

# Distributing Sparse Matrix/Graph Applications in Heterogeneous Clusters – an Experimental Study

Charilaos Tzovas  
Humboldt-Universität zu Berlin  
Department of Computer Science  
Berlin, Germany  
tzovas.charilaos@hu-berlin.de

Maria Predari  
Humboldt-Universität zu Berlin  
Department of Computer Science  
Berlin, Germany  
predarim@hu-berlin.de

Henning Meyerhenke  
Humboldt-Universität zu Berlin  
Department of Computer Science  
Berlin, Germany  
meyerhenke@hu-berlin.de

**Abstract**—Many problems in scientific and engineering applications contain sparse matrices or graphs as main input objects, e. g., numerical simulations on meshes. Large inputs are abundant these days and require parallel processing for memory size and speed. To optimize the execution of such simulations on cluster systems, the input problem needs to be distributed suitably onto the processing units (PUs). More and more frequently, such clusters contain different CPUs or a combination of CPUs and GPUs. This heterogeneity makes the load distribution problem quite challenging. Our study is motivated by the observation that established partitioning tools do not handle such heterogeneous distribution problems as well as homogeneous ones.

In this paper, we first formulate the problem of balanced load distribution for heterogeneous architectures as a multi-objective, single-constraint optimization problem. We then split the problem into two phases and propose a greedy approach to determine optimal block sizes for each PU. These block sizes are then fed into numerous existing graph partitioners, for us to examine how well they handle the above problem. One of the tools we consider is an extension of our own previous work (von Looz et al., ICPP’18) called **Geographer**. Our experiments on well-known benchmark meshes indicate that only two tools under consideration are able to yield good quality. These two are **ParMetis** (both the geometric and the combinatorial variant) and **Geographer**. While **ParMetis** is faster, **Geographer** yields better quality on average.

**Index Terms**—load balancing, graph partitioning, heterogeneous systems

## I. INTRODUCTION

Applications with sparse matrices or sparse graphs are ubiquitous in science and engineering. In case of sparse matrices, applications often model complex problems as discretizations of partial differential equations, e. g., in molecular dynamics [26] or climate simulations [11]. This usually leads to large matrix-vector problems such as sparse linear systems or eigenproblems with significant demands in terms of memory and computation. Hence, such simulations are often executed on parallel systems with distributed memory. Typical algorithms include Krylov subspace solvers such as conjugate gradient (CG), which are iterative and perform sparse matrix-vector product (SpMV) operations in each iteration. Due to the correspondence between matrices and graphs [23], similar computational demands arise in large-

scale graph computations such as the analysis of massive online social networks [9].<sup>1</sup>

To obtain short execution times, it is essential to find a good distribution of the application’s computational tasks onto the available processing units (PUs). This is all the more important for heterogeneous systems, i. e., when the PUs differ in terms of their speed and memory capacity. Such systems become more and more common since GPUs offer a more energy-efficient way to solve certain problems in parallel [34].

A good distribution balances the load (in particular the matrix/graph and the main vectors involved in the computations) between the PUs and leads to a low communication overhead. Since for many of the problems under consideration, GPUs are faster than CPUs, they could in principle receive a larger share of the distributed data. At the same time, however, GPUs usually have smaller memories – a constraint that must be obeyed to guarantee a healthy application.

One popular way of finding a good distribution is to use graphs and/or hypergraphs to model the task interactions and then to employ (hyper)graph/matrix partitioning [6]. Finding an *optimal* distribution this way is  $\mathcal{NP}$ -hard [14]. Thus, in practice heuristics are applied [5], [17], [24]. (Hyper)Graph/matrix partitioning usually leads to a partition of the graph/matrix into so-called *blocks* of vertices/rows, where each block is handled by one PU. Since classic graph partitioning does not take into account how fast the PUs can communicate with each other, explicit algorithms for mapping the blocks onto PUs are sometimes employed in addition, see [2]. This mapping step becomes more important for more complex/irregular scenarios [15] (e. g., in terms of application matrices, heterogeneous and/or hierarchical compute systems).

**Motivation:** From an abstract perspective, the goal is to (re)distribute the input matrix among the available PUs in a way such that the application’s execution time is minimized. This means that (i) no PU receives more data than it can fit into its memory, (ii) each PU receives a part of the input that matches its proportionate computational capabilities, and (iii) the communication costs induced by the distribution are low. Partitioning tools that offer mapping capabilities are

<sup>1</sup>In the remainder, we focus in the description of the underlying application on matrix computations, but graph computations are equally relevant.

addressing requirement (iii) and often result in distributions of low communication costs. However, most of these tools do not provide automatic support for heterogeneous PU characteristics, thus, they do not address requirements (i) and (ii). Indeed, modern systems with CPUs and GPUs working on the same problem together lead to conflicting demands in terms of processing capabilities and available memory. The lack of explicit support by existing partitioning tools for the above problem is what motivates our present study.

*Outline and Contribution:* In Sec. II we formulate the problem of balanced load distribution for heterogeneous PUs as multi-objective optimization problem with a memory constraint. To be able to use existing partitioners (which are reviewed in Sec. III), we divide the problem into two stages: first, we compute the desired block sizes with a greedy method in  $\mathcal{O}(k \log k)$  time for  $k$  PUs, see Sec. IV. There, we also prove these block sizes to be optimal for this first stage (but not for the complete optimization problem, of course). The block sizes are then fed into a wide variety of popular partitioning tools to address the second objective in a second stage. After excluding three tools for not supporting our problem adequately, our experimental study on the remaining three tools with eight algorithms (Sec. VI) shows: only ParMetis and Geographer handle the heterogeneous load distribution with good solution quality. A combination of balanced  $k$ -means and combinatorial refinement, introduced here in Section V, calculates partitions for meshes with 10% better cut value compared to ParMetis. In most cases, we “pay” only with a running time increase of about 50%, and there are several cases where Geographer is both faster *and* has better quality. **Material omitted from the main part can be found in the appendix.**

## II. PROBLEM STATEMENT

As is customary, we exploit the correspondence between a symmetric  $n \times n$  matrix  $A$  and an undirected graph  $G$  with  $n$  vertices [23].  $G$  has an edge  $\{u, v\}$  iff the entry  $A[u, v]$  is different from zero. This allows to see the load distribution problem for matrices through graph lens. A list of important symbols and acronyms that are used in the remainder are given in Table I.

Symbol	Description
GP	graph partitioning
LDHT	load distribution problem for heterogeneous topologies
$k$	number of blocks of a partition
$\varepsilon$	partition imbalance threshold
$cut(\cdot)$	number of edges cut in a partition
$b_i$	block $i$ of a partition
$p_i$	processing unit $i$ of the system (PU)
$c_s(\cdot)$	speed of a PU (number of operations per time unit)
$m_{cap}(\cdot)$	memory capacity of a PU
$C_s$	total computational speed of the system
$M_{cap}$	total memory of the system
$tw(\cdot)$	the target weight of a block

TABLE I: Description of symbols used in the paper.

### A. Graph Partitioning (GP)

We assume an application where each matrix element / graph vertex represents the same computational demand and memory requirement (thus both are normalized to 1). As a first step, let us also assume a homogeneous compute system in which all PUs have the same speed. Moreover, let them communicate among each other with the same speed, regardless of the pair of PUs under consideration.

To make the application as fast as possible, we want to balance the load among PUs and minimize the (expensive) communication between PUs. In such a setting, the balanced load distribution problem is usually modeled by *graph partitioning* (GP). For GP, given an application graph  $G_a = (V, E)$ , an integer  $k$  and an imbalance threshold  $\varepsilon$ , we seek a partition  $\Pi$  into  $k$  blocks  $b_0, b_1, \dots, b_{k-1}$  such that a cost function  $f(\Pi)$  is minimized, while for all blocks  $|b_i| \leq (1 + \varepsilon)n/k$  holds. In case of unweighted vertices or unit weights,  $|b_i|$  denotes the number of nodes in block  $b_i$ .

