Motif conductance vs cluster size.

other hand, communities computed by MAPPR are often in the upper area of the chart and communities computed by LMCHGP are often in the right area of the chart.

5.1.3.3 Additional Comparisons

As mentioned above, we were not able to compare against LCD-Motif [24] explicitly since their code is not available (neither publicly, nor privately) and the data presented in the respective papers does not warrant explicit comparisons (e.g., seed nodes are typically not presented in papers, and in this case instances are directed rather than undirected). Here, we make an attempt at implicit comparisons. Zhang et al. [24] (Table 4 therein) compare motif conductance against MAPPR on three directed instances (cit-HepPh, Slashdot, StanfordWeb) and report an geometric mean improvement of 54% in motif conductance for the triangle motif. As LMCHGP and SOCIAL work for undirected instances, we have build undirected version of those graphs and run both as well as MAPPR for the triangle motif. Respectively, the geometric mean improvement LMCHGP and SOCIAL obtain over MAPPR is 194% and 223%. Both improvements are significantly larger than the improvement of Zhang et al. over MAPPR. Also note that in our experiments from Table 5.1, the geometric mean improvement (using the average motif conductance values) of LMCHGP and SOCIAL over MAPPR in motif conductance is 192% and 219%, respectively.

5.2 Conclusion

In this chapter, we proposed three local algorithms for graph decomposition: LMCHGP and SOCIAL. They were discussed in detail and were experimentally compared against

one another and against the state-of-the-art. Both proposed algorithms solve the local motif clustering problem with sophisticated combinatorial techniques. Our first contribution, LMCHGP (Local Motif Clustering via (Hyper)Graph Partitioning), is an algorithm based on (hyper)graph partitioning. We construct both a graph model and a hypergraph model to represent the motifs around the seed node, and then partition the (hyper)graph model using a multi-level partitioner to minimize the motif conductance in the original graph. Our second contribution, SOCIAL (faSter mOtif Clustering vIa mAXimum fLows), combines the strongly local hypergraph model proposed in the previous algorithm with the *max-flow quotient-cut improvement* algorithm (MQI) to optimize for motif conductance. We construct a flow model and utilize a push-relabel algorithm to find subsets of the initial cluster containing the seed node with lower motif conductance than the whole cluster. In experiments involving triangle motifs, LMCHGP produces communities with lower motif conductance compared to MAPPR, while being up to multiple orders of magnitude faster than it. On the other hand, SOCIAL produces communities with lower motif conductance compared to MAPPR and LMCHGP, while being up to multiple orders of magnitude faster than both. In future work, we intend to conduct experiments with larger motifs and use the Lawler-expansion version of our flow graph, since it is the only one whose quality guarantee holds true for larger motifs. Lastly, we intend to add parallelization to improve the speed on large instances further.

Chapter 6

Multilevel Algorithms

In this chapter, we propose our contributions in the field of multilevel graph decomposition. Firstly, we introduce an integrated multilevel algorithm for solving the process mapping problem, which is designed and implemented in multiple versions with different goals. Experimental results demonstrate that the multiple versions of our algorithm outperform the state-of-the-art in terms of both solution quality and running time. Secondly, we design a multilevel algorithm for solving the signed graph clustering problem and utilize it as a building block to construct a powerful distributed memetic algorithm. Our experimental evaluation demonstrates that our memetic algorithm is capable of converging effectively and significantly outperforms the state-of-the-art with respect to edge-cut.

References. This section is based on [154], which is joint work with Alexander van der Grinten, Henning Meyerhenke, Jesper Larsson Träff, and Christian Schulz, on [155], which is joint work with Felix Hausberger and Christian Schulz, and on the technical report [156], which is also joint work with Felix Hausberger and Christian Schulz.

6.1 Multilevel Process Mapping

In this section, we propose a high-quality algorithm for the process mapping problem. Our algorithm is an integrated approach based on a multilevel scheme coupled with sophisticated initial partitioning and local search techniques. Additionally, we introduce faster techniques to keep precise topology information without the need to store distance matrices. Our experiments indicate that our new integrated algorithm improves mapping quality over other state-of-the-art algorithms. For example, one configuration of our algorithm yields similar solution quality as the previous state-of-the-art in terms of mapping quality for large values of k while being a factor 9.3 faster. Compared to the currently fastest iterated multilevel mapping algorithm Scotch, we obtain 16% better solutions while investing slightly




Figure 6.1. Multilevel scheme used to solve the process mapping problem (adapted from [37]).

more running time. Most importantly, hierarchical multi-section algorithms that take the system hierarchy into account for model creation improve the results of the overall process mapping significantly.

6.1.1 Integrated Mapping

6.1.1.1 Overall Algorithm

We engineered all the components of a multilevel algorithm to solve the process mapping problem in an *integrated* way, as illustrated in Figure 6.1. This algorithm includes coarsening-uncoarsening schemes, methods to construct and refine initial solutions, local search methods, and additional tools to explore trade-offs in memory usage and performance.

6.1.1.2 Coarsening

We use a matching-based coarsening scheme. The *matching-based* coarsening is the most common choice in multilevel partitioning algorithms due to its simplicity, speed, and generality. It has two consecutive steps: An edge rating function and a matching algorithm. Based on local information, the *edge rating function* scores each edge to estimate the benefit of contracting it. We employ the same edge rating function $\exp^*(e) = \omega(e)/(|\Gamma(u)| * |\Gamma(v)|)$ as used in [37]. Then, the *matching algorithm* obtains a maximal match to maximize the sum of the ratings of the contracted edges. As in [37], we computed matchings with the *Global Paths Algorithm* [37], which is a $\frac{1}{2}$ -approximate algorithm.

6.1.1.3 Initial Solution Algorithms

We compute the initial mapping using a two-phase approach. To solve the graph partitioning phase, we compare two multilevel recursive bisection algorithms: (i) *standard bisection* setup, in which we perform a recursive bisection to obtain k blocks; (ii) *multi-section* setup,

in which we perform recursive bisections throughout the hierarchical structure of PEs. To construct a solution for one-to-one mapping, we apply two different construction methods: (i) *identity*, which assigns each block to the PE with the same ID to favor locality; (ii) *hierarchy top down*, which partitions the set of blocks throughout the hierarchical structure of PEs. To refine the one-to-one mapping solution, we perform an N_C^{10} swap neighborhood local search. Hence, the resulting map Π of nodes to PEs becomes our initial solution.

Our *standard bisection* setup for initial partition corresponds to the initial partition step in KaHIP. Moreover, it is a canonical choice to produce an initial partition using multilevel algorithms. On the other hand, the *multi-section* setup draws inspiration from the scheme used in [109]. It is an attempt to specialize the initial partition for the particular case tackled in this paper: a regularly hierarchical distribution of PEs in which the communication cost between two processes (nodes) highly depends on the hierarchy level shared by their corresponding PEs (blocks). Particularly, we apply a recursive partitioning scheme that splits all the nodes in a_ℓ blocks, then splits the nodes in each block in $a_{\ell-1}$ sub-blocks, then splits the nodes in each sub-blocks in $a_{\ell-2}$ sub-sub-blocks, and so forth. Observing that the communication costs decrease as the communicating processes share lower hierarchy levels, the multi-section approach implies a hierarchy of sub-problems that directly reflects the problem cost hierarchy.

In both setups of the partitioning step, we recursively assign consecutive IDs to blocks throughout the process in order to maintain locality. Moreover, the PEs belonging to each hierarchy module are labeled with consecutive IDs, which also promotes locality. Then, the *identity* method is a fast way to construct a solution for one-to-one mapping which takes advantage of this locality: it assigns each block V_i to the PE with the ID i . Note, the *standard bisection* setup conveniently combines with the identity mapping approach when k is a power of 2 since the recursive bisections will be automatically performed throughout the hierarchical topology. For an analogous reason, the *multi-section* setup is a good algorithm to create a coarse model to be mapped by the *identity* mapping approach independently of k . The *hierarchy top down* [92] is a more general approach to construct solutions for one-to-one mapping when the PEs are hierarchically organized. Its mechanism is similar to the idea of multi-section throughout the hierarchy.

6.1.1.4 Uncoarsening

After obtaining an initial solution at the coarsest level, we apply a sequence of four local refinement methods to move nodes between blocks (which are already associated to unique PEs). Then, we undo each of the contractions performed previously, from the coarsest graph until the original input graph. After each uncoarsening step, we repeat our four local refinement methods. The refinements run in a specific order based on their characteristics. First,

a *quotient graph refinement* exhaustively tries to improve solution quality and eliminate imbalance by moving nodes between each pair of blocks connected by an edge in the quotient graph. Second, a *k-way Fiduccia-Mattheyses (FM) algorithm [157] refinement* greedily goes through the boundary nodes trying to relocate them with a more global perspective in order to improve the mapping. Third, a *label propagation refinement* randomly visits all nodes and moves each one to the most appropriate block while not increasing the objective. Finally, a *multi-try FM refinement* is exhaustively applied in rounds with random starting points throughout the graph in order to escape local optima as many times as possible. Before explaining the local search algorithms, we introduce the notion of *gain* for the process mapping problem.

Gain. All our refinement methods are based on the concept of *gain*. We define $\Psi_b(v)$ as the *partial* contribution of a node v to the objective function $J(C, D, \Pi)$ (Equation (2.1)) in case v is assigned to the PE b . More precisely, $\Psi_b(v)$ represents the total cost of the communications involving v if $\Pi(v) = b$ and the neighbors of v remain assigned to their current PEs. The *gain* $g_b(v)$ represents the value that will be subtracted from $J(C, D, \Pi)$ if a node v is moved from its current PE $\Pi(v)$ to PE b . More precisely, $\Psi_b(v) := \sum_{\{v,u\} \in I(v)} C_{v,u} D_{b,\Pi(u)}$ and $g_b(v) := \Psi_{\Pi(v)}(v) - \Psi_b(v)$. Note that $g_{\Pi(v)}(v) \equiv 0$. Observe that a positive (resp., negative) gain indicates improvement (resp., worsening) of the solution. Computing the gains of v to all blocks in $R(v)$ costs $O(|R(v)||I(v)|) = O(|I(v)|^2)$. For comparison purposes, the computation of the same corresponding gains in the context of graph partitioning and edge-cut objective function costs $O(|I(v)|)$.

Quotient Graph Refinement. We implemented an adapted version of the *quotient graph refinement* [37] to incorporate our definition of gains. Within this refinement, we visit each pair of neighboring blocks in the quotient graph Q underlying the current k -partition. Then we apply an FM algorithm [157] to move nodes between the two currently visited blocks, keeping two respective gain-based priority queues of eligible nodes. Each queue is randomly initialized with the boundary in its corresponding block. After a node is moved (which can only happen once during an execution of the local search), its unmoved neighbors become eligible.

K-Way FM Refinement. Our k -way FM refinement was adapted from the implementation in [37]. Unlike the quotient graph refinement, the k -way FM does not restrict the movement of a node to a certain pair of blocks, but performs global-aware movement choices. Our implementation of k -way FM uses only one gain-based priority queue P , which is initialized with the *complete* partition boundary in a random order. Then, the local search

repeatedly looks for the highest-gain node v and moves it to the best $c(v)$ -underloaded neighboring block. When a node is moved, we insert in P all its neighbors that were not in P and have not been moved yet. The k -way local search stops if P is empty (i.e., each node was moved once) or when a stopping criterion based on a random-walk model described in [37] applies. To escape from local optima, this refinement allows some movements with negative gain or to blocks that are not $c(v)$ -underloaded. Afterwards local search is rolled back to the lowest cut fulfilling the balance criterion that occurred.

Label Propagation Refinement. We propose a local search inspired by *label propagation* [140]. The algorithm works in rounds. In each round, the algorithm visits all nodes in a random order, starting with the labels being the current assignment of nodes to blocks. When a node v is visited, it is moved to the $c(v)$ -underloaded neighboring block with highest positive gain. We consider only $c(v)$ -underloaded blocks since this ensures that the target block is not overloaded when the node is moved there. Ties are broken randomly and a 0-gain neighboring block can be occasionally chosen with 50% probability if there is no neighboring $c(v)$ -underloaded block with positive gain. We perform at most ℓ rounds of the algorithm, where ℓ is a tuning parameter.

Multi-Try FM Refinement. We also adapted our gain concept to a localized variant of the k -way local search algorithm similar to that proposed in [37] under the name of *multi-try* FM. Instead of being initialized with all boundary nodes, as in k -way FM, multi-try FM is repeatedly initialized with a single boundary node. This introduces a higher diversification to the search since it is not restricted to movements in boundary nodes with global largest gain. As a result, this local search can escape local optima more easily than k -way FM.

6.1.1.5 Additional Techniques

Implicit Distance Matrix. When the topology matrix D is stored in memory, access time to obtain the distance between a pair of PEs is $O(1)$, but this requires $O(k^2)$ space. From now on, we refer to the algorithm explicitly keeping D in memory as *matrix-based* approach. We implement three alternative approaches to save memory by exploiting the fact that our topology matrix is a hierarchy and the IDs of PEs in each of the hierarchy modules are sequential. For simplification reasons, we call these approaches: (i) *division-based*; (ii) *stored division-based*; and (iii) *binary notation-based*.

In the *division-based* approach, we perform $O(\ell)$ successive integer divisions and comparisons in the ID of two PEs when we need to find out their distance. Here, ℓ is the number of levels in the system hierarchy. As a preprocessing step to be executed only once, we create a vector $h = \left(k / \prod_{t=1}^{\ell} a_t, k / \prod_{t=2}^{\ell} a_t, \dots, k / a_{\ell} \right)$. To find the distance between




Figure 6.2. Section structure of the binary number used to represent PE b .

PEs b and b' with $b \neq b'$, we loop through the hierarchy layers from $i = \ell$ to $i = 1$. In each iteration, we perform the integer division of b and b' by h_i . Whenever the division results differ, then we break the loop and return $D_{b,b'} = d_i$. This approach does not require any additional memory other than a vector with $O(\ell)$ integers and has time complexity $O(\ell)$.

The *stored division-based* approach works in a similar way as the *division-based* one. The only difference is that we avoid repetitive integer divisions of IDs by elements of h by storing the results of all possible divisions in a preprocessing step executed only once. Although we still need $O(\ell)$ running time to perform comparisons in order to obtain the distance between a pair of PEs, the constant factors involved are much lower. This improvement in running time comes at the cost of additional $O(k\ell)$ memory.

The *binary notation-based* approach is a more compact way of decomposing the IDs of PEs. Instead of storing ℓ numbers for each PE, we keep in memory a single binary number per PE. This binary number r consists of ℓ sections r_i , each containing s bits, with s given by Equation (6.1) (see Figure 6.2). To describe the construction of r for a PE b , let a variable t be initialized as $t = b$. Then, we loop through the hierarchy layers, from $i = 1$ to $i = \ell$. In each iteration i , r_i receives the remainder of the division of t by a_i and, then, t is updated to store the integer quotient of t by a_i . Afterwards, it is possible to precisely locate b at the hierarchy by sweeping the sections of r from r_ℓ to r_1 . In particular, r_ℓ specifies its data center, $r_{\ell-1}$ specifies its server among those belonging to its data center, and so forth. Obtaining the distance between distinct PEs b and b' is equivalent to finding which section r_i contains the leftmost nonzero bit in the result of the bit-wise operation $\text{XOR}(b,b')$. The running-time complexity of finding the section of the leftmost nonzero bit is $O(\log(\ell))$. Furthermore, current processors often implement a *count leading zeros* (CLZ) operation in hardware which allows the identification of the leftmost nonzero bit in $O(1)$ time, under the assumption that the size $\log r = O(\log k)$ of the binary numbers is smaller than the size of a machine word.

$$s = \lceil \log_2 \left(\max_{1 \leq t \leq \ell} (a_t) \right) \rceil \quad (6.1)$$

Delta-Gain Updates. Our local searches frequently need to compute *gains* involved in the movement of nodes. A *base approach* to check these gains consists of computing them from scratch whenever they are needed, which can yield many gain recomputations. For this reason, we implement a technique to save running time called *delta-gain updates* [52].