The most common cost function is the edge cut,  $cut(\Pi)$ , i.e., the number (or weight) of edges with endpoints in different blocks, see [6]. Some authors also considered the (more accurate) communication volume and/or the number of boundary vertices, i.e., the number of vertices with at least one neighbor in a different block [18].

### B. Balanced Load Distribution for Heterogeneous PUs

Next, let us model heterogeneous PUs. They can have different processing speeds and different memory capacities. In this scenario, each PU should receive an amount of load that is proportionate to its speed – but never more than its memory capacity. Formally, the input consists of an application graph  $G_a = (V, E)$ , the number  $k$  of blocks, and a representation of the compute system topology. Since most compute systems are hierarchical in one way or another, we make the common assumption [31] that this representation takes the form of a tree  $T$ ; its leaves represent the  $k$  PUs, the inner nodes correspond to sets of PUs (i.e., to their descendants in the tree).

Each PU  $p_i$  has two weights:  $c_s(p_i)$  is its speed (a normalized number of operations per time unit) and  $m_{cap}(p_i)$  its memory capacity. The respective values of an inner node  $v$  of  $T$  are formed by recursively accumulating the values of all children of  $v$ . Thus,  $C_s = \sum_{i=0}^{k-1} c_s(p_i)$  and  $M_{cap} = \sum_{i=0}^{k-1} m_{cap}(p_i)$  are the total computational speed and memory of the system, respectively.

To be able to assess established tools, we first model the balanced load distribution problem by computing a partition  $\Pi$  of  $V$  into  $k$  blocks  $b_0, b_1, \dots, b_{k-1}$  again and then mapping each block  $b_i$  to PU  $p_i$ . This means that block  $b_i$  must take the speed and the memory capacity of  $p_i$  into consideration. For the load distribution problem for heterogeneous topologies (LDHT), the goal is then to find a partition  $\Pi$  such that:

$$\text{minimize } cut(\Pi) \text{ and} \quad (1)$$

$$\text{minimize } \max \frac{tw(b_i)}{c_s(p_i)} \quad (2)$$

$$\text{s. t. } tw(b_i) \leq m_{cap}(p_i), \quad (3)$$

where  $tw(b_i)$  is the target weight for block  $i$ . In this formulation, Eq. (1) is meant to minimize the communication between PUs, Eq. (2) is meant to balance the computational load according to the speeds, while the constraint (3) enforces that no memory overload occurs. The LDHT problem is  $\mathcal{NP}$ -hard since it contains the  $\mathcal{NP}$ -hard graph partitioning problem as a special case.

### III. RELATED WORK

In this section, we focus on aspects closely related to the main motivation of our study. First, we discuss algorithms for the partitioning problem, then we consider the related multi-constraint, multi-objective problem. Finally, we mention existing literature on load balancing for heterogeneous systems.

*a) GP algorithms:* There is a large amount of literature on heuristics for the GP problem. Most of them can be classified as either combinatorial, where the connectivity information of the (hyper)graph is used to drive the partitioning, or geometric, where solely the coordinate information is used. Combinatorial partitioning tools (or simply: (hyper)graph partitioning tools) usually produce solutions, with better quality while geometric ones typically have lower running times. Arguably, one of the most successful heuristics for accelerating any partitioning algorithm is the multilevel approach [17]. The multilevel approach consists of three phases: coarsening, initial partitioning, and uncoarsening. In the coarsening phase, vertices are successively merged together to form a series of smaller graphs until a sufficiently small graph is produced. Then, an initial partition is obtained on that graph and is projected back to the original input, through uncoarsening. In the corresponding uncoarsening level, the merged vertices are split and the partition of the coarser graph is further refined. The multilevel approach is available in a number of well-known graph partitioning tools such as Metis, Scotch and KaHip. The most commonly used methods for graph coarsening are usually based on vertex-matching schemes [22] (as in Metis) and on clusterings from on label propagation (as used by xtraPulp and by some configurations of KaHip). During uncoarsening many tools use local search refinement techniques such FM/KL methods [12], [22].

Geometric partitioning tools include implementations of space-filling curves [33] (SFC), recursive inertial bisection [28] (RIB) and recursive coordinate bisection [16] (RCB). RCB is a recursive bisection scheme that attempts to minimize the boundary between the subdomains by splitting the mesh in half, orthogonal to its longest dimension. RIB uses the principal inertial axis to produce a bisection and does not restrict the splitting to one dimension. Furthermore, a more recent geometric partitioning, called MultiJagged [10], computes partitions by recursively multi-sectioning the dataset, which can also be viewed as a generalization of RCB. Implementations of all the above methods can be found in the Zoltan public library [3]. We also use Geographer, a scalable algorithm for geometric partitioning, recently proposed by a subset of the authors [32]. Geographer applies a balanced version of the  $k$ -means algorithm to obtain blocks with compact shapes.

Finally, parallel partitioning algorithms are crucial for modern distributed systems, especially since graph sizes are increasing dramatically. In general, geometric partitioning tools tend to be easier to parallelize than multilevel graph partitioners. However, parallel versions of all the algorithms/tools presented above are available. Parallel algorithms exist in ParMetis, ParHip, xtraPulp, PT-Scotch for graph partitioning and the Zoltan library for the geometric ones. The classic partitioning algorithms do not explicitly support the LDHT problem. However, to solve the LDHT problem, we need a classic partitioning tool that accepts specific weights per block. In the experiments, we use the above tools (parallel versions), excluding only those who do not allow specific weights as input. Some of these tools are equipped with additional capabilities, such as solutions for the multi-constraint, multi-objective problem as ParMetis and xtraPulp.

*b) Multi-constraint, multi-objective partitioning:* Multi-constraint, multi-objective graph partitioning algorithms are used to model problems with several balance constraints and/or several optimization objectives. Example problems arise in multi-physics or multi-mesh simulations [29]. Then, one assigns a weight vector of size  $m$  to each vertex and a weight vector of size  $l$  to each edge. The problem then becomes that of finding a partition that minimizes the edge cut with respect to all  $l$  weights, subject to the constraints that each of the  $m$  weights is balanced across the subdomains. ParMetis was one of the first tools to include algorithms for the multi-constraint, multi-objective paradigm. In Ref. [21] the problem of balancing computation and memory constraints is modeled as a two-constraint problem by associating a vector of size two for each vertex, where the elements of the vector represent the computation and memory requirements associated with that vertex. Later in Ref. [27] the multi-constraint multi-objective algorithm of ParMetis was used to devise an architecture-aware partitioning algorithm. Although solutions for the multi-constraint, multi-objective problem are relevant, the existing tools do not explicitly address the LDHT problem; instead, they treat *all* weights as constraints (upper bounds). In LDHT, the goal is to balance one weight without exceeding the second one. We address this problem by describing a two-stage approach. First we compute the optimal block weights to be given to these tools, and thus our partial problem leads to single-constraint, single-objective graph partitioning.

*c) Load balancing for heterogeneous systems:* Related work on load balancing for heterogeneous clusters focuses mainly on solving the mapping problem (or process/task placement), for which the communication costs between PUs are added in the classic model of graph partitioning. The majority of those algorithms do not consider heterogeneous clusters in terms of different speed or memory capabilities (rather only in terms of communication costs) and thus do not solve the LDHT problem. More details can be found in a survey [19] on algorithms and software for mapping. A recent work that focuses on partitioning meshes for heterogeneous systems appeared in [8].

#### IV. DETERMINING OPTIMAL BLOCK SIZES

Let us assume we are faced with a load distribution problem and use one of the established tools that accept (nearly) arbitrary target block sizes/weights  $tw(b_0), \dots, tw(b_{k-1})$  for the blocks of the application graph's vertex set. Here, we discuss how to calculate these  $tw(\cdot)$  values for the LDHT problem, which is more difficult than standard GP. For GP, one typically passes only  $G_a$ ,  $k$ , and  $\varepsilon$  to the respective tool, which then computes a partition with block sizes  $\leq (1+\varepsilon)\lceil \frac{n}{k} \rceil$ .

For LDHT, the calculation of the  $tw(\cdot)$  values is not supported inside the respective partitioning tools, though, and needs to happen beforehand. We thus strive for a two-phase process in which we compute the optimal block sizes first. These sizes are then passed on to the partitioner. Of course, a unified process with just one phase would be preferable (regarding quality), because a two-phase process neglects the objective function edge cut in the first phase completely. But it would certainly overburden the users of the load distribution tools to make the necessary deep changes inside the tools (which would need to be on the algorithmic level). Thus, the cut is optimized by the tools in a second black-box phase and we focus now on the calculation of the  $tw(\cdot)$  values for the LDHT Problem (unweighted  $G_a$ , heterogenous topology).

We denote the optimal target block size without any constraints for a block  $b_i$  by  $tw(b_i)^*$ . In the trivial case where all PUs have sufficient memory, this optimal solution (to a much simpler problem) is for every block to have a weight proportional to the corresponding PU's speed:

$$tw(b_i)^* = n \cdot \frac{c_s(p_i)}{C_s}. \quad (4)$$

Of course, in reality each PU has a certain memory capacity that must not be exceeded. If the ‘‘previously optimal size’’ (without constraint) does not fit into the memory of a PU, then one should load it as much as possible to obtain the ‘‘new optimum’’ (with memory constraint). Note that a PU with insufficient memory (in comparison to its speed) has an implication on the load of other PUs as well, since the remaining total load differs from what we assumed in Eq. (4).

For the calculation of the  $tw(\cdot)$ , we propose a greedy method, see Algorithm 1. We sort the PUs in decreasing order by the criterion  $c_s(p_i)/m_{cap}(p_i)$  (Line 1). This way, we want to ensure that the fastest PUs get as much load as they can handle. Lines 2 and 3 initialize the joint computational load posed by  $G_a$  and the normalized joint speed of the PUs encoded in the topology tree  $T$ , respectively. The main loop iterates over the PUs in sorted order. In Line 5, in analogy with Eq. (4), we compute the desirable target weight,  $desW(b_i)$ , of block  $b_i$  for PU  $p_i$ . If  $desW(b_i)$  is larger than the memory capacity of PU  $p_i$ , then  $p_i$  is assigned this memory capacity and is called *saturated* (Line 8). Otherwise,  $p_i$ 's memory capacity exceeds its desirable target load, so that  $p_i$  is assigned exactly this target load, and is called *non-saturated* (Line 10). Finally, in Lines 11 and 12, we subtract the load assigned in this iteration from  $jLoad$  and the current PU's speed from  $jSpeed$ , respectively.

---

**Algorithm 1:** Calculate target block sizes for the LDHT Problem.

---

**Data:** Applic. graph  $G_a$ ,  $k$ , topology info as tree  $T$   
**Result:** Target block sizes  $tw(b_0), \dots, tw(b_{k-1})$

- 1 Sort PUs such that  $c_s(p_0)/m_{cap}(p_0) \geq c_s(p_1)/m_{cap}(p_1) \geq \dots \geq c_s(p_{k-1})/m_{cap}(p_{k-1})$
- 2  $jLoad \leftarrow |V|$   $\triangleright$  computational load of the graph
- 3  $jSpeed \leftarrow C_s$   $\triangleright$  normalized speed of all PUs from  $T$
- 4 **for**  $i \leftarrow 0, 1, \dots, k-1$  **do**
- 5      $desW(b_i) \leftarrow c_s(p_i) \cdot jLoad / jSpeed$   
         $\triangleright$  desirable weight proportionate to speed
- 6     **if**  $desW(b_i) > m_{cap}(p_i)$  **then**
- 7          $tw(b_i) \leftarrow m_{cap}(p_i)$   $\triangleright$  load = mem capacity
- 8     **else**
- 9          $tw(b_i) \leftarrow desW(b_i)$   $\triangleright$  more mem than load
- 10      $jLoad \leftarrow jLoad - tw(b_i)$
- 11      $jSpeed \leftarrow jSpeed - c_s(p_i)$
- 12 **end for**
- 13 **return**  $tw(b_0), tw(b_1), \dots, tw(b_{k-1})$

---

Proceeding greedily in this sorted order ensures that we first fill the ‘‘fast enough’’ PUs and they receive as much as they can handle. Moreover, we prove in the following that one actually obtains the optimal solution of Eq. (2) under the memory constraints. The algorithm's running time of  $O(k \log k)$  is no contradiction to the  $\mathcal{NP}$ -hardness of LDHT – due to the separation into two disjoint phases, we do not solve the LDHT problem in an optimal way.

It is absolutely natural to assume that non-saturated PUs exist – otherwise no valid solution exists. For the optimality proof, we first show that all saturated PUs appear consecutively at the top of the sorted sequence produced in Line 1 (the proof can be found in Appendix A).

**Lemma 1.** *No saturated PU appears after a non-saturated PU in the sorted sequence produced in Line 1.*

**Theorem 1.** *Algorithm 1 computes the optimal solution for the objective function (2) under the memory constraint (3). It runs in  $O(k \log k)$  time.*

In the inductive optimality proof (Appendix B) we exploit the fact that a change from a saturated PU to a non-saturated one in the course of Algorithm 1 happens only once – and one never changes back. This is a consequence from Lemma 1.

#### V. EXTENSIONS TO GEOGRAPHER

In this section we describe a new version of Geographer, called Geographer-R, in which we combine the geometric partitioning with combinatorial techniques. More precisely, we employ a multilevel approach and use a parallel variant of the FM algorithm for local refinement. Contrary to the classic multilevel approach, we obtain a partition before even starting to coarsen the graph. This is done to obtain a good initial data distribution to PUs and we use Geographer to do so. Each PU receives one block for the remainder of the procedure, and is

responsible for coarsening its local subgraph using heavy-edge matching techniques.

For the uncoarsening, we create the communication graph  $G_c$  (a.k.a. quotient graph: each vertex of  $G_c$  corresponds to a block of  $G_a$ , weighted edges in  $G_c$  model the communication volume exchanged between the blocks of  $G_a$ ) first. Inspired by Ref. [20], we compute a maximum edge coloring algorithm to determine the communication rounds and the communicating PU pairs in each round. As a reminder, the classic FM algorithm sorts vertices by gain and moves the vertex with the highest gain if it does not violate the balance constraint. In our case, we do not consider all vertices but only a smaller number of those in each block, located in the extended neighborhood of the boundary nodes. Those vertices are identified by performing a number of BFS rounds starting from the boundary nodes. Then, after exchanging candidate nodes among communicating PUs, in each round we apply pairwise FM refinements between the corresponding PU pairs in parallel. Each PU in the pair performs the local refinement independently and the best of the two solutions is kept. Uncoarsening and local refinement are repeated until we obtain a solution for the original graph.

We also extended the balanced  $k$ -means algorithm [32] from Geographer to better address hierarchical compute systems. The hierarchy is provided as an implicit tree by a list of numbers  $k_1, k_2, \dots, k_h$  for  $h$  hierarchy levels – where  $k_i$  denotes the tree fan-out on that level. Thus, the total number of blocks is  $k = \prod k_i$ . To account for this hierarchy, we partition on level  $i$  each block into  $k_{i+1}$  blocks, respectively. Proceeding in this hierarchical fashion usually leads to solutions with slightly larger edge cut – direct  $k$ -way approaches often yield better quality than their recursive counterparts in this respect. The benefit is appreciated, however, from a mapping quality perspective: blocks that share a border (and thus communicate across this border) will likely be mapped to nearby PUs. As a fast post-processing step, we do a global repartitioning step that “smooths” the border and improves the cut.

Our experiments for the scenario that we emphasize on in this paper indicate that the quality in terms of edge cut is very close, usually within  $\pm 1\%$ , see Fig. 1. Although the communication volume can be reduced more in numerous cases, we decided to provide only the results for the original algorithm in detail to avoid cluttering the plots.