In *delta-gain updates*, we store a vector R of length $|R(v)| = O(|\Gamma(v)|)$ for each node v . In this vector, we keep the gains $g_b(v)$ for all PEs b containing neighbors of v . Additionally, we store an n -sized vector h to keep flags that indicate whether a node has up-to-date gains in memory. Asymptotically speaking, these vectors represent $O(n+m)$ extra memory. Each flag is initialized with an inactive seed and is considered active if its value equals a counter that is increased after each uncoarsening steps. When we need to check a gain of some node v , we look at h_v to verify if the gains of v are up-to-date. If they are not, we compute all gains $g_b(v)$ from scratch, which costs $O(|I(v)|^2)$, and activate h_v . Otherwise, we just access the required gain from memory in $O(1)$ time.

If a node v moves from its current PE to another one, we have to update all delta gains of v and $u \in \Gamma(v)$ with h_u being active. Assume that h_v and h_u are active and v moves from PE 1 to PE 2 during some local refinement. After this movement, we should change the delta gains of u and v in memory. For v , it suffices to subtract $g_2(v)$ from all other gains of v and then set $g_2(v)$ to 0. For u , it is slightly trickier, but we do not need to recalculate all its gains from scratch since their only source of change is the edge e that connects u and v . Hence, we respectively subtract and add to $g_b(u)$ the corresponding contribution of e before and after the movement of v . We end up doing the update in time $O(|I(v)| + |I(v)| * \overline{|R(u)|})$, where $\overline{|R(u)|}$ is the average of $|R(u)|$, $\forall \{v, u\} \in I(v)$.

6.1.2 Experimental Evaluation

Setup. We performed our implementations using the KaHIP framework (using C++) and compiled them using gcc 8.3 with full optimization turned on (-O3 flag). All of our experiments were run on a single core of a Machine E.

Baselines. For experiments based on the two-phase approach for tackling the process mapping problem, we solve the graph partitioning phase using KaHIP [37]. We use its configurations *fast*, *eco* and *strong* which are described in [37] – we respectively refer to them as K(Fast), K(Eco) and K(Strong). We select the most successful one-to-one mapping algorithms from [92] and also Scotch for our comparison: (i) *Top down* with N_C^d local search (TopDownN), which represent the state-of-the-art for one-to-one mapping when k is not a power of 2; (ii) *identity* mapping, which (when coupled with the KaHIP multilevel partitioning algorithm) represents the state-of-the-art for the process mapping problem via two-phase approach when k is a power of 2; (iii) the algorithm of *Müller-Merbach* [94]

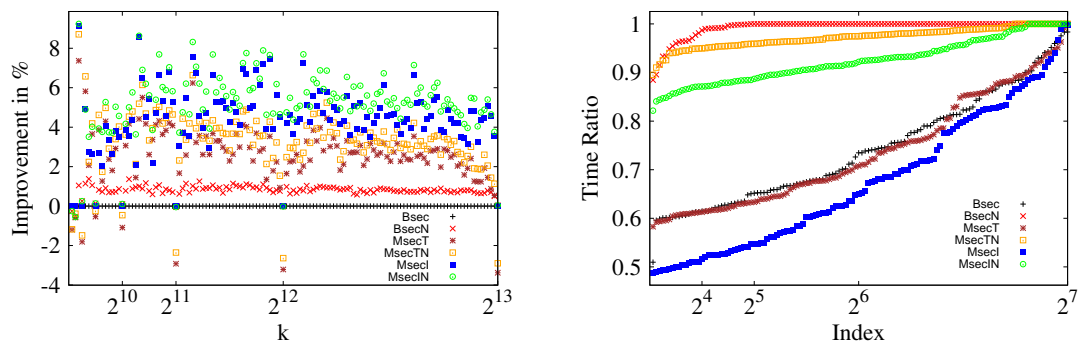
(Müller-Merbach), whose results are also used as a reference algorithm to calculate solution improvements in [92]; and (iv) Scotch [98]. We run the two-phase approaches TopDownN, Identity, and Müller-Merbach coupled with K(Fast), K(Eco) and K(Strong) as a partitioning algorithm. We also compare against global multi-section [109], in which the graph partitioning step is already using *hierarchical multi-section* itself. In particular, we use its configurations strong, econ, and fast, which we refer to as Gmsec(Strong), Gmsec(Eco) and Gmsec(Fast). Recall that this algorithm is also non-integrated: it uses different quality configurations of KaHIP to partition the graph, compute the coarser communication model, and then apply TopDownN to solve one-to-one mapping. We also run Scotch [98] configured to only use recursive bipartitioning methods using the quality setting. Scotch is among the algorithms with best running times in our experiments. Hence, we add an algorithm (ScotchTC) which reports the best solution out of multiple runs of Scotch with different random seeds when given the same amount of time to compute a solution as our *strong* configuration has used. We contacted Christopher Walshaw, who informed us that Jostle [158] is not available anymore.

Instances. Our instances come from various sources. Basic properties of the graphs under consideration can be found in Table 2.2. To keep the evaluation simple, we use the following hierarchy configurations for all the experiments: $D = 1 : 10 : 100$, $S = 4 : 16 : r$, with $r \in \{1, 2, 3, \dots, 128\}$. Hence, $k = 64 \cdot r$.

Methodology. Depending on the focus of the experiment, we measure running time and/or the process mapping communication cost defined in Equation (2.1). Some of our plots are *modified* performance profiles. These plots relate the running times of all algorithms to the slowest algorithm on a per-instance basis. For each algorithm, these ratios are sorted in increasing order. The plots show $\left(\frac{\sigma_A}{\sigma_{\text{slowest}}}\right)$ on the y-axis. A point close to zero indicates that the algorithm was considerably faster than the slowest algorithm.

6.1.2.1 Parameter Study

In this section, we present a sequence of experiments to test the performance of our algorithmic components regarding solution quality and running time. Our general goal consists of individually evaluating the effectiveness and significance of each component. Our specific goal consists of tuning three different configurations of the algorithm based on different principles: (i) a *strong* configuration, mostly concerned with maximizing solution quality; (ii) a *fast* configuration, mostly concerned with minimizing running time; and (iii) an *eco* configuration, which seeks to balance running time and solution quality.



(a) Improvements in objective function over Bsec. Higher is better.

(b) Performance profile for running time (ordered running time ratios). Lower is better.

Figure 6.3. Comparing initial mapping algorithms from Table 6.1.

Our experimental strategy consists of defining a single *focus* aspect of the algorithm for each experiment. Then, this specific aspect is tested with different components or setups while other parameters of the algorithm are kept constant. Initially we use standard components. Then we use the best component found in an experiment the next section.

We begin by focusing on a representative component of the multilevel scheme: (i) initial mapping; (ii) local search. Then, we evaluate algorithmic aspects which only affect running time and memory consumption: the distance matrix representation. The *standard configuration* consists of the matching-based contraction, all local search methods, explicit storage of distance matrix, and no delta-gains updates. All experiments in this section ran for the six *tuning graphs* from Table 2.2.

Initial Mapping. For the computation of initial mappings, we consider the six configurations listed in Table 6.1. Observe that Bsec and BsecN should apply either *identity* or *top down* depending on k . This choice is based on results obtained in [92] comparing these two construction algorithms for one-to-one mapping. Figure 6.3 plots the results regarding solution quality and running time for our six configurations.

Looking at solution quality, the configurations using multisection dominate those using standard bisection except for instances having k as a power of 2. This exception was expected since the standard bisection naturally performs a multisection partition for these instances. Among the configurations using multisection, *identity* produces overall better solutions than *hierarchy top down*, which is explained by the inherent locality of the multisection approach. Finally, the N_C^{10} local search is the least significant factor for solution quality, although it slightly improves solution compared to the similar configurations that skip local search.

Config.	Std.Bisec.	Multisec.	Identity	Top Down	N_C^{10}
Bsec	yes	no	if k power of 2	if k not power of 2	no
BsecN	yes	no	if k power of 2	if k not power of 2	yes
MsecT	no	yes	no	yes	no
MsecTN	no	yes	no	yes	yes
MsecI	no	yes	yes	no	no
MsecIN	no	yes	yes	no	yes

Table 6.1. Various configurations for the evaluation of different initial mapping algorithms.

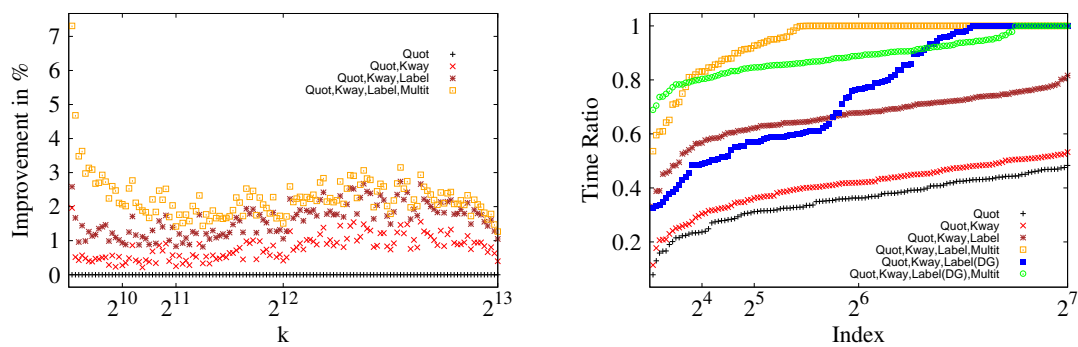
The N_C^{10} local search is the dominant factor regarding running time. Observe that the configurations using *identity* are always the fastest ones among those algorithms that either use N_C^{10} local search or among those that don't. Hence, the construction algorithm for one-to-one mapping is the second most relevant factor for running time. Finally, the partitioning algorithm has little influence over running time, which reflects the rather small time difference between each of the pairs {BsecN, MsecTN} and {Bsec, MsecT}.

Since MsecIN has the best overall solution quality results, it is the natural choice for *strong*. Notice that MsecI has the best overall running times, which makes it the perfect choice for *fast*. Nevertheless, it is also the second best regarding solution quality, which suffices to make it also the best choice for *eco*.

Local Search. For local search experiments, we start looking at the *fast* algorithm. To obtain a fast algorithm, we restrict its number of local search methods to one. Experiments with single local search algorithms do not yield much insight except that label propagation with delta-gain updates yields a very good trade-off for running time and solution quality.

For the *eco* configuration of our algorithm, we build four configurations by incrementally inserting the local search methods. Additionally, we consider two extra configurations equipped with delta-gain updates during label propagation. Figure 6.4 summarizes the results concerning these six configurations. Since the behavior of *strong* in this experiment is equivalent, we omit its results without loss of completeness.

Figure 6.4 shows that solution quality and running time consistently increase after each consecutive addition of local refinement methods. Regarding delta gains, running times decrease for some values of k but increase considerably and irregularly for others. Since this behavior is undesirable for *eco*, we drop delta gains for it. We also drop delta gains for *strong* since it does not affect solution quality and has negligible influence on running time compared to the N_C^{10} refinement. The clear choice for *strong* is the configuration with the four local searches since all of them contribute to incrementally improve solution quality. For *eco*, we drop only the multi-try FM local search since it adds little to solution quality but significantly increases the running time.



(a) Improvements in objective function over Quot. Higher is better.

(b) Performance profile for running time (ordered running time ratios). Lower is better.

Figure 6.4. Results for local search experiment of the *eco* algorithm. It comprises four scenarios that represent successive additions of the respective local search methods. In (b), we show two additional scenarios in which delta-gain updates are used (represented by DG).

Distance Matrix. As the objective function is not influenced by the representation of the distance matrix, we only evaluate running time. We test four configurations: one with each of the three techniques that imply the distance matrix and a reference scenario in which we store the full distance matrix. Since all configurations of our algorithm display equivalent behavior, we focus on *strong* without loss of generality. Figure 6.5 plots a running time ratio chart for *strong*. It is easily understandable that the *binary notation* technique is faster than the *stored division*-based approach, and also that the latter is faster than the *division*-based approach. On the other hand, the *binary notation* outperforms the full-distance matrix approach. While both approaches allow $O(1)$ distance calculations on our *x86_64* architecture, accessing the distance matrix incurs a memory access. This leads to frequent cache misses since the $O(k^2)$ -sized distance matrix does not fit into the cache of our machine. Lastly, not using the distance matrix significantly improves the memory footprint of the algorithm. This especially pronounced if the number of blocks gets very large. For example, for 2^{15} blocks, not using the distance matrix saves roughly an order of magnitude in memory.

After the tuning step, the three configurations of our algorithm ended up as follows: (i) *fast* applies MsecI, label propagation with delta-gain updates, and binary notation; (ii) *eco* applies MsecI, quotient graph refinement, k -way FM, label propagation, and binary notation; and (iii) *strong* applies MsecIN, quotient graph refinement, k -way FM, label propagation, multi-try FM, and binary notation. To improve speed even more, we also include a configuration called *fastest* which applies MsecI as initial mapping, does not use any local search during uncoarsening, and never needs to use information from the distance matrix.


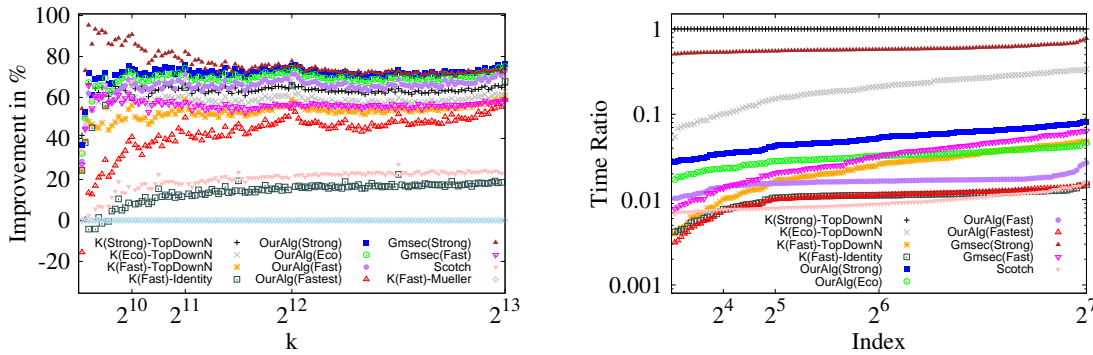


Figure 6.5. Performance plot for running times of different algorithm configurations (ordered running time ratios for different distance matrix implementations). Four alternative configurations are considered: implicit representations of the distance matrix based on division, stored division, and binary notation, and a reference scenario in which the matrix is explicitly stored. Lower is better.