## VI. EXPERIMENTAL EVALUATION

In this section, we describe our experiments to evaluate the effectiveness of different partitioners under the load distribution model of Sections II and IV, for various (simulated) heterogeneous systems. Regarding the requirements of the application, we only treat the LDHT scenario, where the computational and memory requirements of the application are equal, corresponding to two unit weights for all vertices of the application graph. For the purpose of our experiments, we simulate modern systems with different levels of heterogeneity. To this end, we consider three different topologies; TOPO1

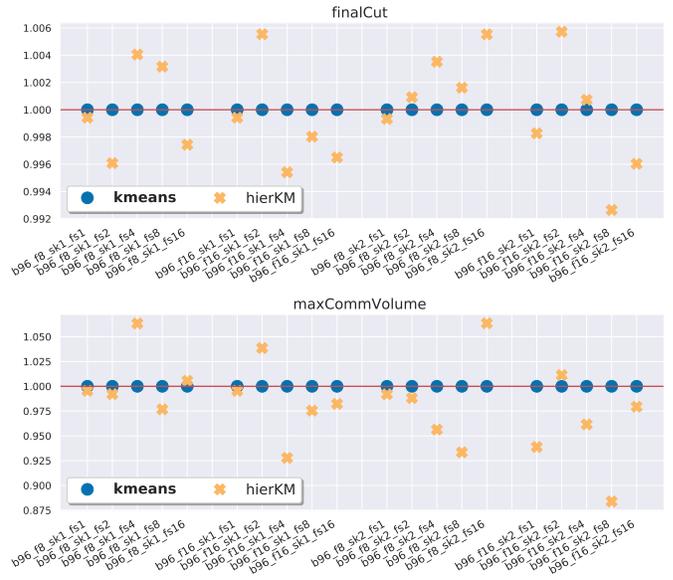


Fig. 1: Comparison regarding relative quality (top: edge cut, bottom: max. communication volume) between balanced  $k$ -means and the hierarchical version (lower is better).

(see VI-A), TOPO2 (see VI-B) and TOPO3 (see VI-C), which are further explained in the corresponding subsections.

All relevant material, tool versions and guidelines on how to reproduce the experimental results are publicly available in [https://github.com/hu-macsy/geographer/tree/Dev/publication\\_results](https://github.com/hu-macsy/geographer/tree/Dev/publication_results).

*a) Metrics:* Combinatorial quality metrics considered here are the edge cut and the maximum communication volume. Regarding actual application performance, we conduct simulations on a cluster and report running time results for common HPC kernels such as SpMV (sparse matrix-vector produce) and linear system solves with CG (conjugate gradient). We apply the CG solver from the numerical library LAMA [4] to systems derived from the graph’s Laplacian matrix. To ensure that the linear systems have a solution, we shift the diagonal of the Laplacian slightly to make the matrix positive definite. For SpMV and for the CG solver, the Laplacian of the input matrix is distributed according to the partition provided by the respective partitioning tool.

*b) Tool Selection:* For all our experiments we consider a variety of distributed tools for the partitioning process. More precisely, in the competitors set we include two versions of ParMetis, termed ParMetisGeom and ParMetisGraph. Their difference is that ParMetisGeom uses an SFC to obtain the initial partition of the graph. From the tools in the zoltan2 package [3], we consider the geometric partitioning methods RCB, RIB, and SFC. The graph partitioner xtraPulp is not included in our set because it targets complex networks and preliminary tests showed insufficient quality (high cut values and unbalanced parts) for our data sets. The current implementation of MultiJagged and ParHip do not accept sufficiently imbalanced block weights, so that they are excluded as well.

name	nodes $n$	edges
NACA0015	1,039,183	3,114,818
M6	3,501,776	10,501,936
333SP	3,712,815	11,108,633
AS365	3,799,275	11,368,076
NLR	4,163,763	12,487,976
hugetic-00020	7,122,792	10,680,777
hugetrace-00020	16,002,413	23,998,813
hugebubbles-00020	21,198,119	31,790,179
alyaTestCaseA	9,938,375	39,338,978
alyaTestCaseB	30,959,144	122,951,408
refinetrace	578,551,252	867,786,528
rdg_2d_2 <sup>x</sup>	2 <sup>x</sup> , $x = 25, \dots, 29$	$\approx 3n$
rgg_2d_2 <sup>x</sup>	2 <sup>x</sup> , $x = 25, \dots, 29$	$\approx 3n$
rgg_3d_2 <sup>x</sup>	2 <sup>x</sup> , $x = 25, \dots, 29$	$\approx 3n$

TABLE II: The graphs used in the experiments with the number of nodes and edges.

Regarding Geographer, geoKM refers to the balanced  $k$ -means method, geoRef when we additionally perform local refinement as a postprocessing step. The version geoPM-Ref consists of the balanced  $k$ -means method plus the local refinement routine from ParMetis.

*c) Instances:* To test the partitioners, we use geometric graphs that stem from or resemble scientific simulations. The graphs are displayed in Table II. The larger graphs were generated using the generator for random geometric graphs (rgg) and Delaunay triangulations (rdg) from KaGen [13]. We also include two graphs from the PRACE benchmark suite [30] that represent the respiratory system. Moreover, we use the generator from Ref. [25] to create a large adaptive triangular 2D mesh (from the *refinetrace* series) with 578M vertices. The rest of the graphs come from the 10th DIMACS challenge [1]. In our detailed presentation of the results, we use aggregate values for representative instances. Due to the different nature of the individual instances, an aggregation over all instances would not lead to meaningful results.

*d) Test Systems:* We use two test systems for our experiments. The small one is a local cluster with 16 compute nodes, each with 4 6-core Intel Xeon X7460 CPUs and 128 GB of RAM. This system was used for the smaller inputs from DIMACS and PRACE. The main volume of the experiments were carried out on the HLRN IV system Lise (<https://www.hlrn.de/>). Each compute node of Lise has two Intel Cascade Lake Platinum 9242 CPUs with 384 GB RAM and 96 cores. Using two different systems here does not affect the evaluation, since the running time comparison depends on *relative* values. For exact values on certain graph and topologies see Table IV. For the experiments in TOPO3, we used the local cluster and tuned down the CPU speed of certain compute nodes to simulate heterogeneity – as described in Section VI-C. In our experiments we assign one MPI process per PU. The factors between different PU speeds reflect recent results [7] on CPU-GPU comparisons.

#### A. Experiments for TOPO1

In the first category, we have 2 sets of PUs, set  $S$  (for slow) and set  $F$  (for fast). This category is further divided into two parts. Using  $k$  as the total number of PUs, we set  $|F| := k/12$

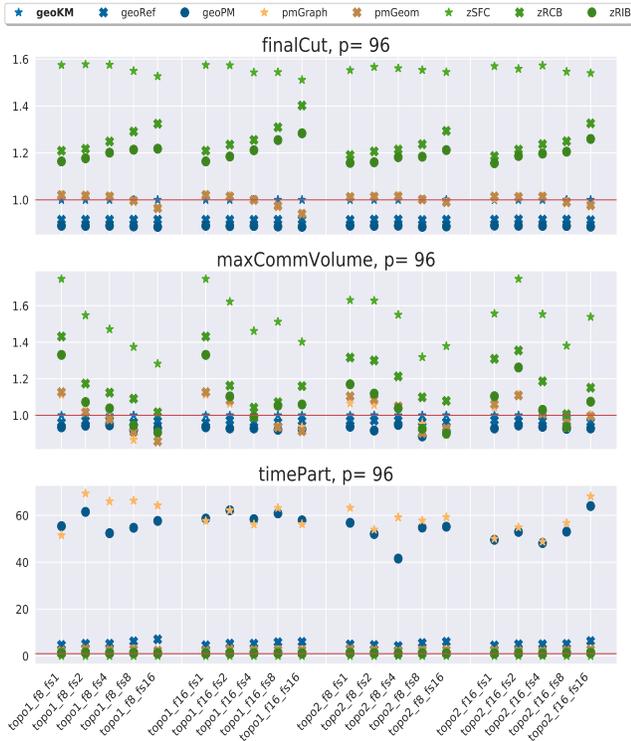
exp	speed	memory	$tw(\text{fast})/tw(\text{slow})$
1	1	2	1 - 1
2	2	3.2	2 - 2
3	4	5.2	3.2 - 3.5
4	8	8.5	5.5 - 6.1
5	16	13.8	9.4 - 11.5

TABLE III: Speed and memory of the fast PUs used for TOPO1 and TOPO2. The slow PUs have computational speed 1 and memory 2 for all the experiments. The last column show approximately how bigger is the target block weight for the fast PUs for the experiments where  $|F| = p/12$  and  $p/6$ .

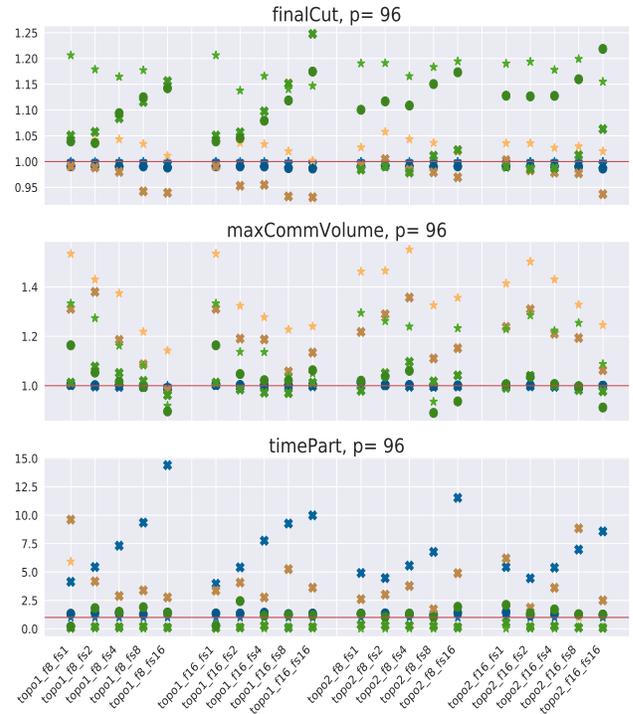
in the first part and  $|F| := k/6$  in the second one. The specs of the PUs in  $S$  remain always constant and we increase only the memory capacity and computational speed of the PUs in  $|F|$  as follows. We start with the same computational speed and memory as the slow PUs; this corresponds to a homogeneous system since all target block weights are equal. Then, we increase the computational speed of the fast PUs by a factor of 2 and the memory by a factor of 1.6 for 4 steps. The speed and memory of the PUs in  $F$ , for each experiment, are shown in Table III. In the first heterogeneous step, the PUs in  $F$  have enough memory to get the desirable block size, but in the remaining three steps, their memory is saturated. In total, this gives 10 experiments per graph. These experiments are performed for  $k = 24 \cdot 2^i$ , where  $i \in \{2, 3\}$ .

Fig. 2(a) shows results aggregated over the three hugeX graphs. The  $x$ -axis shows the different variations of TOPO1, the  $y$ -axis relative quality or partitioning time. A first observation is that heterogeneity plays a role, but not a big one. In particular the geometric methods from Zoltan are negatively affected in terms of quality with increasing heterogeneity (from left to right). The geometric-only approaches provide solutions with comparably high cut and max communication volume values, especially for two dimensions. Among these, the balanced  $k$ -means approach has the best quality in most cases with more than 15% improvement. Quality-wise, in most cases, ParMetis offers solution quality very close to the quality of balanced  $k$ -means. But the best solutions come from the two version geoRef and geoPMRef, which combine balanced  $k$ -means with local refinement.

For the 3D *alya* graphs in Fig. 2(b), the geometric methods narrow the gap somewhat. These application graphs are presented to show results against the general trend. In a few cases, the geometric methods even give here the solution with the lowest cut or volume, e. g., for *topol\_f8\_fs16*. ParMetis solutions have good cut values – on the other hand, the max communication volume is worse, even compared to geometric methods. As for Geographer, geoKM is better than the other geometric methods, but the combination with refinement offers some improvement. As expected, geometric methods have the lowest running time, less than a few seconds. The refinement of geoRef is faster than the refinement of ParMetis for the 2D graphs. For the 3D *alya* graphs, our refinement algorithm is affected by heterogeneity; the more heterogeneous the system, the slower the algorithm becomes. This is probably because,



(a) hugeX graphs representing numerical simulation meshes (2D)



(b) Two *alya* graphs (3D)

Fig. 2: Results for (a) three hugeX graphs from the 10th DIMACS Challenge and (b) two *alya* graphs using 96 PUs and 16 different topologies. The data points are the geometric mean of the values for each graph and are relative to the respective balanced  $k$ -means value. For the names in the  $x$ -axis: topoX indicates TOPO1 or TOPO2, ‘f’ indicates the number of fast PUs and ‘fs’ indicates the speed of the fast PUs. The slow PUs have speed 1 in all cases. These experiments were conducted on the local cluster. Values are relative to balanced  $k$ -means (lower is better).

as speed and memory of the fast PUs increase, the size of the blocks that correspond to the fast PUs increases, too. As we consider a percentage of vertices from each block for the local refinement, more vertices are eligible to move.

### B. Experiments for TOPO2

The second category consists of three sets of PUs,  $F$ ,  $S_1$  and  $S_2$ . It is designed to model systems with two kinds of CPUs and one kind of GPU. Again,  $|F| = k/6$  or  $|F| = k/12$ , but now, the slow PUs are divided in two equally sized groups  $|S_1| = |S_2|$ . The PUs in  $S_1$  have constant memory and speed, but the speed of the PUs in  $S_1$  increases as shown in Eq. (5):

$$\frac{c_s(s_1)}{m_{cap}(s_1)} = \frac{1}{2} \frac{c_s(f)}{m_{cap}(f)}, \quad s_1 \in S_1, f \in F \quad (5)$$

This ensures that PUs in  $S_1$  will be assigned their block weights after the PUs in  $F$  but before the PUs in  $S_2$ . The specs of the PUs in  $F$  are increased as before, see Table III. These experiments are performed for  $k = 24 \cdot 2^i$ ,  $i = 1, 2, \dots, 8$ .

Experiments with high numbers of PUs were conducted on Lise; their results are shown in Fig. 3 for the *refinedtrace* graph, which shows representative results: the algorithms geoRef and geoPMRef yield partitions with the lowest cut and communication volume. As the number of PUs increases, we

observe that ParMetisGraph and geoPMRef suffer from high running times, in contrast to our refinement implementation. Geometric approaches are steadily worse in quality but very fast in running time. Heterogeneity has again a negative effect on the edge cut produced by the geometric tools from Zoltan, whereas ParMetis seems to benefit slightly. An effect on the different algorithms in Geographer is hardly visible.

In Fig. 4, we see experiments with 3D random geometric and Delaunay triangulation graphs. Regarding quality, the same pattern as before emerges: geoKM is better than other geometric methods and the combinatorial algorithms all give better solutions (of similar quality). Again, geoRef seems to suffer from high running times when the topologies become more heterogeneous.

### C. Experiments for TOPO3

The third category consists of heterogeneous topologies simulated in the local compute cluster. We cannot change the specs of an individual core, but we change all 24 cores of a node. The experiments involve 4 and 8 compute nodes (96 and 182 PUs) where, each time, 1 or 2 nodes (24 or 48 PUs) are unchanged and the rest have their speed and memory lowered in order to represent the slow PUs. For these experiments we are able to get meaningful running times for SpMV and CG.