(a) Improvements in objective function over k (Fast)-Müller-Merbach. Higher is better. (b) Performance profile for running time. Lower is better.

Figure 6.6. Comparisons against state-of-the-art approaches for process mapping.

6.1.2.2 State of the Art

In this section, we compare our algorithms against the best alternative algorithms in the literature. We report our experiments on all graphs listed in Table 2.2 – excluding the graphs from the tuning set. Figure 6.6 gives an overview over our results.

Overall, Scotch has the lowest average running time, directly followed by our algorithm *fastest*, $K(\text{Fast})$ -Identity, and our algorithm *fast* (respectively 9%, 10%, and 73% slower than Scotch on average). The average running time of $K(\text{Fast})$ -TopDownN and $Gmsec(\text{Fast})$ are respectively a factor 2.3 and 3.1 higher than Scotch. For our algorithms *eco* and *strong*, this factor is respectively 3.3 and 5.4. By definition ScotchTC is also a factor 5.4 higher than Scotch. Next, $Gmsec(\text{Eco})$ and $K(\text{Eco})$ -TopDownN have much higher runtimes (9.3 and 20.4 times slower than Scotch, respectively). $Gmsec(\text{Strong})$ and $K(\text{Strong})$ -TopDownN are the slowest ones (62 and 107 times slower than Scotch, respectively).

We now highlight the comparison of various configurations/algorithms. Gmsec(Strong) is the algorithm with best overall mapping quality. It is 2.5% on average better compared to our *strong* configuration, however our *strong* configuration is more than an order of magnitude faster on average – a factor 11.5 faster. Better quality of Gmsec(Strong) stems from the fact that global multi-section itself already takes the system hierarchy into account and hence yields good models to be mapped. Moreover, the graph partitioning approach itself which is used to compute a communication graph also uses more (time-consuming) sophisticated local search algorithms. This includes methods such as flow-based methods which particularly work well for small values of k as well as global search methods such as V-cycles. In particular for $k > 2^{11}$, our *strong* has the same average quality as Gmsec(Strong) but is 9.3 times faster. Our algorithm computes similar solutions in much less time since the multilevel algorithm directly optimizes the correct objective. For $k > 2^{11}$, our *eco* is 2.4% better and 2.8 times faster than Gmsec(Eco), and our *fast* is 9% better and 2.6 times faster than Gmsec(Fast). Overall, our *strong* configuration improves solution quality over K(Strong)-TopDownN by 5.1% while being a factor 20 on average faster. Our *eco* configuration has roughly 3.6% better quality than K(Strong)-TopDown but is a factor 32 faster on average. Our *fast* configuration still yields 1.3% better solutions than K(Strong)-TopDown on average, and is a factor 62 faster. Here, improvements stem from the fact that the new algorithms are integrated and not two-phase as well as the fact that these algorithms do not perform multi-sections. Lastly, our *fastest* algorithm is on average 9% slower than Scotch but also improves solution quality over Scotch by 16%. Our *strong* algorithm is 40% better than Scotch and consumes a factor 5.4 more running time.

6.2 Multilevel Signed Graph Clustering

In this section, we make a first attempt at solving the signed graph clustering problem by leveraging some of the most effective techniques from graph partitioning that minimize edge-cut. We introduce and thoroughly engineer a memetic algorithm specifically designed for this problem. The memetic algorithm features a multilevel approach, which encompasses a coarsening-uncoarsening process and efficient local search methods. The memetic algorithm is further enhanced with natural multilevel recombination and mutation operations. We also parallelize our approach using a scalable coarse-grained island-based strategy which has already shown to be scalable in practice. Experimental results show that our memetic algorithm is able to converge effectively and demonstrate that it outperforms the state-of-the-art with respect to edge-cut, producing significantly better solutions in 1 hour of processing on 16 elements.

6.2.1 Multilevel Algorithm

We now describe a novel multilevel algorithm for signed graph clustering which is an important component of our memetic algorithm. We start with our multilevel strategy, then we discuss each of its components.

6.2.1.1 Overall Multilevel Strategy

In this section, we describe our overall multilevel strategy for signed graph clustering. Our multilevel approach for graph partitioning differs from others in that it lacks a separate algorithm for initial solution computation. Instead, our scheme starts with coarsening to produce a hierarchy of coarser graphs. Upon completion, the coarsest graph is assigned individual clusters for each node, resulting in an initial clustering. The uncoarsening phase maps the current best clustering all the way back to the original graph while running local search approaches on each level. Two consecutive cycles of this multilevel process are performed, as respectively shown in Figures 2.2 and 6.7. In the second cycle, the previously obtained clustering is forced as the initial solution on the coarsest level, to be further refined during uncoarsening. Our overall multilevel algorithm is outlined in Algorithm 6.1.

Algorithm 6.1 Our multilevel algorithm

Input graph $G = (V, E)$

Output clustering $\Pi : V \rightarrow \mathbb{N}$

```

1:  $G_0 \leftarrow \emptyset; G_1 \leftarrow G; i \leftarrow 0$ 
2: while  $G_i \neq G_{i+1}$  do
3:    $i \leftarrow i + 1$ 
4:   ComputeClustering( $G_i$ )
5:    $G_{i+1} \leftarrow \text{Contract}(G_i)$ 
6:  $\Pi_i \leftarrow \text{MapNodesToClusters}(G_i)$ 
7: for  $j \leftarrow i - 1, \dots, 1$  do
8:    $\Pi_j \leftarrow \text{RemapToFiner}(\Pi_{j+1})$ 
9:    $\Pi_j \leftarrow \text{RefineClustering}(G_j, \Pi_j)$ 
10:  $\Pi \leftarrow \text{GlobalSearch}(\Pi_1)$ 

```

6.2.1.2 Coarsening

In this section, we explain our *coarsening* phase. As shown in Algorithm 6.1, it consists of two consecutive steps which are iteratively repeated on the current coarsest graph until no more contraction can be performed without increasing the overall edge-cut value. In the

first step, a graph clustering is computed. In the second step, each cluster contained in this clustering is contracted in order to produce a coarser graph.

To be self-contained, we briefly outline our clustering algorithm, which is based on *label propagation* [140]. In a graph with n nodes and m edges, size-constrained label propagation can be implemented in $O(n+m)$ time. The algorithm starts with each node in its own cluster, i.e., its initial cluster ID is set to its node ID. The algorithm then works in rounds. In each round, all nodes are traversed, and each node u is moved to the cluster with the largest connection weight with u in case it is strictly positive, i.e., it is only moved to the cluster V_i with maximum connection weight $\omega(\{(u, v) \mid v \in N(u) \cap V_i\})$ if it is strictly positive. Ties are broken randomly. This greedy approach ensures that the edge-cut is monotonically decreased. The algorithm is repeated at most L times, where L is a tuning parameter. The intuition of the overall procedure is to force positive edges inside clusters and negative edges between clusters, which is favorable for the edge-cut objective.

After the label propagation, each cluster is replaced by a single node (as exemplified in Figure 2.3), creating a hierarchy of coarser graphs. The contraction ensures a clustering of the coarse graph corresponds to the same edge-cut in finer graphs. This aggressive graph-contraction strategy allows us to shrink irregular networks, and the coarsening stops when no further contraction can be done without increasing the edge-cut. The final output is the coarsest graph, which implies an initial clustering of assigning each node to its own cluster. This coarsening phase serves as a method to construct an initial clustering, starting with n clusters and decreasing the number of clusters and the edge-cut at each hierarchy level.

6.2.1.3 Uncoarsening

In this section, we explain our *uncoarsening* phase. It is executed directly after the coarsening phase has been completed and every node of the coarsest graph has been assigned to its own cluster. As shown in Algorithm 6.1, our uncoarsening phase consists of two consecutive steps which are iteratively repeated for each graph in our contraction hierarchy, from the coarsest one to the finest one. In the first step, the current clustering is mapped to the graph contained in the current level of the contraction hierarchy. In the second step, we apply a sequence of local refinement methods to optimize our clustering by moving nodes between clusters.

At each level of uncoarsening, we remap the best clustering to the next finer graph using our contraction approach, maintaining edge-cut as in the example shown in Figure 2.3. We apply two refinement methods: First, a *label propagation* refinement randomly visits all nodes and greedily moves each one to the cluster where edge-cut is minimized. Second, a *k-way Fiduccia-Mattheyses* (FM) [157] refinement greedily goes through the boundary nodes trying to relocate them with a more global perspective in order to improve

solution quality. Both methods are based on the concept of *Gain*, defined as the increase in edge-cut caused by node movement.

Label Propagation Refinement. During uncoarsening, one of our local search procedures is based on *label propagation* [140]. The algorithm works in rounds. In each round, the algorithm visits all nodes in a random order, starting with the labels being the current assignment of nodes to clusters. When a node v is visited, it is moved to the neighboring cluster with highest positive gain. Ties are broken randomly and a 0-gain neighboring cluster can be occasionally chosen with 50% probability if there is no neighboring cluster with positive gain. We perform at most ℓ rounds of the algorithm, where ℓ is a tuning parameter.

K-Way FM Refinement. Our k -way FM refinement was adapted from the implementation in [37], which is an adapted version of the FM algorithm [157] to move nodes between clusters. Unlike the original algorithm, FM algorithm, our refinement moves nodes globally, rather than restricting movements to pairs of clusters. We use a gain-based priority queue initialized with the *complete* partition boundary. We repeatedly move the highest-gain node to its best neighboring cluster, updating the queue by adding its neighbors not yet moved. The search stops when the queue is empty (i.e., each node was moved once) or a random-walk stopping criterion described in [37] is met. The refinement escapes local optima by allowing negative-Gain movements and rolling back to the lowest cut that satisfies the balance criterion.

6.2.1.4 Global Search

After the standard multilevel cycle described in Algorithm 6.1, a global search is run using an extra multilevel cycle almost identical to the standard one, with the previous clustering used as the initial solution. Figure 6.7 illustrates this global search procedure. Its difference to the standard multilevel cycle is that the previously computed clustering is forced to be the initial solution in the extra multilevel cycle. This is achieved by blocking cut edges from contraction and stopping coarsening when the coarser graph equals the quotient graph of the clustering. This global search strategy was introduced in [159] and has been successfully used for many optimization problems. This approach is extended in Section 6.2.2 to form the operators of our evolutionary algorithm.

6.2.2 Distributed Evolutionary Algorithm

In this section, we extend our multilevel algorithm to a distributed evolutionary framework. Our approach is directly inspired by the evolutionary algorithm proposed in [160] for graph




Figure 6.7. Global Search (adapted from [37]).

partitioning. We start by providing a broad description of our overall evolutionary strategy. Then, we discuss each of its algorithmic components and show how we use our multilevel algorithm as a tool to implement evolutionary operators such as recombination and mutation.

6.2.2.1 Evolutionary Strategy

In this section, we describe the overall strategy of our distributed evolutionary algorithm, which is summarized in Algorithm 6.2. As a typical evolutionary approach, our algorithm is a population-based heuristic which mimics the biological process of evolution to optimize for our problem. Given a signed graph G , we define an *individual* as a particular clustering on G . Each processing element initializes a population P of α individuals using different random seeds. Over many rounds or *generations*, the population evolves via either *mutation* (with probability $\beta = 10\%$) or *recombination* (with probability $1 - \beta = 90\%$). In particular, the process continues until a desired time limit t_f is reached. The *recombination* procedure is executed on two *good* individuals selected via *tournament* and then used as *parents* to produce *offspring* which partially inherits their characteristics. The *mutation* procedure is run on a randomly-picked individual which is used as reference for building another individual from scratch. The next phase is called *replacement*. New individuals are inserted into the population and worst ones are *evicted* to maintain α population size. Each running instance of our algorithm generates only one individual per generation, hence it is a *steady-state* evolutionary algorithm [39]. After each generation, the best individual found so far by each instance is shared among processing elements.

6.2.2.2 Notation, Population, and Fitness

In this section, we describe how the evolutionary concepts of notation, fitness function, and population are defined and operated within our algorithm. We define an individual as a particular clustering Π on the graph G and represent it by its set of cut edges. This way,

Algorithm 6.2 Our distributed evolutionary algorithm**Input** graph $G = (V, E)$ **Output** clustering $\Pi : V \rightarrow \mathbb{N}$

```

1:  $P \leftarrow \{\Pi_1, \dots, \Pi_\alpha\}$ 
2: while running time  $< t_f$  do
3:    $i \leftarrow \text{Random}([0, 1])$ 
4:   if  $i > \beta$  then
5:      $\Pi_a, \Pi_b \leftarrow \text{Selection}(P)$ 
6:      $\Pi_c \leftarrow \text{Recombination}(P)$ 
7:   else
8:      $\Pi_c \leftarrow \text{Mutation}(\Pi_a)$ 
9:    $P \leftarrow \text{Replacement}(P, \Pi_c)$ 
10:   $\text{CommunicateBest}(P)$ 
11:  $\Pi \leftarrow \text{BestIndividualOverall}()$ 

```

each clustering has a unique notation, i.e., there is no symmetry to represent individuals. Note that some sets of edges cannot be the set of cut edges of any feasible clustering, e.g., a single edge of a cycle. However, feasibility is *implicitly* ensured throughout the algorithm, as will become clear later. As a consequence, we do not need any penalty function, hence the *fitness* function of an individual is its *edge-cut*, which is the objective to be minimized. The initial population is built by running our multilevel algorithm (Algorithm 6.1) from scratch α times, where $\alpha \geq 3$ is automatically determined such that roughly 10% of the allowed time limit t_f is spent on the construction of the initial population.

6.2.2.3 Recombination

In this section, we explain the operation of the recombination operator in our evolutionary algorithm. We select two *good* individuals, Π_a and Π_b , from the population to serve as parents using a *tournament* [161], i.e., two distinct individuals are drawn and the one with best edge-cut is selected to be a parent. To ensure distinct parents, the second parent may be the loser of the tournament if the winner is already the first parent. To create an offspring, we run Algorithm 6.1 on G and block the cut edges of Π_a and Π_b from contracting during the first multilevel cycle. Figure 6.8 demonstrates which edges are blocked for a particular instance. This leads to a coarsening phase that stops when no further contraction is possible unless the edge-cut increases or blocked edges are contracted. We then choose the best edge-cut clustering among three options: the two parents and a clustering where each node of the coarsest graph is assigned its own cluster. Our algorithm then uncoarsens this new clustering, applies local searches, and further improves the clustering with the global search

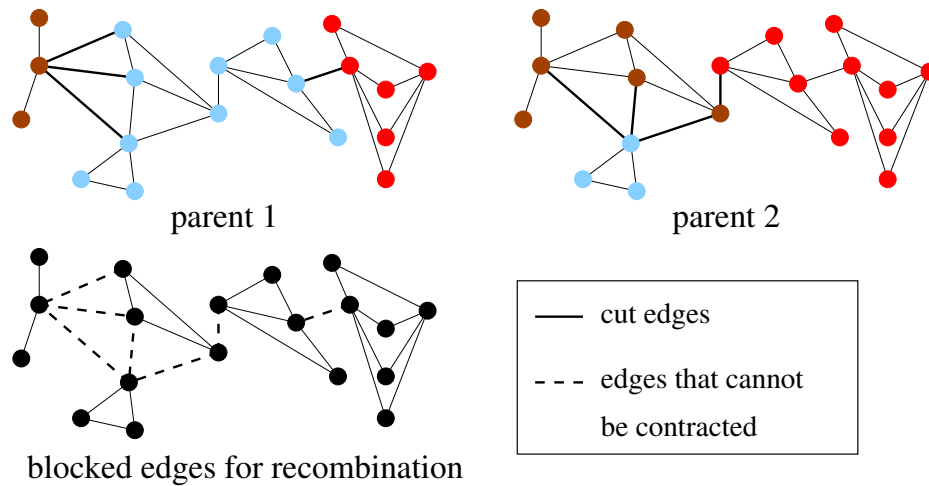


Figure 6.8. In the recombination, the cut edges of both parents cannot be contracted during the first multilevel cycle to build an offspring.

in Figure 6.7. This approach directly ensures non-increasing edge-cut compared to Π_a and Π_b and combines the characteristics of both parents, as expected from recombination.