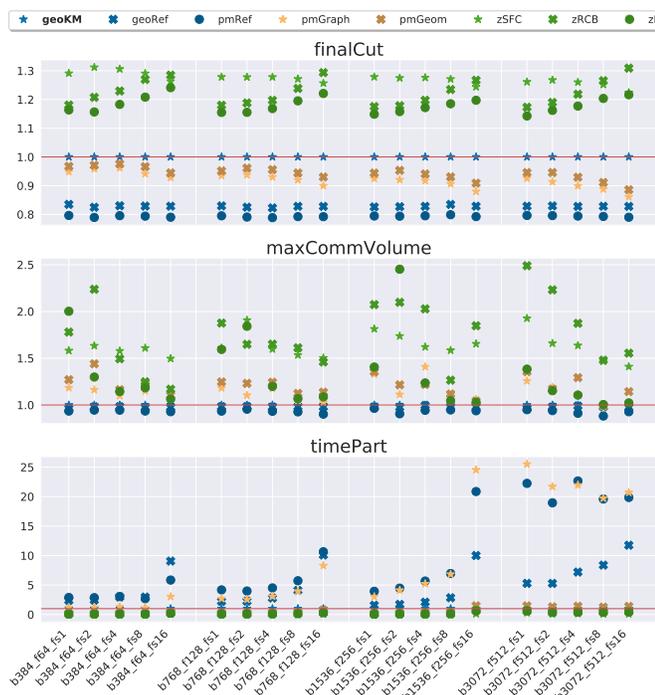


Fig. 3: Results for the refinetrace graph with  $24 \cdot 2^i, i = 4, \dots, 9$  PUs. All experiments belong to TOPO2. For the names in the x axis: ‘b’ is the number of blocks/PUs, next, ‘f’ indicated the number of fast PUs and ‘fs’ indicates the speed of the fast PUs. The slow PUs have speed 1 in all cases. These experiments were conducted in the Lise cluster.

For the benchmarks to be meaningful, we use a sufficiently large graph that still fits into the memory, `rdg_2d_29`. In Fig. 5 we see the edge cut and the running time of the CG solver per iteration for the linear systems (the SpMV results are similar to CG and thus omitted). The previous trend regarding the cut quality continues. However, while the quality in terms of edge cut differs across different tools, the communication volume and the running time per CG iteration show smaller deviations. We conjecture that a more severe heterogeneity will also translate into higher real-world differences.

## VII. CONCLUSIONS

So far, when working with established partitioning tools for load distribution in sparse matrix / graph problems on compute systems with different PUs and memory capacities, an additional preprocessing step is required. This preprocessing computes the different block sizes in a partition, which are then fed into the partitioners. For the first phase of this two-stage process, we proposed a greedy algorithm and proved it to be optimal. Our experiments for the second phase indicate that a combination of geometric and combinatorial methods yield the most promising results regarding the tradeoff between quality and running time. Overall, the results suggest that more work is necessary to support strongly heterogeneous computations out of the box and in a scalable way. This particularly includes

a one-phase approach for even more complex scenarios with different computational weights and communication costs.

*Acknowledgments:* This work is partially supported by German Research Foundation (DFG) grant ME 3619/4-1. We also acknowledge support by the North-German Supercomputing Alliance (HLRN).

## REFERENCES

- [1] David A. Bader, Henning Meyerhenke, Peter Sanders, and Dorothea Wagner, editors. *Graph Partitioning and Graph Clustering - 10th DIMACS Implementation Challenge Workshop, Georgia Institute of Technology, Atlanta, GA, USA, February 13-14, 2012. Proceedings*, volume 588 of *Contemporary Mathematics*. American Mathematical Society, 2013. VI-0c
- [2] C.E. Bichot and P. Siarry. *Graph Partitioning*. ISTE. Wiley, 2013. I
- [3] E. G. Boman, U. V. Catalyurek, C. Chevalier, and K. D. Devine. The Zoltan and Isorropia parallel toolkits for combinatorial scientific computing: Partitioning, ordering, and coloring. *Scientific Programming*, 20(2):129–150, 2012. III-0a, VI-0b
- [4] Thomas Brandes, Eric Schricker, and Thomas Soddemann. The lama approach for writing portable applications on heterogeneous architectures. In Michael Griebel, Anton Schüller, and Marc Alexander Schweitzer, editors, *Projects and Products of Fraunhofer SCAI*. Springer Berlin Heidelberg, 2017. VI-0a
- [5] Thang Nguyen Bui and Curt Jones. A heuristic for reducing fill-in in sparse matrix factorization. In Richard F. Sincovec, David E. Keyes, Michael R. Leuze, Linda R. Petzold, and Daniel A. Reed, editors, *Proceedings of the Sixth SIAM Conference on Parallel Processing for Scientific Computing, PPSC 1993, Norfolk, Virginia, USA, March 22-24, 1993*, pages 445–452. SIAM, 1993. I
- [6] Aydin Buluç, Henning Meyerhenke, Ilya Safro, Peter Sanders, and Christian Schulz. Recent advances in graph partitioning. In Lasse Kliemann and Peter Sanders, editors, *Algorithm Engineering - Selected Results and Surveys*, volume 9220 of *Lecture Notes in Computer Science*, pages 117–158. 2016. I, II-A
- [7] Jen-Hao Chen, Ren-Chuen Chen, and Jinn-Liang Liu. A gpu poisson-fermi solver for ion channel simulations. *Computer Physics Communications*, 229:99 – 105, 2018. VI-0d
- [8] Cédric Chevalier, Franck Ledoux, and Sébastien Morais. *A Multilevel Mesh Partitioning Algorithm Driven by Memory Constraints*, pages 85–95. SIAM, 2020. III-0c
- [9] Avery Ching, Sergey Edunov, Maja Kabiljo, Dionysios Logothetis, and Sambavi Muthukrishnan. One trillion edges: Graph processing at facebook-scale. *Proc. VLDB Endow.*, 8(12):1804–1815, 2015. I
- [10] Mehmet Deveci, Sivasankaran Rajamanickam, Karen D. Devine, and Ümit V. Çatalyürek. Multi-jagged: A scalable parallel spatial partitioning algorithm. *IEEE Trans. Parallel Distrib. Syst.*, 27, 2016. III-0a
- [11] Florent Duchaine, Sandrine Berger, Gabriel Staffelbach, and Gicquel L.Y.M. Partitioned high performance code coupling applied to cfd. pages 3–12, 02 2017. I
- [12] C.M. Fiduccia and R.M. Mattheyses. A linear-time heuristic for improving network partitions. *Design Automation, 1982. 19th Conference on*, pages 175–181, June 1982. III-0a
- [13] Daniel Funke, Sebastian Lamm, Peter Sanders, Christian Schulz, Darren Strash, and Moritz von Looz. Communication-free massively distributed graph generation. In *2018 IEEE International Parallel and Distributed Processing Symposium, IPDPS 2018, Vancouver, BC, Canada, May 21 – May 25, 2018*, 2018. VI-0c
- [14] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., 1979. I
- [15] Roland Glantz, Maria Predari, and Henning Meyerhenke. Topology-induced enhancement of mappings. In *Proceedings of the 47th International Conference on Parallel Processing, ICPP 2018, Eugene, OR, USA, August 13-16, 2018*, pages 9:1–9:10. ACM, 2018. I
- [16] Michael T. Heath and Padma Raghavan. A cartesian parallel nested dissection algorithm, 1994. III-0a
- [17] B. Hendrickson and R. Leland. A multi-level algorithm for partitioning graphs. In *Proceedings Supercomputing '95*. ACM Press, 1995. I, III-0a
- [18] Bruce Hendrickson and Tamara G. Kolda. Graph partitioning models for parallel computing. *Parallel Comput.*, 26(12), November 2000. II-A

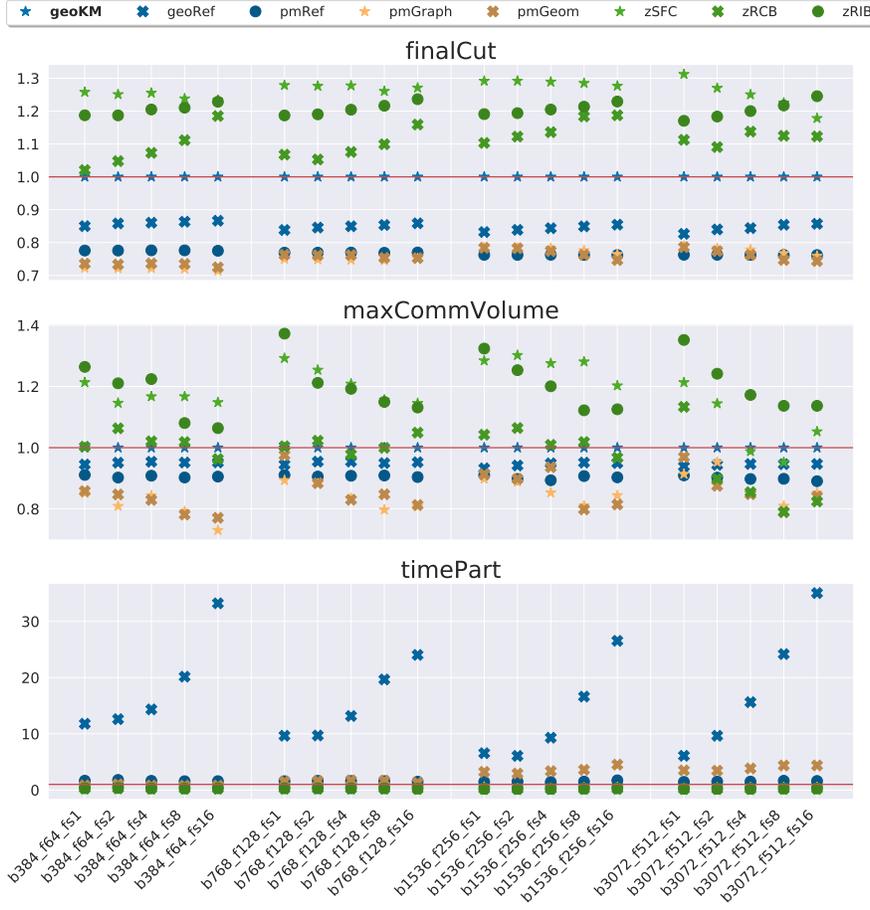


Fig. 4: Results for the 3D *rgg* and *rdg* graphs with  $2^{29}$  vertices, using  $24 \cdot 2^i$ ,  $i = 4, \dots, 9$ , PUs. The data points show the geometric mean of the value for each graph and are relative to the corresponding balanced  $k$ -means value. All experiments belong to TOPO2. Regarding the names on the  $x$  axis: ‘b’ is the number of blocks/PUs, ‘f’ indicates the number of fast PUs, and ‘fs’ indicates the speed of the fast PUs. The slow PUs have speed 1 in all cases. These experiments were run on Lise.

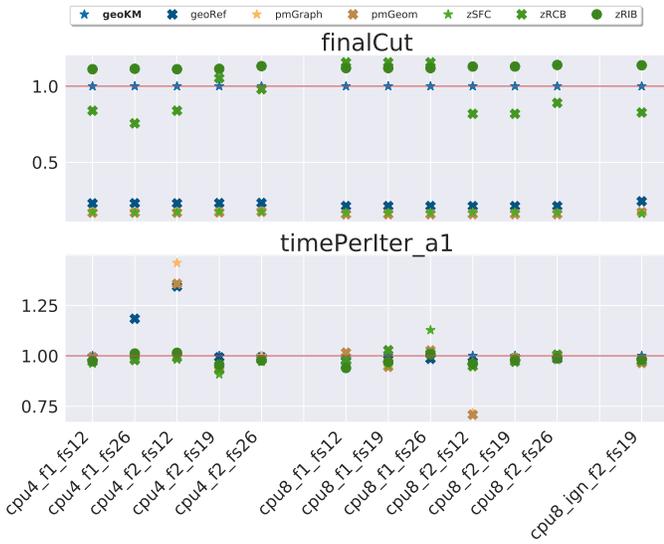


Fig. 5: Cut values and time per CG iteration using different settings within TOPO3 for the *rdg\_2d\_29* graph.

- [19] Torsten Hoefler, Emmanuel Jeannot, and Guillaume Mercier. An Overview of Process Mapping Techniques and Algorithms in High-Performance Computing. In Emmanuel Jeannot and Julius Zilinskas, editors, *High Performance Computing on Complex Environments*, pages 75–94. Wiley, June 2014. III-0c
- [20] 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 2010, Atlanta, Georgia, USA, April 2010 - Conference Proceedings*. IEEE, 2010. V
- [21] G. Karypis and V. Kumar. Multilevel algorithms for multi-constraint graph partitioning. In *SC '98: Proceedings of the 1998 ACM/IEEE Conference on Supercomputing*, pages 28–28, 1998. III-0b
- [22] G. Karypis and V. Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on Scientific Computing*, 20(1):359, 1999. III-0a, III-0a
- [23] Jeremy Kepner, Peter Aaltonen, David A. Bader, Aydin Buluç, Franz Franchetti, John R. Gilbert, Dylan Hutchison, Manoj Kumar, Andrew Lumsdaine, Henning Meyerhenke, Scott McMillan, Carl Yang, John D. Owens, Marcin Zalewski, Timothy G. Mattson, and José E. Moreira. Mathematical foundations of the graphblas. In *2016 IEEE High Performance Extreme Computing Conference, HPEC 2016, Waltham, MA, USA, September 13-15, 2016*, pages 1–9. IEEE, 2016. I, II
- [24] B. W. Kernighan and S. Lin. An efficient heuristic for partitioning graphs. *Bell Systems Technical Journal*, 49:291–308, 1970. I
- [25] O. Marquardt and S. Schamberger. Open benchmarks for load balancing heuristics in parallel adaptive finite element computations. In *Proc. of the Intern. Conf. on Parallel and Distributed Processing Techniques and*