6.2.2.4 Mutation

In this section, we explain the mutation operator in our algorithm. We randomly select an individual Π_a from the population and create a new individual by running Algorithm 6.1 on G while blocking Π_a 's cut edges from contracting only in the *first level* of the first multilevel cycle. The initial solution obtained at the finer level of this multilevel cycle can differ from Π_a but inherits some of its characteristics due to this edge-blocking constraint. The new individual is obtained by optimizing this initial solution during uncoarsening as well as the global search cycle. This operator does not guarantee non-increasing edge-cut compared to Π_a , although it increases population variability and optimizes the new individual as much as possible.

6.2.2.5 Replacement

In this section, we explain the replacement phase in our algorithm. It inserts the newly generated individual Π_c from mutation or recombination into the population P and removes an individual to keep only α individuals alive. We first check if Π_c has worse edge-cut than all individuals in P . If so, we do not insert Π_c into P . Otherwise, we remove the individual in P most similar to Π_c , where a high similarity is defined as a low-cardinality symmetric difference between the two individuals' edge sets. This approach helps maintain high diversity in P , which is beneficial for the evolutionary process.

6.2.2.6 Parallelization

In this section, we explain how our evolutionary algorithm can be executed across distributed processing elements. We follow the *coarse-grained* parallelization approach used in [160] for a similar problem. Each instance follows the same steps outlined in Algorithm 6.2. Communication is done through a variation of the *randomized rumor spreading* protocol [162]. The communication is organized in rounds, in which each process tries to send and receive individuals. In each round, a process sends the best individual from its local population to a randomly selected process it has not yet sent this individual to. Afterwards, it tries to receive incoming individuals and inserts them into its local population using the described replacement strategy. The overall Algorithm 6.2 is performed asynchronously, i.e., without global synchronization. The communication rounds can also happen only every x iterations to save communication time (we use $x = 1$ on our very fast communication network).

6.2.3 Experimental Evaluation

Setup. We performed the implementation of our algorithms and competing algorithms inside the KaHIP framework (using C++) and compiled them using gcc 9.3 with full optimization turned on (-O3 flag). We have used Machine C, which has sixteen cores.

Baselines. As the representative of the state-of-the-art for signed graph clustering, we use the general framework GASP [135], which is publicly available. In particular, the two best algorithms implemented in it with respect to the minimization of edge-cut, namely HCC-Sum [135] and GAEC [134]. As both of these algorithms have deterministic implementations, we only repeat each experiment with them once. We requested implementations from the authors of [136], but we have received no response by the time this manuscript was submitted. However, based on the z_value reported in [136] and preliminary experiments with GAEC, GAEC outperforms this method. Hence, we do not include it here.

Instances. In our experiments, we use all the real-world signed graphs listed in Table 2.4. Before running the experiments, we converted each of the graphs in Table 2.4 to an undirected signed graph without parallel or self edges. In particular, we achieve this by simply removing self edges and substituting all parallel and opposite arcs by a single undirected edge whose weight equals the sum of the weights of these arcs.

Methodology. Depending on the focus of the experiment, we measure running time and/or edge-cut. An alternative metric for edge-cut is the z_value [136], which can be mathematically defined as $1 - \frac{edgecut}{\omega(E^-)}$ where lower values are better and a value of 0 means a perfectly




Figure 6.9. Convergence plot for comparison of our memetic algorithm against repeated executions of our multilevel algorithm for all graphs in Table 2.4.

balanced clustering. Unless explicitly mentioned otherwise, we run our experiments three times on each of the graphs listed Table 2.4 with different random seeds, except for deterministic algorithms, which are run only once per graph. Since all values of edge-cut computed in our experiments are negative, we define its geometric mean as the geometric mean of its absolute value multiplied afterwards by -1 .

We present convergence plots, which are computed in the following way. Whenever a processing element creates a clustering it reports a pair (t, cut) , where the timestamp t is the currently elapsed time (in seconds) on the particular processing element and cut refers to the edge-cut of the computed clustering. After the completion of our algorithms on P processing elements, we are left with P sequences of pairs (t, cut) which we now merge into one sequence based on the timestamp. Then, we scan this sequence in ascending order of timestamp and remove all pairs with an edge-cut higher than then minimum edge-cut scanned so far. The resulting sequence defines a *step function* which indicates the best edge-cut found up to each timestamp. Since multiple repetitions are executed for each graph, we simply combine the step functions of all repetitions by computing the arithmetic mean of the best edge-cut found during all repetitions at each timestamp. When presenting convergence plots over multiple graphs, we normalize the elapsed time for each instance and then combine all step functions in the same way described above, but based on geometric mean rather than arithmetic mean. The normalization of the elapsed time is done as in [160] by dividing all timestamps by a normalization factor representing the approximate running time necessary to compute a clustering for the respective graph in one processing element. We define this normalization factor as the first (smallest) timestamp reported for an instance, which represents the the actual running time invested to create its first clustering.

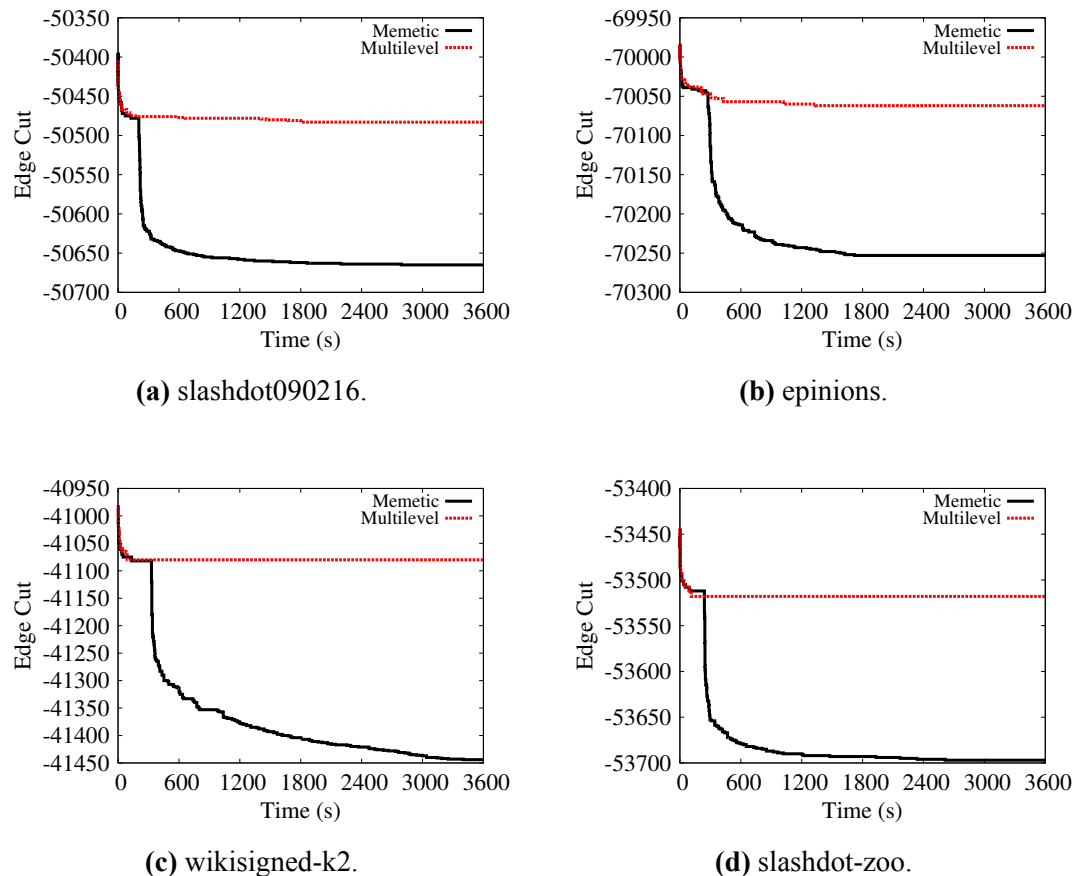


Figure 6.10. Convergence plot for comparison of our memetic algorithm against repeated executions of our multilevel algorithm for individual graphs.

6.2.3.1 Convergence

In this section, we evaluate the convergence of our memetic algorithm. For this experiment, we run our memetic algorithm for 1 hour on 16 processing elements without hyper-threading. As a baseline, we repeatedly run our multilevel algorithm (Algorithm 6.1) for the same amount of time in parallel on the same number of processing elements with different random seeds. Our objective is to check whether our evolutionary strategy is directly improving the global optimization process, or if it has no advantage over simply repeating our multilevel algorithm multiple times.

Figure 6.9 shows a convergence plot over all instances. The baseline improves edge-cut at first, but then stabilizes without major improvement. Our memetic algorithm performs similarly to the baseline for the first 10% of the running time, then outperforms it and continues to improve edge-cut. This is as expected, as we allocate roughly 10% of the running time to construct an initial population using multiple repetitions of our multilevel algorithm. After that, recombination and mutation operators are applied to the population. Roughly a

Graph	HCC-Sum	GAEC	Memetic
bitcoinalpha	-5 476	-5 529	-5 564
bitcoinote	-20 318	-20 391	-20 440
elec	-7 723	-7 717	-7 735
chess	-4 293	-4 312	-4 796
slashdot081106	-48 442	-48 454	-49 805
slashdot-zoo	-51 213	-51 806	-53 703
slashdot090216	-49 831	-49 797	-50 668
slashdot090221	-49 789	-49 798	-50 926
wikiconflict	-2 166 604	-2 166 608	-2 167 344
epinions	-68 759	-69 156	-70 323
wikisigned-k2	-41 377	-41 458	-41 623
Overall	-36 202	-36 320	-37 126

Table 6.2. Average edge-cut comparison.

similar pattern is observed for all graphs, as exemplified in Figure 6.10. The convergence plots of individual graphs show another pattern not visible in Figure 6.9. After about 10% of the running time, our memetic algorithm computes better solutions quickly, then continues to improve edge-cut at a slower pace. This shows that our recombination and mutation operators quickly improve solution quality and continue to do so at a slower pace over time even after this abrupt improvement.

6.2.3.2 Comparison against State-of-the-Art

In this section, we experimentally verify how much our evolutionary algorithm can improve over the state-of-the-art. For this experiment, we also run our memetic algorithm for 1 hour on 16 processing elements. In order to further diversify our initial population and since HCC-Sum and GAEC are fast algorithms, we include the solutions computed by HCC-Sum and GAEC in the initial population and compute its remaining individuals using our multilevel algorithm as usual. As a baseline, we use HCC-Sum and GAEC. Note that these algorithms are not randomized and thus can not be run multiple times to improve the result. Table 6.2 shows our overall results, in which our memetic algorithm computes better solutions than the state-of-the-art for every instance. Note that our memetic strategy is able to improve the initial population (which also contains results of the competing algorithms) significantly. Our best edge-cut gain over HCC-Sum and GAEC is obtained for slashdot-zoo, for which our average edge-cut is respectively 2 490 and 1 897 units smaller. In proportional terms, our largest increase in the absolute value of the edge-cut over HCC-Sum and GAEC is obtained for chess: 11.9% and 11.4%, respectively. A natural lower bound for the edge-cut of a graph is given by $\omega(E^-)$. Based on this bound, our algorithm gets the smallest absolute optimality gap for bitcoinalpha (3 786 units) and the smallest proportional

optimality gap for wikiconflict (8.35%). We conclude that our memetic strategy is effective in improving results of the initial population and overall our algorithm can be seen as the state-of-the-art.

6.3 Conclusion

In this chapter, we proposed two multilevel algorithms for graph decomposition. Both were engineered in all their details and were subjected to extensive experimental evaluation against the state-of-the-art.

First, we proposed a multilevel algorithm aimed at solving the process mapping problem. Our algorithm integrates graph partitioning and process mapping and comprises multiple versions with varying speed/quality trade-offs. Key ingredients of our algorithm include fast label propagation, more localized local search, initial partitioning, as well as a compressed data structure for computing processor distances without the need for storing a distance matrix. Experimental results demonstrate that our algorithm represents the new state-of-the-art for process mapping. Specifically, our algorithm generates superior or comparable overall solutions compared to any of the competing algorithms while being orders of magnitude faster than the previous best algorithm in terms of quality. Our improvements are primarily attributable to the integrated multilevel approach coupled with high-quality local search algorithms and initial mapping algorithms that divide the initial network based on the specified system hierarchy. Important future work entails parallelization, as well as the integration of global search schemes and different types of coarsening to further enhance solution quality. Additionally, we plan to investigate the impact of our algorithm on the real performance of applications such as sparse matrix-vector multiplications.

Second, we engineered a memetic algorithm for the signed graph clustering problem. Building upon a novel multilevel algorithm, we derive highly intuitive recombination and mutation operations. Experimental results show that our memetic strategy produces substantially superior solutions compared to the present state-of-the-art. In future work, we aim to enhance our algorithm by incorporating local searches based on maximum flows.

Chapter 7

Discussion

7.1 Conclusion

(Hyper)graph decomposition refers to a set of problems that focus on dividing large (hyper)graphs into smaller sub(hyper)graphs, which facilitates their analysis. The significance of this lies in its capability to enable efficient computation on large and intricate (hyper)graphs, including chemical compounds, social networks, and computer networks. In this dissertation, we propose multiple algorithmic contributions in the field of (hyper)graph decomposition, which utilize various techniques like buffered streaming, shared-memory parallelism, and efficient data structures to enhance the performance and quality of solutions. We provide a thorough description for each algorithm along with experimental results that demonstrate their superiority over existing state-of-the-art algorithms.

The first algorithm we propose is a buffered streaming algorithm for graph partitioning, which loads a batch of nodes and builds a model representing the loaded subgraph and already present partition structure. This model allows us to apply multilevel algorithms and compute high-quality solutions of massive graphs on inexpensive machines. We develop a multilevel algorithm to partition the model that optimizes an objective function, which improves on the previous state-of-the-art by removing the dependence on the number of blocks from the running time. Our algorithm computes better solutions than state-of-the-art using a small buffer size, and for larger numbers of blocks, it becomes faster than state-of-the-art.

The second algorithm is a shared-memory parallel streaming algorithm for the process mapping problem, which maps a streamed communication graph onto a hierarchical topology by performing recursive multi-sections on-the-fly. This algorithm can also be used to solve the graph partitioning problem as a general tool. Our algorithm has considerably lower running time complexity than existing state-of-the-art techniques and produces better process mappings.

The third algorithm is a streaming algorithm for hypergraph partitioning, which adapts the state-of-the-art streaming algorithm for graph partitioning by using an efficient data structure that makes the overall running time linearly dependent on the pin-count of the hypergraph and the memory consumption linearly dependent on the numbers of nets and blocks. Our algorithm outperforms all existing (buffered) streaming algorithms and even an in-memory algorithm with respect to both weighted number of cut hyperedges and connectivity measures.

Next, we propose algorithms for the local motif clustering problem, which builds a (hyper)graph model representing the motif-distribution around the seed node on the original graph. The first algorithm partitions the (hyper)graph model using a multi-level hypergraph or graph partitioner to minimize the motif conductance of the corresponding partition in the original graph. The second algorithm transforms the hypergraph model into a flow model based on the max-flow quotient-cut improvement algorithm to obtain a superior solution automatically. Our algorithms produce better communities than existing state-of-the-art techniques, while being up to multiple orders of magnitude faster.

We also propose multilevel algorithms for the process mapping problem, which include fast label propagation, more localized local search, initial partitioning, and a compressed data structure to compute processor distances without storing a distance matrix. Our algorithms are able to exploit a given hierarchical structure of the distributed system under consideration and obtain better solutions than existing state-of-the-art techniques.

Finally, we propose algorithms to solve the signed graph clustering problem by using some of the most effective techniques from graph partitioning that minimize edge-cut. Our multilevel algorithm includes a coarsening-uncoarsening process and efficient local search methods. We also introduce a distributed memetic algorithm that utilizes our multilevel algorithm and further enhances it with natural multilevel recombination and mutation operations. Experimental results demonstrate that our memetic algorithm outperforms the state-of-the-art with respect to edge-cut, producing significantly better solutions.

Overall, our contributions demonstrate the effectiveness of various algorithmic techniques, such as buffered streaming, shared-memory parallelism, and efficient data structures, in solving a range of important problems in graph and hypergraph decomposition. Our algorithms achieve improved performance and better solutions compared to state-of-the-art methods in various metrics, making them highly promising for practical applications.

7.2 Future Work

We have included specific future work in Chapters 4, 5, and 6. To enhance the performance of our algorithms, a general approach that could be beneficial is to parallelize them using

shared and/or distributed memory. This would lead to faster execution and enable them to handle even larger instances. In addition, it would be desirable to conduct experiments with more specialized benchmark sets. In particular, it would be worthwhile to verify the effectiveness of our streaming, local, and multilevel (hyper)graph decomposition algorithms within specific classes of (hyper)graphs, such as those that are used to model physical simulations, social networks, road networks, and other related phenomena. For motif clustering, more extensive experiments could be conducted, including larger and more diverse motifs, which could provide new insights into the structure of the problem. Another important objective is to investigate the impact of our algorithms on real-world applications. For example, our process mapping and (hyper)graph partitioning algorithms can be tested as preprocessing steps for the distributed computation of sparse matrix vector multiplication. Furthermore, extending our algorithms to dynamic (hyper)graphs, which is a current trend in the literature, could be accomplished in an easy manner, particularly for streaming algorithms. Finally, machine learning techniques could be applied to (hyper)graph decomposition problems, given their successful implementation in literature to solve other combinatorial problems.

Bibliography

- [1] Michael R. Garey, David S. Johnson, and Larry Stockmeyer. Some Simplified NP-Complete Problems. In *Proc. of the 6th ACM Sym. on Theory of Computing*, (STOC), pages 47–63. ACM, 1974. doi:10.1145/800119.803884.
- [2] Thang Nguyen Bui and Curt Jones. Finding Good Approximate Vertex and Edge Partitions is NP-Hard. *Information Processing Letters*, 42(3):153–159, 1992. doi:10.1016/0020-0190(92)90140-Q.
- [3] Charalampos Tsourakakis, Christos Gkantsidis, Bozidar Radunovic, and Milan Vojnovic. Fennel: Streaming graph partitioning for massive scale graphs. In *Proc. of the 7th ACM international conference on Web search and data mining*, pages 333–342, 2014. doi:10.1145/2556195.2556213.
- [4] Amel Awadelkarim and Johan Ugander. Prioritized restreaming algorithms for balanced graph partitioning. In *Proc. of the 26th ACM SIGKDD Intl. Conf. on Knowledge Discovery & Data Mining*, pages 1877–1887, 2020. doi:10.1145/3394486.3403239.
- [5] Nazanin Jafari, Oguz Selvitopi, and Cevdet Aykanat. Fast shared-memory streaming multilevel graph partitioning. *Journal of Parallel and Distributed Computing*, 147: 140–151, 2021. doi:10.1016/j.jpdc.2020.09.004.
- [6] Christian Mayer, Ruben Mayer, Muhammad Adnan Tariq, Heiko Geppert, Larissa Laich, Lukas Rieger, and Kurt Rothermel. Adwise: Adaptive window-based streaming edge partitioning for high-speed graph processing. In *2018 IEEE 38th Intl. Conf. on Distributed Computing Systems (ICDCS)*, pages 685–695. IEEE, 2018. doi:10.1109/ICDCS.2018.00072.
- [7] Loc Hoang, Roshan Dathathri, Gurbinder Gill, and Keshav Pingali. Cusp: A customizable streaming edge partitioner for distributed graph analytics. In *2019 IEEE Intl. Parallel and Distributed Processing Sym. (IPDPS)*, pages 439–450. IEEE, 2019. doi:10.1109/IPDPS.2019.00054.

- [8] Dan Alistarh, Jennifer Iglesias, and Milan Vojnovic. Streaming min-max hypergraph partitioning. In *Advances in Neural Information Processing Systems*, pages 1900–1908, 2015. doi:10.5555/2969442.2969452.
- [9] Fatih Taşyaran, Berkay Demireller, Kamer Kaya, and Bora Uçar. Streaming Hypergraph Partitioning Algorithms on Limited Memory Environments. In *HPCS 2020 - Intl. Conf. on High Performance Computing & Simulation*, pages 1–8. IEEE, 2021. URL <https://hal.archives-ouvertes.fr/hal-03182122>.
- [10] Marcelo Fonseca Faraj and Christian Schulz. Buffered streaming graph partitioning. *ACM J. Exp. Algorithmics*, 27:1.10:1–1.10:26, 2022. doi:10.1145/3546911. URL <https://doi.org/10.1145/3546911>.
- [11] Marcelo Fonseca Faraj and Christian Schulz. Recursive multi-section on the fly: Shared-memory streaming algorithms for hierarchical graph partitioning and process mapping. In *2022 IEEE Intl. Conf. on Cluster Computing (CLUSTER)*, pages 473–483, 2022. doi:10.1109/CLUSTER51413.2022.00057. URL <https://ieeexplore.ieee.org/document/9912716>.
- [12] Kamal Eyubov, Marcelo Fonseca Faraj, and Christian Schulz. FREIGHT: Fast Streaming Hypergraph Partitioning. In *Intl. Sym. on Experimental Algorithms (SEA)*, volume 265 of *LNCS*. Springer, to appear.
- [13] Ravi Kannan, Santosh Vempala, and Adrian Vetta. On clusterings: Good, bad and spectral. *JACM*, 51(3):497–515, 2004. doi:10.1145/990308.990313. URL <https://doi.org/10.1145/990308.990313>.
- [14] Alessandro Epasto, Jon Feldman, Silvio Lattanzi, Stefano Leonardi, and Vahab Mirrokni. Reduce and aggregate: similarity ranking in multi-categorical bipartite graphs. In *WWW*, pages 349–360, 2014. doi:10.1145/2566486.2568025. URL <https://doi.org/10.1145/2566486.2568025>.
- [15] Lucas G. S. Jeub, Prakash Balachandran, Mason A. Porter, Peter J. Mucha, and Michael W. Mahoney. Think locally, act locally: Detection of small, medium-sized, and large communities in large networks. *Phys. Rev. E*, 91:012821, Jan 2015. doi:10.1103/PhysRevE.91.012821. URL <https://link.aps.org/doi/10.1103/PhysRevE.91.012821>.
- [16] Konstantin Voevodski, Shang-Hua Teng, and Yu Xia. Spectral affinity in protein networks. *BMC systems biology*, 3(1):1–13, 2009. doi:10.1186/1752-0509-3-112. URL <https://doi.org/10.1186/1752-0509-3-112>.

- [17] Dorothea Wagner and Frank Wagner. Between Min Cut and Graph Bisection. In *Proc. of the 18th Intl. Sym. on Mathematical Foundations of Computer Science*, pages 744–750. Springer, 1993. doi:10.1007/3-540-57182-5_65.
- [18] Reid Andersen, Fan Chung, and Kevin Lang. Local graph partitioning using pagerank vectors. In *FOCS*, pages 475–486, 2006. doi:10.1109/FOCS.2006.44. URL <https://doi.org/10.1109/FOCS.2006.44>.
- [19] Jure Leskovec, Kevin J. Lang, Anirban Dasgupta, and Michael W. Mahoney. Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters. *Internet Mathematics*, 6(1):29–123, 2009. doi:10.1080/15427951.2009.10129177. URL <https://doi.org/10.1080/15427951.2009.10129177>.
- [20] Fan Chung and Olivia Simpson. Solving linear systems with boundary conditions using heat kernel pagerank. In *Intl. Workshop on Algorithms and Models for the Web-Graph*, pages 203–219. Springer, 2013. doi:10.1007/978-3-319-03536-9_16. URL https://doi.org/10.1007/978-3-319-03536-9_16.
- [21] Yixuan Li, Kun He, David Bindel, and John E. Hopcroft. Uncovering the small community structure in large networks: A local spectral approach. In *WWW*, pages 658–668, 2015. doi:10.1145/2736277.2741676. URL <https://doi.org/10.1145/2736277.2741676>.
- [22] Kimon Fountoulakis, Meng Liu, David F. Gleich, and Michael W. Mahoney. Flow-based algorithms for improving clusters: A unifying framework, software, and performance. *SIAM Rev.*, 65(1):59–143, 2023. doi:10.1137/20m1333055. URL <https://doi.org/10.1137/20m1333055>.
- [23] Hao Yin, Austin R. Benson, Jure Leskovec, and David F. Gleich. Local higher-order graph clustering. In *Proc. of the 23rd ACM SIGKDD Intl. Conf. on Knowledge Discovery and Data Mining*, pages 555–564. ACM, 2017. doi:10.1145/3097983.3098069. URL <https://doi.org/10.1145/3097983.3098069>.
- [24] Yunlei Zhang, Bin Wu, Yu Liu, and Jinna Lv. Local community detection based on network motifs. *Tsinghua Science and Technology*, 24(6):716–727, 2019. doi:10.26599/TST.2018.9010106. URL <https://doi.org/10.26599/TST.2018.9010106>.
- [25] Tao Meng, Lijun Cai, Tingqin He, Lei Chen, and Ziyun Deng. Local higher-order community detection based on fuzzy membership functions. *IEEE Access*, 7:128510–

- 128525, 2019. doi:10.1109/ACCESS.2019.2939535. URL <https://doi.org/10.1109/ACCESS.2019.2939535>.
- [26] Mrudula Murali, Katerina Potika, and Chris Pollett. Online local communities with motifs. In *2020 Second Intl. Conf. on Transdisciplinary AI (TransAI)*, pages 59–66. IEEE Computer Society, sep 2020. doi:10.1109/TransAI49837.2020.00014.
- [27] Ulrik Brandes, Daniel Delling, Marco Gaertler, Robert Gorke, Martin Hofer, Zoran Nikoloski, and Dorothea Wagner. On modularity clustering. *IEEE transactions on knowledge and data engineering*, 20(2):172–188, 2007. doi:10.1109/TKDE.2007.190689.
- [28] Arzum Karataş and Serap Şahin. Application areas of community detection: A review. In *2018 Intl. Congress on Big Data, Deep Learning and Fighting Cyber Terrorism (IBIGDELFT)*, pages 65–70, 2018. doi:10.1109/IBIGDELFT.2018.8625349.
- [29] Yoshiko Wakabayashi. *Aggregation of binary relations: algorithmic and polyhedral investigations*. PhD thesis, University of Augsburg, Germany, 1986. URL <https://d-nb.info/870902490>.
- [30] Kevin J. Lang and Satish Rao. A flow-based method for improving the expansion or conductance of graph cuts. In *Integer Programming and Combinatorial Optimization, 10th Intl. IPCO Conf.*, volume 3064 of *LNCS*, pages 325–337. Springer, 2004. doi:10.1007/978-3-540-25960-2_25. URL https://doi.org/10.1007/978-3-540-25960-2_25.
- [31] Austin R. Benson, David F. Gleich, and Jure Leskovec. Higher-order organization of complex networks. *Science*, 353(6295):163–166, 2016. doi:10.1126/science.aad9029. URL <https://ui.adsabs.harvard.edu/abs/2016Sci...353..163B/abstract>.
- [32] Barbara Brandfass, Thomas Alrutz, and Thomas Gerhold. Rank reordering for MPI communication optimization. *Computers & Fluids*, 80:372–380, 2013. URL <https://doi.org/10.1016/j.compfluid.2012.01.019>.
- [33] Lester Randolph Ford and Delbert Ray Fulkerson. Maximal flow through a network. *Canadian Journal of Mathematics*, 8:399–404, 1956. doi:10.4153/CJM-1956-045-5.
- [34] Andrew V. Goldberg and Robert Endre Tarjan. A new approach to the maximum-flow problem. *J. ACM*, 35(4):921–940, 1988. doi:10.1145/48014.61051. URL <https://doi.org/10.1145/48014.61051>.

- [35] Nate Veldt, Austin R. Benson, and Jon Kleinberg. Hypergraph cuts with general splitting functions. *SIAM Review*, 64(3):650–685, 2022. doi:10.1137/20M1321048. URL <https://doi.org/10.1137/20M1321048>.
- [36] Eugene L. Lawler. Cutsets and partitions of hypergraphs. *Networks*, 3(3):275–285, 1973. doi:10.1002/net.3230030306. URL <https://doi.org/10.1002/net.3230030306>.
- [37] Peter Sanders and Christian Schulz. Engineering Multilevel Graph Partitioning Algorithms. In *Proc. of the 19th European Symp. on Algorithms*, volume 6942 of *LNCS*, pages 469–480. Springer, 2011. doi:10.1007/978-3-642-23719-5_40.
- [38] Henning Meyerhenke, Peter Sanders, and Christian Schulz. Partitioning complex networks via size-constrained clustering. In *Experimental Algorithms - 13th International Symposium, SEA*, volume 8504 of *LNCS*, pages 351–363. Springer, 2014. doi:10.1007/978-3-319-07959-2_30. URL https://doi.org/10.1007/978-3-319-07959-2_30.
- [39] Efrén Mezura-Montes. *Evolutionary Computation. A Unified Approach*. kenneth a. de jong. (2006, MIT press.) 256 pages. *Artif. Life*, 13(4):423–426, 2007. doi:10.1162/artl.2007.13.4.423. URL <https://doi.org/10.1162/artl.2007.13.4.423>.
- [40] Kumara Sastry, David E. Goldberg, and Graham Kendall. Genetic algorithms. In *Search methodologies*, pages 93–117. 2014. URL http://eprints.ukh.ac.id/id/eprint/271/1/2014_Book_SearchMethodologies.pdf#page=106. Springer.
- [41] David E. Goldberg, Bradley Korb, and Kalyanmoy Deb. Messy genetic algorithms: Motivation, analysis, and first results. *Complex Syst.*, 3(5), 1989. URL http://www.complex-systems.com/abstracts/v03_i05_a05.html.
- [42] Zbigniew Michalewicz. *Genetic Algorithms + Data Structures = Evolution Programs, Third Revised and Extended Edition*. Springer, 1996. ISBN 978-3-540-60676-5. doi:10.1007/978-3-662-03315-9. URL <https://doi.org/10.1007/978-3-662-03315-9>.
- [43] Jure Leskovec and Andrej Krevl. SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>, June 2014.
- [44] Ryan A. Rossi and Nesreen K. Ahmed. The network data repository with interactive graph analytics and visualization. <http://networkrepository.com>, 2015.

- [45] David A. Bader, Henning Meyerhenke, Peter Sanders, Christian Schulz, Andrea Kappes, and Dorothea Wagner. Benchmarking for graph clustering and partitioning. In *Encyclopedia of Social Network Analysis and Mining*, pages 73–82. 2014. doi:10.1007/978-1-4614-6170-8_23. URL https://doi.org/10.1007/978-1-4614-6170-8_23.
- [46] Manuel Holtgrewe, Peter Sanders, and Christian Schulz. Engineering a scalable high quality graph partitioner. In *24th IEEE International Symposium on Parallel and Distributed Processing, IPDPS*, pages 1–12. IEEE, 2010. doi:10.1109/IPDPS.2010.5470485. URL <https://doi.org/10.1109/IPDPS.2010.5470485>.
- [47] Daniel Funke, Sebastian Lamm, Ulrich Meyer, Manuel Penschuck, Peter Sanders, Christian Schulz, Darren Strash, and Moritz von Looz. Communication-free massively distributed graph generation. *Journal of Parallel and Distributed Computing*, 131:200–217, 2019.
- [48] Manuel Penschuck, Ulrik Brandes, Michael Hamann, Sebastian Lamm, Ulrich Meyer, Ilya Safro, Peter Sanders, and Christian Schulz. Recent advances in scalable network generation. In David A. Bader, editor, *Massive Graph Analytics*. Chapman and Hall/CRC, 2022.
- [49] Alan J. Soper, Chris Walshaw, and Mark Cross. A combined evolutionary search and multilevel optimisation approach to graph-partitioning. *J. Glob. Optim.*, 29(2):225–241, 2004. doi:10.1023/B:JOGO.0000042115.44455.f3. URL <https://doi.org/10.1023/B:JOGO.0000042115.44455.f3>.
- [50] Timothy A. Davis and Yifan Hu. The university of florida sparse matrix collection. *ACM Trans. Math. Softw.*, 38(1):1:1–1:25, 2011. doi:10.1145/2049662.2049663. URL <https://doi.org/10.1145/2049662.2049663>.
- [51] Daniel Delling, Peter Sanders, Dominik Schultes, and Dorothea Wagner. Engineering route planning algorithms. In *Algorithmics of Large and Complex Networks - Design, Analysis, and Simulation [DFG priority program 1126]*, volume 5515 of *LNCS*, pages 117–139. Springer, 2009. doi:10.1007/978-3-642-02094-0_7. URL https://doi.org/10.1007/978-3-642-02094-0_7.
- [52] Sebastian Schlag, Vitali Henne, Tobias Heuer, Henning Meyerhenke, Peter Sanders, and Christian Schulz. k -way hypergraph partitioning via n -level recursive bisection.

- In *Proceedings of the Eighteenth Workshop on Algorithm Engineering and Experiments, ALENEX*, pages 53–67. SIAM, 2016. doi:10.1137/1.9781611974317.5. URL <https://doi.org/10.1137/1.9781611974317.5>.
- [53] Charles J. Alpert. The ISPD98 circuit benchmark suite. In *Proc. of the 1998 Intl. Sym. on Physical Design, ISPD 1998, Monterey, CA, USA, April 6-8, 1998*, pages 80–85. ACM, 1998. doi:10.1145/274535.274546. URL <https://doi.org/10.1145/274535.274546>.
- [54] Anton Belov, Daniel Diepold, Marijn Heule, and Matti Järvisalo. The sat competition 2014. <http://www.satcompetition.org/2014/>, 2014.
- [55] Jure Leskovec. Stanford Network Analysis Package (SNAP), 2013.
- [56] Jérôme Kunegis. KONECT: the koblenz network collection. In *22nd Intl. World Wide Web Conf., WWW '13, Rio de Janeiro, Brazil, May 13-17, 2013, Companion Volume*, pages 1343–1350. Intl. World Wide Web Conferences Steering Committee / ACM, 2013. doi:10.1145/2487788.2488173. URL <https://doi.org/10.1145/2487788.2488173>.
- [57] Ümit V. Çatalyürek and Cevdet Aykanat. Patoh (partitioning tool for hypergraphs). In *Encyclopedia of Parallel Computing*, pages 1479–1487. Springer, 2011. doi:10.1007/978-0-387-09766-4_93. URL https://doi.org/10.1007/978-0-387-09766-4_93.
- [58] George Karypis and Vipin Kumar. Parallel multilevel k -way partitioning scheme for irregular graphs. In *Proc. of the 1996 ACM/IEEE Conf. on Supercomputing, Supercomputing '96*. IEEE Computer Society, 1996. doi:10.1145/369028.369103.
- [59] George Karypis and Vipin Kumar. Multilevel k -way hypergraph partitioning. In *Proceedings of the 36th Conference on Design Automation*, pages 343–348. ACM Press, 1999. doi:10.1145/309847.309954. URL <https://doi.org/10.1145/309847.309954>.
- [60] François Pellegrini and Jean Roman. Experimental analysis of the dual recursive bipartitioning algorithm for static mapping. Technical report, TR 1038-96, LaBRI, URA CNRS 1304, Univ. Bordeaux I, 1996.
- [61] Christian Mayer, Ruben Mayer, Sukanya Bhowmik, Lukas Epple, and Kurt Rothermel. HYPE: massive hypergraph partitioning with neighborhood expansion. In *IEEE Intl. Conf. on Big Data (IEEE BigData 2018)*, pages 458–467. IEEE, 2018.

- doi:10.1109/BigData.2018.8621968. URL <https://doi.org/10.1109/BigData.2018.8621968>.
- [62] Peter Sanders and Christian Schulz. Think Locally, Act Globally: Highly Balanced Graph Partitioning. In *12th Intl. Sym. on Experimental Algorithms (SEA)*, LNCS. Springer, 2013. doi:10.1007/978-3-642-38527-8_16.
- [63] Lars Gottesbüren, Tobias Heuer, Peter Sanders, Christian Schulz, and Daniel Seemaier. Deep multilevel graph partitioning. In *29th Annual European Sym. on Algorithms, ESA 2021, September 6-8, 2021, Lisbon, Portugal (Virtual Conf.)*, volume 204 of *LIPICs*, pages 48:1–48:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.ESA.2021.48. URL <https://doi.org/10.4230/LIPICs.ESA.2021.48>.
- [64] Lars Gottesbüren, Tobias Heuer, Peter Sanders, and Sebastian Schlag. Scalable Shared-Memory Hypergraph Partitioning. In *Proc. of the Sym. on Algorithm Engineering and Experiments ALENEX*, pages 16–30, 2021. doi:10.1137/1.9781611976472.2.
- [65] Yaroslav Akhremtsev, Peter Sanders, and Christian Schulz. High-quality shared-memory graph partitioning. In *Euro-Par 2018: Parallel Processing - 24th Intl. Conf. on Parallel and Distributed Computing*, pages 659–671, 2018. doi:10.1007/978-3-319-96983-1_47. URL https://doi.org/10.1007/978-3-319-96983-1_47.
- [66] Ümit V. Çatalyürek, Karen D. Devine, Marcelo Fonseca Faraj, Lars Gottesbüren, Tobias Heuer, Henning Meyerhenke, Peter Sanders, Sebastian Schlag, Christian Schulz, Daniel Seemaier, and Dorothea Wagner. More recent advances in (hyper)graph partitioning. *ACM Computing Surveys*, 55:1–38, 2023. doi:doi.org/10.1145/3571808.
- [67] Aydın Buluç, Henning Meyerhenke, Ilya Safro, Peter Sanders, and Christian Schulz. *Recent Advances in Graph Partitioning*, pages 117–158. Springer Intl. Publishing, Cham, 2016. doi:10.1007/978-3-319-49487-6_4.
- [68] Christian Schulz and Darren Strash. Graph partitioning: Formulations and applications to big data. In *Encyclopedia of Big Data Technologies*. Springer, 2019. doi:10.1007/978-3-319-63962-8_312-2. URL https://doi.org/10.1007/978-3-319-63962-8_312-2.
- [69] Isabelle Stanton and Gabriel Kliot. Streaming graph partitioning for large distributed graphs. In *Proc. of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1222–1230, 2012. doi:<https://doi.org/10.1145/2339530.2339722>.

- [70] Isabelle Stanton. Streaming balanced graph partitioning algorithms for random graphs. In *Proc. of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1287–1301. SIAM, 2014. doi:10.1137/1.9781611973402.95.
- [71] Fan Chung, Shirin Handjani, and Doug Jungreis. Generalizations of polya’s urn problem. *Annals of combinatorics*, 7(2):141–153, 2003. URL <https://link.springer.com/content/pdf/10.1007/s00026-003-0178-y.pdf>.
- [72] Wei Zhang, Yong Chen, and Dong Dai. Akin: A streaming graph partitioning algorithm for distributed graph storage systems. In *2018 18th IEEE/ACM Intl. Sym. on Cluster, Cloud and Grid Computing (CCGRID)*, pages 183–192. IEEE, 2018. doi:10.1109/CCGRID.2018.00033.
- [73] Lieve Hamers et al. Similarity measures in scientometric research: The jaccard index versus salton’s cosine formula. *Information Processing and Management*, 25(3):315–18, 1989.
- [74] Joel Nishimura and Johan Ugander. Restreaming graph partitioning: simple versatile algorithms for advanced balancing. In *Proc. of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1106–1114, 2013. doi:10.1145/2487575.2487696.
- [75] Md Anwarul Kaium Patwary, Saurabh Garg, and Byeong Kang. Window-based streaming graph partitioning algorithm. In *Proc. of the Australasian Computer Science Week Multiconference*, pages 1–10, 2019. doi:10.1145/3290688.3290711.
- [76] Austin R. Benson, David F. Gleich, and Jure Leskovec. Tensor spectral clustering for partitioning higher-order network structures. In *Proc. of the 2015 SIAM Intl. Conf. on Data Mining*, pages 118–126. SIAM, 2015. doi:10.1137/1.9781611974010.14. URL <https://doi.org/10.1137/1.9781611974010.14>.
- [77] Christine Klymko, David Gleich, and Tamara G Kolda. Using triangles to improve community detection in directed networks. *arXiv preprint arXiv:1404.5874*, 2014. URL <https://arxiv.org/abs/1404.5874>.
- [78] Natasa Przulj. Biological network comparison using graphlet degree distribution. *Bioinformatics*, 23(2):e177–e183, 2007. URL <https://doi.org/10.1093/bioinformatics/btl301>.
- [79] Charalampos E. Tsourakakis, Jakub Pachocki, and Michael Mitzenmacher. Scalable motif-aware graph clustering. In *WWW*, pages 1451–1460, 2017.

- doi:10.1145/3038912.3052653. URL <https://doi.org/10.1145/3038912.3052653>.
- [80] Kyle Kloster and David F. Gleich. Heat kernel based community detection. In *ACM SIGKDD*, pages 1386–1395, 2014. doi:10.1145/2623330.2623706. URL <https://doi.org/10.1145/2623330.2623706>.
- [81] Michael W. Mahoney, Lorenzo Orecchia, and Nisheeth K. Vishnoi. A local spectral method for graphs: With applications to improving graph partitions and exploring data graphs locally. *Journal of Machine Learning Research*, 13(1):2339–2365, 2012. URL <http://jmlr.org/papers/v13/mahoney12a.html>.
- [82] Wanyun Cui, Yanghua Xiao, Haixun Wang, and Wei Wang. Local search of communities in large graphs. In *ACM SIGMOD Intl. Conf. on Management of data*, pages 991–1002, 2014. doi:10.1145/2588555.2612179. URL <https://doi.org/10.1145/2588555.2612179>.
- [83] Mauro Sozio and Aristides Gionis. The community-search problem and how to plan a successful cocktail party. In *ACM SIGKDD*, pages 939–948, 2010. doi:10.1145/1835804.1835923. URL <https://doi.org/10.1145/1835804.1835923>.
- [84] Karl Rohe and Tai Qin. The blessing of transitivity in sparse and stochastic networks. *arXiv preprint arXiv:1307.2302*, 2013. URL <https://arxiv.org/abs/1307.2302>.
- [85] Xin Huang, Hong Cheng, Lu Qin, Wentao Tian, and Jeffrey Xu Yu. Querying k-truss community in large and dynamic graphs. In *ACM SIGMOD*, pages 1311–1322, 2014. doi:10.1145/2588555.2610495. URL <https://doi.org/10.1145/2588555.2610495>.
- [86] Jonathan Cohen. Trusses: Cohesive subgraphs for social network analysis. *National security agency Tech. report*, 16(3.1), 2008. URL <https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.505.7006&rep=rep1&type=pdf>.
- [87] Dawei Zhou, Si Zhang, Mehmet Yigit Yildirim, Scott Alcorn, Hanghang Tong, Hasan Davulcu, and Jingrui He. High-order structure exploration on massive graphs: A local graph clustering perspective. *ACM TKDD*, 15(2):1–26, 2021. doi:10.1145/3425637. URL <https://doi.org/10.1145/3425637>.
- [88] Daniel A. Spielman and Shang-Hua Teng. A local clustering algorithm for massive graphs and its application to nearly linear time graph partitioning. *SIAM J. Comp.*,

- 42(1):1–26, 2013. doi:10.1137/080744888. URL <https://doi.org/10.1137/080744888>.
- [89] Ronghua Shang, Weitong Zhang, Jingwen Zhang, Jie Feng, and Licheng Jiao. Local community detection based on higher-order structure and edge information. *Physica A: Statistical Mechanics and its Applications*, 587:126513, 2022. doi:10.1016/j.physa.2021.126513. URL <https://doi.org/10.1016/j.physa.2021.126513>.
- [90] Chris Walshaw and Mark Cross. Multilevel mesh partitioning for heterogeneous communication networks. *Future Gener. Comput. Syst.*, 17(5):601–623, 2001. doi:10.1016/S0167-739X(00)00107-2. URL [https://doi.org/10.1016/S0167-739X\(00\)00107-2](https://doi.org/10.1016/S0167-739X(00)00107-2).
- [91] François Pellegrini and Jean Roman. SCOTCH: A software package for static mapping by dual recursive bipartitioning of process and architecture graphs. In *High-Performance Computing and Networking*, volume 1067 of *LNCS*, pages 493–498. Springer, 1996. doi:10.1007/3-540-61142-8_588.
- [92] Christian Schulz and Jesper Larsson Träff. Better process mapping and sparse quadratic assignment. In *16th Intl. Sym. on Experimental Algorithms*, volume 75 of *LIPICs*, pages 4:1–4:15, 2017. doi:10.4230/LIPICs.SEA.2017.4.
- [93] Charles H. Heider. A computationally simplified pair-exchange algorithm for the quadratic assignment problem. Technical report, DTIC Document, 1972.
- [94] Heiner Müller-Merbach. *Optimale reihenfolgen*, volume 15 of *Ökonometrie und Unternehmensforschung*. Springer-Verlag, 1970.
- [95] Roland Glantz, Hening Meyerhenke, and Alexander Noe. Algorithms for mapping parallel processes onto grid and torus architectures. In *2015 23rd Euromicro Intl. Conf. on Parallel, Distributed, and Network-Based Processing*, pages 236–243. IEEE, 2015. doi:10.1109/PDP.2015.21.
- [96] Charles-Edmond Bichot and Patrick Siarry, editors. *Graph Partitioning*. Wiley, 2013. doi:10.1002/9781118601181.
- [97] Mehmet Deveci, Kamer Kaya, Bora Uçar, and Ümit V. Çatalyürek. Fast and high quality topology-aware task mapping. In *2015 IEEE Intl. Parallel and Distributed Processing Sym., IPDPS 2015, Hyderabad, India, May 25-29, 2015*, pages 197–206. IEEE Computer Society, 2015. doi:10.1109/IPDPS.2015.93. URL <https://doi.org/10.1109/IPDPS.2015.93>.

- [98] François Pellegrini. Scotch Home Page. <http://www.labri.fr/pelegrin/scotch>.
- [99] Chris Walshaw and Mark Cross. JOSTLE: Parallel Multilevel Graph-Partitioning Software – An Overview. In *Mesh Partitioning Techniques and Domain Decomposition Techniques*, pages 27–58. 2007. ISBN 978-1-874672-29-6.
- [100] Mehmet Deveci, Karen D. Devine, Kevin Pedretti, Mark A. Taylor, Sivasankaran Rajamanickam, and Ümit V. Çatalyürek. Geometric mapping of tasks to processors on parallel computers with mesh or torus networks. *IEEE Transactions on Parallel and Distributed Systems*, 30(9):2018–2032, 2019. doi:10.1109/TPDS.2019.2900043.
- [101] Mehmet Deveci, Sivasankaran Rajamanickam, Karen D. Devine, and Ümit V. Çatalyürek. Multi-jagged: A scalable parallel spatial partitioning algorithm. *IEEE Trans. Parallel Distributed Syst.*, 27(3):803–817, 2016. doi:10.1109/TPDS.2015.2412545. URL <https://doi.org/10.1109/TPDS.2015.2412545>.
- [102] Joshua T. Vogelstein, John M. Conroy, Vince Lyzinski, Louis J. Podrazik, Steven G. Kratzer, Eric T. Harley, Donniell E. Fishkind, R. Jacob Vogelstein, and Carey E. Priebe. Fast approximate quadratic programming for graph matching. *PLOS one*, 10(4):e0121002, 2015. doi:10.1371/journal.pone.0121002.
- [103] Shinji Umeyama. An eigendecomposition approach to weighted graph matching problems. *IEEE transactions on pattern analysis and machine intelligence*, 10(5):695–703, 1988. doi:10.1109/34.6778.
- [104] Rohit Singh, Jinbo Xu, and Bonnie Berger. Pairwise global alignment of protein interaction networks by matching neighborhood topology. In *Annual international conference on research in computational molecular biology*, pages 16–31. Springer, 2007. doi:10.1007/978-3-540-71681-5_2. URL https://doi.org/10.1007/978-3-540-71681-5_2.
- [105] Mikhail Zaslavskiy, Francis Bach, and Jean-Philippe Vert. A path following algorithm for the graph matching problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(12):2227–2242, 2008. doi:10.1109/TPAMI.2008.245.
- [106] Rainer E Burkard, Stefan E Karisch, and Franz Rendl. Qaplib—a quadratic assignment problem library. *Journal of Global optimization*, 10(4):391–403, 1997. doi:10.1023/A:1008293323270. URL <https://doi.org/10.1023/A:1008293323270>.

- [107] Roland Glantz, Maria Predari, and Henning Meyerhenke. Topology-induced enhancement of mappings. In *Proc. of the 47th Intl. Conf. on Parallel Processing*, pages 1–10, 2018. doi:10.1145/3225058.3225117.
- [108] Peter Sanders and Christian Schulz. KaHIP – Karlsruhe High Quality Partitioning Homepage. <http://algo2.iti.kit.edu/documents/kahip/index.html>.
- [109] Konrad Von Kirchbach, Christian Schulz, and Jesper Larsson Träff. Better process mapping and sparse quadratic assignment. *Journal of Experimental Algorithmics (JEA)*, 25:1–19, 2020. doi:10.1145/3409667.
- [110] Maria Predari, Charilaos Tzovas, Christian Schulz, and Henning Meyerhenke. An mpi-based algorithm for mapping complex networks onto hierarchical architectures. In *Euro-Par 2021: Parallel Processing - 27th Intl. Conf. on Parallel and Distributed Computing, Lisbon, Portugal, September 1-3, 2021, Proc.*, volume 12820 of *LNCS*, pages 167–182. Springer, 2021. doi:10.1007/978-3-030-85665-6_11. URL https://doi.org/10.1007/978-3-030-85665-6_11.
- [111] Henning Meyerhenke, Peter Sanders, and Christian Schulz. Parallel graph partitioning for complex networks. *IEEE Trans. Parallel Distrib. Syst.*, 28(9):2625–2638, 2017. doi:10.1109/TPDS.2017.2671868. URL <https://doi.org/10.1109/TPDS.2017.2671868>.
- [112] Konrad von Kirchbach, Markus Lehr, Sascha Hunold, Christian Schulz, and Jesper Larsson Träff. Efficient process-to-node mapping algorithms for stencil computations. In *IEEE Intl. Conf. on Cluster Computing (CLUSTER)*, pages 1–11. IEEE, 2020. doi:10.1109/CLUSTER49012.2020.00011.
- [113] Christian Schulz and Jesper Larsson Träff. VieM – Vienna Process Mapping and Sparse Quadratic Assignment. <http://viem.taa.univie.ac.at/>.
- [114] Bo Yang, William K. Cheung, and Jiming Liu. Community mining from signed social networks. *IEEE Trans. Knowl. Data Eng.*, 19(10):1333–1348, 2007. doi:10.1109/TKDE.2007.1061. URL <https://doi.org/10.1109/TKDE.2007.1061>.
- [115] Shiwei Che, Wu Yang, and Wei Wang. A memetic algorithm for community detection in signed networks. *IEEE Access*, 8:123585–123602, 2020. doi:10.1109/ACCESS.2020.3006108. URL <https://doi.org/10.1109/ACCESS.2020.3006108>.

- [116] Vincent A Traag and Jeroen Bruggeman. Community detection in networks with positive and negative links. *Physical Review E*, 80(3):036115, 2009. doi:10.1103/PhysRevE.80.036115. URL <https://doi.org/10.1103/PhysRevE.80.036115>.
- [117] Nikhil Bansal, Avrim Blum, and Shuchi Chawla. Correlation clustering. *Mach. Learn.*, 56(1-3):89–113, 2004. doi:10.1023/B:MACH.0000033116.57574.95. URL <https://doi.org/10.1023/B:MACH.0000033116.57574.95>.
- [118] Erik D. Demaine, Dotan Emanuel, Amos Fiat, and Nicole Immorlica. Correlation clustering in general weighted graphs. *Theor. Comput. Sci.*, 361(2-3):172–187, 2006. doi:10.1016/j.tcs.2006.05.008. URL <https://doi.org/10.1016/j.tcs.2006.05.008>.
- [119] Martin Grötschel and Yoshiko Wakabayashi. A cutting plane algorithm for a clustering problem. *Math. Program.*, 45(1-3):59–96, 1989. doi:10.1007/BF01589097. URL <https://doi.org/10.1007/BF01589097>.
- [120] Thang N. Dinh and My T. Thai. Toward optimal community detection: From trees to general weighted networks. *Internet Math.*, 11(3):181–200, 2015. doi:10.1080/15427951.2014.950875. URL <https://doi.org/10.1080/15427951.2014.950875>.
- [121] Atsushi Miyauchi and Noriyoshi Sukegawa. Redundant constraints in the standard formulation for the clique partitioning problem. *Optim. Lett.*, 9(1):199–207, 2015. doi:10.1007/s11590-014-0754-6. URL <https://doi.org/10.1007/s11590-014-0754-6>.
- [122] Atsushi Miyauchi, Tomohiro Sonobe, and Noriyoshi Sukegawa. Exact clustering via integer programming and maximum satisfiability. In *Proc. of the Thirty-Second AAAI Conf. on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Sym. on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 1387–1394. AAAI Press, 2018. URL <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16176>.
- [123] Jérôme Kunegis, Stephan Schmidt, Andreas Lommatzsch, Jürgen Lerner, Ernesto William De Luca, and Sahin Albayrak. Spectral analysis of signed graphs for clustering, prediction and visualization. In *Proc. of the SIAM Intl. Conf. on Data Mining, SDM 2010, April 29 - May 1, 2010, Columbus, Ohio, USA*, pages 559–570. SIAM,

2010. doi:10.1137/1.9781611972801.49. URL <https://doi.org/10.1137/1.9781611972801.49>.
- [124] Jean Gallier. Spectral theory of unsigned and signed graphs. applications to graph clustering: a survey. *CoRR*, abs/1601.04692, 2016. URL <http://arxiv.org/abs/1601.04692>.
- [125] Andrew Knyazev. On spectral partitioning of signed graphs. In *Proc. of the Eighth SIAM Workshop on Combinatorial Scientific Computing, CSC 2018, Bergen, Norway, June 6-8, 2018*, pages 11–22. SIAM, 2018. doi:10.1137/1.9781611975215.2. URL <https://doi.org/10.1137/1.9781611975215.2>.
- [126] Pedro Mercado, Francesco Tudisco, and Matthias Hein. Spectral clustering of signed graphs via matrix power means. In *Proc. of the 36th Intl. Conf. on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proc. of Machine Learning Research*, pages 4526–4536. PMLR, 2019. URL <http://proceedings.mlr.press/v97/mercado19a.html>.
- [127] Dongxiao He, Jie Liu, Dayou Liu, Di Jin, and Zhengxue Jia. Ant colony optimization for community detection in large-scale complex networks. In *Seventh Intl. Conf. on Natural Computation, ICNC 2011, Shanghai, China, 26-28 July, 2011*, pages 1151–1155. IEEE, 2011. doi:10.1109/ICNC.2011.6022234. URL <https://doi.org/10.1109/ICNC.2011.6022234>.
- [128] Xiaoyu Zhu, Yinghong Ma, and Zhiyuan Liu. A novel evolutionary algorithm on communities detection in signed networks. *Physica A: Statistical Mechanics and its Applications*, 503:938–946, 2018. doi:10.1016/j.physa.2018.08.112. URL <https://doi.org/10.1016/j.physa.2018.08.112>.
- [129] Michael J. Brusco and Patrick Doreian. Partitioning signed networks using relocation heuristics, tabu search, and variable neighborhood search. *Soc. Networks*, 56:70–80, 2019. doi:10.1016/j.socnet.2018.08.007. URL <https://doi.org/10.1016/j.socnet.2018.08.007>.
- [130] Mayasa M Abdulrahman, Amenah Dahim Abood, and Baraa A Attea. An enhanced multi-objective evolutionary algorithm with decomposition for signed community detection problem. In *2020 2nd Annual Intl. Conf. on Information and Sciences (AiCIS)*, pages 45–50. IEEE, 2020. doi:10.1109/AiCIS51645.2020.00017. URL <https://doi.org/10.1109/AiCIS51645.2020.00017>.

- [131] Yiran Chen, Qinma Kang, Wenqiang Duan, Yunfan Shan, Ran Xiao, and Yunfan Kang. An iterated local search algorithm for community detection in signed networks. *Intl. Journal of Modern Physics C*, page 2250105, 2022. doi:10.1142/S0129183122501054. URL <https://doi.org/10.1142/S0129183122501054>.
- [132] Jiliang Tang, Yi Chang, Charu C. Aggarwal, and Huan Liu. A survey of signed network mining in social media. *ACM Comput. Surv.*, 49(3):42:1–42:37, 2016. doi:10.1145/2956185. URL <https://doi.org/10.1145/2956185>.
- [133] Maria Tomasso, Lucas J. Rusnak, and Jelena Tesic. Advances in scaling community discovery methods for signed graph networks. *J. Complex Networks*, 10(3), 2022. doi:10.1093/comnet/cnac013. URL <https://doi.org/10.1093/comnet/cnac013>.
- [134] Margret Keuper, Evgeny Levinkov, Nicolas Bonneel, Guillaume Lavoué, Thomas Brox, and Bjoern Andres. Efficient decomposition of image and mesh graphs by lifted multicuts. In *2015 IEEE Intl. Conf. on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 1751–1759. IEEE Computer Society, 2015. doi:10.1109/ICCV.2015.204. URL <https://doi.org/10.1109/ICCV.2015.204>.
- [135] Alberto Bailoni, Constantin Pape, Nathan Hütsch, Steffen Wolf, Thorsten Beier, Anna Kreshuk, and Fred A. Hamprecht. GASP, a generalized framework for agglomerative clustering of signed graphs and its application to instance segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 11635–11645. IEEE, 2022. doi:10.1109/CVPR52688.2022.01135. URL <https://doi.org/10.1109/CVPR52688.2022.01135>.
- [136] Jialin Hua, Jian Yu, and Miin-Shen Yang. Fast clustering for signed graphs based on random walk gap. *Soc. Networks*, 60:113–128, 2020. doi:10.1016/j.socnet.2018.08.008. URL <https://doi.org/10.1016/j.socnet.2018.08.008>.
- [137] Sergio Gómez, Pablo Jensen, and Alex Arenas. Analysis of community structure in networks of correlated data. *Physical Review E*, 80(1):016114, 2009. doi:10.1103/PhysRevE.80.016114. URL <https://doi.org/10.1103/PhysRevE.80.016114>.
- [138] Miyuki Koshimura, Emi Watanabe, Yuko Sakurai, and Makoto Yokoo. Concise integer linear programming formulation for clique partitioning problems. *Constraints*

- An Int. J.*, 27(1):99–115, 2022. doi:10.1007/s10601-022-09326-z. URL <https://doi.org/10.1007/s10601-022-09326-z>.
- [139] George Karypis and Vipin Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM J. Sci. Comput.*, 20(1):359–392, 1998. doi:10.1137/S1064827595287997. URL <https://doi.org/10.1137/S1064827595287997>.
- [140] Usha Nandini Raghavan, Réka Albert, and Soundar Kumara. Near linear time algorithm to detect community structures in large-scale networks. *Phys. Rev. E*, 76:036106, Sep 2007. doi:10.1103/PhysRevE.76.036106. URL <https://link.aps.org/doi/10.1103/PhysRevE.76.036106>.
- [141] Sebastian Schlag, Tobias Heuer, Lars Gottesebüren, Yaroslav Akhremtsev, Christian Schulz, and Peter Sanders. High-quality hypergraph partitioning. *ACM Journal of Experimental Algorithms (JEA)*, 2022. doi:10.1145/3529090. URL <https://doi.org/10.1145/3529090>.
- [142] Adil Chhabra, Marcelo Fonseca Faraj, and Christian Schulz. Local motif clustering via (hyper)graph partitioning. In *2023 Proc. of the Sym. on Algorithm Engineering and Experiments (ALENEX)*, pages 96–109. doi:10.1137/1.9781611977561.ch9. URL <https://epubs.siam.org/doi/abs/10.1137/1.9781611977561.ch9>.
- [143] Adil Chhabra, Marcelo Fonseca Faraj, and Christian Schulz. Local motif clustering via (hyper)graph partitioning (extended abstract). In *Proc. of the Fifteenth Intl. Sym. on Combinatorial Search, SOCS 2022, Vienna, Austria, July 21-23, 2022*, pages 261–263. AAAI Press, 2022. doi:10.1609/socs.v15i1.21779. URL <https://ojs.aaai.org/index.php/SOCS/article/view/21779>.
- [144] Adil Chhabra, Marcelo Fonseca Faraj, and Christian Schulz. Faster local motif clustering via maximum flows. *CoRR*, abs/2301.07145, 2023. doi:10.48550/arXiv.2301.07145. URL <https://doi.org/10.48550/arXiv.2301.07145>.
- [145] Ronald C Read and Derek G Corneil. The graph isomorphism disease. *Journal of graph theory*, 1(4):339–363, 1977. doi:10.1002/jgt.3190010410. URL <https://doi.org/10.1002/jgt.3190010410>.
- [146] Mark Ortman and Ulrik Brandes. Triangle listing algorithms: Back from the diversion. In *ALENEX*, pages 1–8. SIAM, 2014. doi:10.1137/1.9781611973198.1. URL <https://doi.org/10.1137/1.9781611973198.1>.

- [147] Petter Holme and Beom Jun Kim. Growing scale-free networks with tunable clustering. *Physical review E*, 65(2):026107, 2002. doi:10.1103/PhysRevE.65.026107. URL <https://doi.org/10.1103/PhysRevE.65.026107>.
- [148] Vladimir Batagelj and Matjaž Zaveršnik. Short cycle connectivity. *Discrete Mathematics*, 307(3-5):310–318, 2007. doi:10.1016/j.disc.2005.09.051. URL <https://doi.org/10.1016/j.disc.2005.09.051>.
- [149] Arnau Prat-Pérez, David Dominguez-Sal, Josep M Brunat, and Josep-Lluís Larriba-Pey. Shaping communities out of triangles. In *Proc. of the 21st ACM Intl. Conf. on Information and knowledge management*, pages 1677–1681, 2012. doi:10.1145/2396761.2398496. URL <https://doi.org/10.1145/2396761.2398496>.
- [150] Raphael Kimmig, Henning Meyerhenke, and Darren Strash. Shared memory parallel subgraph enumeration. In *IPDPSW*, pages 519–529. IEEE, 2017. doi:10.1109/IPDPSW.2017.133. URL <https://doi.org/10.1109/IPDPSW.2017.133>.
- [151] Norishige Chiba and Takao Nishizeki. Arboricity and subgraph listing algorithms. *SIAM J. Comp.*, 14(1):210–223, 1985. doi:10.1137/0214017. URL <https://doi.org/10.1137/0214017>.
- [152] Lars Gottesbüren, Tobias Heuer, Peter Sanders, and Sebastian Schlag. Scalable shared-memory hypergraph partitioning. In *2021 Proc. of the Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 16–30. SIAM, 2021. doi:10.1137/1.9781611976472.2. URL <https://doi.org/10.1137/1.9781611976472.2>.
- [153] Lars Gottesbüren, Tobias Heuer, Peter Sanders, Christian Schulz, and Daniel Seemaier. Deep multilevel graph partitioning. In *29th Annual European Symp. on Algorithms, ESA 2021, Sep. 6-8, 2021, Lisbon, Portugal*, volume 204 of *LIPICs*, pages 48:1–48:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.ESA.2021.48. URL <https://doi.org/10.4230/LIPICs.ESA.2021.48>.
- [154] Marcelo Fonseca Faraj, Alexander van der Grinten, Henning Meyerhenke, Jesper Larsson Träff, and Christian Schulz. High-quality hierarchical process mapping. In *18th Intl. Sym. on Experimental Algorithms, SEA*, volume 160 of *LIPICs*, pages 4:1–4:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

- doi:10.4230/LIPIcs.SEA.2020.4. URL <https://doi.org/10.4230/LIPIcs.SEA.2020.4>.
- [155] Felix Hausberger, Marcelo Fonseca Faraj, and Christian Schulz. A distributed multilevel memetic algorithm for signed graph clustering [short paper]. In *Genetic and Evolutionary Computation Conf. Companion, (GECCO)*, to appear. doi:10.1145/3583133.3590537.
- [156] Felix Hausberger, Marcelo Fonseca Faraj, and Christian Schulz. A distributed multilevel memetic algorithm for signed graph clustering. *CoRR*, abs/2208.13618, 2022. doi:10.48550/arXiv.2208.13618. URL <https://doi.org/10.48550/arXiv.2208.13618>.
- [157] Charles M. Fiduccia and Robert M. Mattheyses. A linear-time heuristic for improving network partitions. In *Proceedings of the 19th Design Automation Conference, DAC*, pages 175–181. ACM/IEEE, 1982. doi:10.1145/800263.809204. URL <https://doi.org/10.1145/800263.809204>.
- [158] Chris Walshaw and Mark Cross. Mesh partitioning: A multilevel balancing and refinement algorithm. *SIAM J. Sci. Comput.*, 22(1):63–80, 2000. doi:10.1137/S1064827598337373. URL <https://doi.org/10.1137/S1064827598337373>.
- [159] Chris Walshaw. Multilevel refinement for combinatorial optimisation problems. *Ann. Oper. Res.*, 131(1-4):325–372, 2004. doi:10.1023/B:ANOR.0000039525.80601.15. URL <https://doi.org/10.1023/B:ANOR.0000039525.80601.15>.
- [160] Peter Sanders and Christian Schulz. Distributed evolutionary graph partitioning. In *Proc. of the 14th Meeting on Algorithm Engineering & Experiments, ALENEX 2012, The Westin Miyako, Kyoto, Japan, January 16, 2012*, pages 16–29. SIAM / Omnipress, 2012. doi:10.1137/1.9781611972924.2. URL <https://doi.org/10.1137/1.9781611972924.2>.
- [161] Brad L. Miller and David E. Goldberg. Genetic algorithms, tournament selection, and the effects of noise. *Complex Syst.*, 9(3), 1995. URL http://www.complex-systems.com/abstracts/v09_i03_a02.html.
- [162] Benjamin Doerr and Mahmoud Fouz. Asymptotically optimal randomized rumor spreading. In *Automata, Languages and Programming - 38th Intl. Colloquium, ICALP*, volume 6756 of *LNCS*, pages 502–513. Springer, 2011. doi:10.1007/978-3-642-22012-8_40. URL https://doi.org/10.1007/978-3-642-22012-8_40.

Curriculum Vitae

Education

- 2011-2017 Bachelor, Computer Engineering
Federal Center of Technological Education of Minas Gerais, Brazil
Thesis: A Local Search Genetic Algorithm Approach to Optimize
Roadside Unit Placement in Vehicular Networks based on Gamma
Deployment Metric
- 2014-2015 Nondegree International Visiting Program
Colorado State University, United States
- 2017-2019 Master, Computer Science
Federal University of Minas Gerais, Brazil
Thesis: Gamma Deployment Problem in Grids: Complexity and a
new Integer Linear Programming Formulation

Work Experience

- 2013-2014 Intern, Technological Initiation
Federal Center of Technological Education of Minas Gerais, Brazil
- 2017 Intern, Software Engineer
LYNX Process, Brazil
- 2019-2020 Academic Employee
University of Vienna, Austria
- 2020- Academic Employee
Heidelberg University, Germany

Honors

- 2010 Silver Medal in Mathematics
Brazilian Olympics of Mathematics of Public High School, Brazil
- 2014 Scholarship at Colorado State University
CAPES Foundation, Brazil
- 2017 Honorable Mention for Best Final Bachelor GPA
Federal Center of Technological Education of Minas Gerais, Brazil

List of Publications

Journal Papers

- [1] Marcelo Fonseca Faraj, Sebastián Urrutia, and João F. M. Sarubbi. Gamma Deployment Problem in Grids: Hardness and New Integer Linear Programming Formulation. *International Transactions in Operational Research*, Volume 27, Article No 6, pages 2740–2759, 2020.
- [2] Marcelo Fonseca Faraj and Christian Schulz. Buffered Streaming Graph Partitioning. *ACM Journal of Experimental Algorithmics*, Volume 27, pages 1.10:1–1.10:26, 2022.
- [3] Ümit V. Çatalyürek, Karen D. Devine, Marcelo Fonseca Faraj, Lars Gottlieb, Tobias Heuer, Henning Meyerhenke, Peter Sanders, Sebastian Schlag, Christian Schulz, Daniel Seemaier, and Dorothea Wagner. More Recent Advances in (Hyper)Graph Partitioning. *ACM Computing Surveys*, Volume 55, Issue 12, Article No 253, pages 1–38, 2022.

Conference Full Papers

- [1] Marcelo Fonseca Faraj, João F. M. Sarubbi, Cristiano M. Silva, Marcelo F. Porto, and Nilson T. R. Nunes. Estudo de Caso: o Problema do Transporte Escolar Rural em Minas Gerais. In *Simpósio Brasileiro de Pesquisa Operacional (SBPO)*, pages 1889–1900, 2014.
- [2] Marcelo Fonseca Faraj, João F. M. Sarubbi, Cristiano M. Silva, Marcelo F. Porto, and Nilson T. R. Nunes. A Real Geographical Application for the School Bus Routing Problem. In *IEEE Conference on Intelligent Transportation Systems (ITSC)*, pages 2762–2767, 2014.
- [3] Marcelo Fonseca Faraj, Sebastián Urrutia, and João F. M. Sarubbi. Problema da Disposição Gamma: Prova de NP-Completeness e um Novo Modelo de Programação Linear Inteira. In *Simpósio Brasileiro de Pesquisa Operacional (SBPO)*, 2018.

- [4] Marcelo Fonseca Faraj, João F. M. Sarubbi, Cristiano M. Silva, and Flávio V. C. Martins. A Memetic Algorithm Approach to Deploy RSUs Based on the Gamma Deployment Metric. In *IEEE Congress on Evolutionary Computation (CEC)*, pages 1–8, 2018.
- [5] Marcelo Fonseca Faraj, Alexander van der Grinten, Henning Meyerhenke, Jesper L. Träff, and Christian Schulz. High-Quality Hierarchical Process Mapping. In *Proceedings of the 18th Symposium on Experimental Algorithms (SEA)*, volume 160 of LIPIcs, pages 4:1–4:15, 2020.
- [6] Marcelo Fonseca Faraj and Christian Schulz. Recursive Multi-Section on the Fly: Shared-Memory Streaming Algorithms for Hierarchical Graph Partitioning and Process Mapping. In *IEEE International Conference on Cluster Computing (CLUSTER)*, volume 9411, pages 473–483, 2022.
- [7] Adil Chhabra, Marcelo Fonseca Faraj, and Christian Schulz. Local Motif Clustering via (Hyper)Graph Partitioning. In *Proceedings of the Symposium on Algorithm Engineering and Experiments (ALENEX)*, pages 96–109. SIAM, 2023.
- [8] Kamal Eyubov, Marcelo Fonseca Faraj, and Christian Schulz. FREIGHT: Fast Streaming Hypergraph Partitioning. In *Proceedings of the 21st Symposium on Experimental Algorithms (SEA)*, to appear, 2023.

Conference Short Papers and Poster Papers

- [1] Marcelo Fonseca Faraj, João F. M. Sarubbi, Cristiano M. Silva, and Flávio V. C. Martins. A Hybrid Genetic Algorithm for Deploying RSUs in VANETs based on Inter-Contact Time (Extended Abstract). In *Genetic and Evolutionary Computation Conference Companion (GECCO '17 Companion)*, pages 193–194, 2017.
- [2] Marcelo Fonseca Faraj, Sebastián Urrutia, and João F. M. Sarubbi. O Problema da Deposição Gamma é NP-Completo (Extended Abstract). In *Encontro de Teoria da Computação (ETC)*, 2018.
- [3] Adil Chhabra, Marcelo Fonseca Faraj, and Christian Schulz. Local Motif Clustering via (Hyper)Graph Partitioning (Extended Abstract). In *Proceedings of the 15th International Symposium on Combinatorial Search (SoCS)*, pages 261–263, AAAI Press, 2022.
- [4] Felix Hausberger, Marcelo Fonseca Faraj, and Christian Schulz. A Distributed Multilevel Memetic Algorithm for Signed Graph Clustering (Short Paper). In *Genetic*

and Evolutionary Computation Conference Companion (GECCO '23 Companion), to appear, 2023.

Technical Reports

- [1] Marcelo Fonseca Faraj, Alexander van der Grinten, Henning Meyerhenke, Jesper L. Träff, and Christian Schulz. High-Quality Hierarchical Process Mapping. Technical Report, University of Vienna, Humboldt Universität zu Berlin, and Technical University of Vienna, 2020. (arXiv:2001.07134v2)
- [2] Marcelo Fonseca Faraj and Christian Schulz. Buffered Streaming Graph Partitioning. Technical Report, Heidelberg University, 2021. (arXiv:2102.09384)
- [3] Marcelo Fonseca Faraj and Christian Schulz. Recursive Multi-Section on the Fly: Shared-Memory Streaming Algorithms for Hierarchical Graph Partitioning and Process Mapping. Technical Report, Heidelberg University, 2021. (arXiv:2202.00394)
- [4] Ümit V. Çatalyürek, Karen D. Devine, Marcelo Fonseca Faraj, Lars Gottlieb, Tobias Heuer, Henning Meyerhenke, Peter Sanders, Sebastian Schlag, Christian Schulz, Daniel Seemaier, and Dorothea Wagner. More Recent Advances in (Hyper)Graph Partitioning. Technical Report, 2022. (arXiv:2205.13202)
- [5] Adil Chhabra, Marcelo Fonseca Faraj, and Christian Schulz. Local Motif Clustering via (Hyper)Graph Partitioning. Technical Report, Heidelberg University, 2022. (arXiv:2205.06176)
- [6] Felix Hausberger, Marcelo Fonseca Faraj, and Christian Schulz. A Distributed Multi-level Memetic Algorithm for Signed Graph Clustering. Technical Report, Heidelberg University, 2022. (arXiv:2208.13618)
- [7] Kamal Eyubov, Marcelo Fonseca Faraj, and Christian Schulz. FREIGHT: Fast Streaming Hypergraph Partitioning. Technical Report, Heidelberg University, 2023. (arXiv:2302.06259)
- [8] Adil Chhabra, Marcelo Fonseca Faraj, and Christian Schulz. Faster Local Motif Clustering via Maximum Flows. Technical Report, Heidelberg University, 2023. (arXiv:2301.07145)