Graphs	competitors	finalCut				maxCommVolume				timePart			
		t1_f8	t1_f16	t2_f8	t2_f16	t1_f8	t1_f16	t2_f8	t2_f16	t1_f8	t1_f16	t2_f8	t2_f16
333SP	geoKM	43428	40230	41978	40108	2692	1740	2395	1422	1.96	2.43	2.96	28
	geoRef	42393	39294	40904	39025	2623	1701	2241	1342	14.12	13.26	14.87	12.76
	geoPM	41349	38591	<b>39967</b>	<b>38365</b>	2561	1671	1822	<b>1321</b>	114.13	97.47	101.5	102.77
	pmGraph	<b>40291</b>	38812	44078	41311	<b>1839</b>	1537	<b>1514</b>	1567	51.92	51.59	55.36	534
	pmGeom	40562	<b>38249</b>	43109	41635	1897	<b>1474</b>	1554	1477	3.75	10.18	8.72	8.15
	zSFC	96465	86767	99712	94397	2722	3366	3486	3135	<b>0.04</b>	<b>0.04</b>	<b>0.03</b>	<b>0.03</b>
	zRCB	90126	89530	80907	75655	6373	6133	4814	3236	2.46	2.85	2.46	2.77
	zRIB	52620	50356	51820	52117	2474	2168	1270	1556	2.48	3.49	4.94	2.83
NLR	geoKM	64093	61001	66647	65304	2549	2110	1943	2017	3.18	3.3	2.49	2.1
	geoRef	62438	59539	64986	63680	2489	2066	1873	1957	18.42	16	16.63	12.94
	geoPM	<b>61211</b>	<b>58267</b>	<b>63853</b>	<b>62544</b>	2443	1978	1832	<b>1838</b>	163.79	147.11	145.51	1483
	pmGraph	68397	63406	71938	68521	2513	2160	1980	2002	183.81	263.75	299.36	130.46
	pmGeom	67995	61760	71583	68238	2302	<b>1923</b>	2094	1886	7	8.86	7.58	4.57
	zSFC	93975	89967	97693	95187	3488	2686	2455	2664	<b>0.03</b>	<b>0.03</b>	<b>0.03</b>	<b>0.04</b>
	zRCB	78400	79242	78031	79429	<b>2061</b>	1970	<b>1797</b>	1929	2.81	3.13	3.56	2.63
	zRIB	77679	77681	80737	79637	2252	2059	2028	1875	3.3	3	3.27	2.76
hugebubbles-00020	geoKM	55607	53167	57706	55905	3567	3241	2924	2673	4.73	4.22	4.43	3.67
	geoRef	46059	43728	47666	46039	3459	3124	2798	2609	37.52	33.37	31	26.77
	geoPM	<b>43796</b>	<b>41536</b>	<b>45379</b>	<b>43754</b>	3259	2910	<b>2623</b>	<b>2399</b>	249.95	182.26	197.8	181.28
	pmGraph	49102	45169	51523	49960	3004	2930	3165	2778	319.49	1486	171.87	241.42
	pmGeom	48998	44915	51663	49568	<b>2832</b>	<b>2758</b>	3120	2562	7.42	5.77	6.57	7.86
	zSFC	68678	62895	71609	68779	4085	3675	3532	3467	<b>0.14</b>	<b>0.14</b>	<b>0.13</b>	<b>0.13</b>
	zRCB	72584	73623	78029	74534	3287	3269	3283	2890	3.31	3.28	3.38	3.73
	zRIB	69821	69446	71269	72760	3529	3996	3026	3464	3.77	3.16	3.51	3.7
hugegrace-00020	geoKM	49064	46688	51352	48714	2957	2605	2509	2093	2.37	2.93	2.75	2.34
	geoRef	40330	38600	42378	39911	2864	2531	2447	2029	26.6	19.73	20.84	209
	geoPM	<b>38063</b>	<b>36282</b>	<b>40150</b>	<b>37823</b>	2667	2359	2318	<b>1950</b>	153	226.38	154.79	1431
	pmGraph	41824	38749	44311	42069	2831	2368	<b>2096</b>	2127	142.34	204.86	180.49	208.21
	pmGeom	42079	39037	45245	42701	<b>2450</b>	<b>2310</b>	2125	2252	6.36	6.43	79	99
	zSFC	56457	52976	59297	56859	3574	3184	3374	2711	<b>0.09</b>	<b>0.13</b>	<b>0.09</b>	<b>0.09</b>
	zRCB	61547	62981	64210	62905	3331	2804	2755	2524	2.99	3.21	3.33	35
	zRIB	58929	57615	61530	60848	2776	2590	2500	2323	2.95	37	3.34	3.33
rdg_2d_23	geoKM	169940	159700	177770	170440	6540	5144	5777	4590	0.74	0.62	0.8	0.58
	geoRef	128020	122760	132780	129850	5689	4448	4899	4017	10.47	6.71	7.56	5.55
	geoPM	<b>102700</b>	<b>97184</b>	<b>107750</b>	<b>104440</b>	<b>4289</b>	<b>3445</b>	<b>3519</b>	<b>3151</b>	1.19	1.15	1.15	0.98
	pmGraph	111690	100240	115650	111950	6005	3988	3763	3826	0.41	0.4	0.39	0.37
	pmGeom	110190	103660	114760	112070	4331	4114	3783	3453	96.25	26.48	26.6	41.46
	zSFC	121460	114430	134600	128600	3270	3506	3987	3488	<b>0.04</b>	<b>0.05</b>	<b>0.04</b>	<b>0.04</b>
	zRCB	172660	184170	183960	193760	5413	5181	4020	4887	0.1	0.08	0.09	0.07
	zRIB	191780	181910	197910	196270	5647	4810	4921	4451	0.35	0.2	0.45	0.26
alyaTestCaseB	geoKM	2066100	1949600	2171400	2081900	93325	74778	72500	68538	5.13	5.64	4.97	5.59
	geoRef	2055800	1939300	2162500	2070500	92616	<b>74530</b>	<b>72225</b>	68407	74.41	60.75	53.32	49.3
	geoPM	2046800	1932000	2155500	2061100	92653	74847	72485	68530	7.62	7.42	6.85	6.82
	pmGraph	2085100	1956800	2234200	2140100	103480	100010	98172	85985	0.9	0.87	0.87	0.88
	pmGeom	<b>1955700</b>	<b>1815200</b>	<b>2121700</b>	<b>1949500</b>	101460	84799	78720	71147	6.85	7.27	9.98	16.22
	zSFC	2358300	2213100	2579300	2393200	82725	78263	84616	72471	<b>0.15</b>	<b>0.16</b>	<b>0.16</b>	<b>0.17</b>
	zRCB	2382200	2429400	2185500	2190800	89201	76273	74508	66581	0.6	0.61	0.74	0.48
	zRIB	2262400	2226600	2299900	2391100	<b>76607</b>	76252	74954	<b>61899</b>	51	3.28	4.83	5.47

TABLE IV: Exact values for a subset of topologies and graphs for 3 metrics: edge cut, maximum communication volume and partitioning time. The results are for 96 PUs and we include 4 topologies, 2 from TOPO1 and 2 from TOPO2 with 8 or 16 fast PUs. For the presented results, all fast PUs have a speed of 16 (fs16). Best results are marked in bold font.

Applications, (PDPTA'05). CSREA Press, 2005. VI-0c

[26] S. M. Mniszewski, M. J. Cawkwell, M. E. Wall, J. Mohd-Yusof, N. Bock, T. C. Germann, and A. M. N. Niklasson. Efficient parallel linear scaling construction of the density matrix for born-oppenheimer molecular dynamics. *Journal of Chemical Theory and Computation*, 11(10):4644–4654, 2015. PMID: 26574255. 1

[27] Irene Moulitsas and George Karypis. Architecture aware partitioning algorithms. In Anu G. Bourgeois and S. Q. Zheng, editors, *Algorithms and Architectures for Parallel Processing, pages 42–53*, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg. III-0b

[28] B. Nour-omid, A. Raefsky, and G. Lyzenga. Solving finite element equations on concurrent computers. 1986. III-0a

[29] S. Plimpton, B. Hendrickson, and J. Stewart. A parallel rendezvous algorithm for interpolation between multiple grids. In *SC '98: Proc. 1998 ACM/IEEE Conf. on Supercomputing*, pages 53–53, 1998. III-0b

[30] PRACE. Unified european applications benchmark suite, 2016. Accessed: 2017-09-30. VI-0c

[31] Christian Schulz and Jesper Larsson Träff. Better process mapping and sparse quadratic assignment. In *16th Intl. Symp. on Experimental Algorithms, SEA 2017*, volume 75 of *LIPIcs*, pages 4:1–4:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. II-B

[32] Moritz von Looz, Charilaos Tzovas, and Henning Meyerhenke. Balanced k-means for parallel geometric partitioning. In *Proc. 47th International Conference on Parallel Processing, ICPP 2018, Eugene, OR, USA, August, 2018*, pages 52:1–52:10. ACM, 2018. III-0a, V

[33] M. S. Warren and J. K. Salmon. A parallel hashed oct-tree n-body algorithm. In *Proceedings of the 1993 ACM/IEEE Conference on Supercomputing*, Supercomputing '93, page 12–21, New York, NY, USA, 1993. Association for Computing Machinery. III-0a

[34] Tao Zhang, Jingjie Zhang, Wei Shu, Min-You Wu, and Xiaoyao Liang. Efficient graph computation on hybrid cpu and gpu systems. *The Journal of Supercomputing*, 71(4):1563–1586, 2015. 1

## A. Proof of Lemma 1

*Proof.* We will prove that if some arbitrary PU in position  $i$  in the sorted sequence is non-saturated, then all PUs in positions greater than  $i$  are also non-saturated. Towards a contradiction, let  $d$  be the first iteration in which the first saturated PU appears *after* a non-saturated one. Consequently, when denoting the state of  $j\text{Load}$  and  $j\text{Speed}$  at the beginning of iteration  $i$ ,  $0 \leq i \leq k-1$ , by  $j\text{Load}^{(i)}$  and  $j\text{Speed}^{(i)}$  respectively, we get:

$$\begin{aligned} & \text{desW}(b_d) > m_{cap}(p_d) \\ \implies c_s(p_d) \cdot \frac{j\text{Load}^{(d)}}{j\text{Speed}^{(d)}} & > m_{cap}(p_d) \\ \implies \frac{c_s(p_d)}{m_{cap}(p_d)} & > \frac{j\text{Speed}^{(d)}}{j\text{Load}^{(d)}} \end{aligned} \quad (6)$$

By construction PU  $p_{d-1}$  is not saturated. Thus:

$$\begin{aligned} & \text{desW}(b_{d-1}) \leq m_{cap}(p_{d-1}) \\ \implies \frac{c_s(p_{d-1})}{m_{cap}(p_{d-1})} & \leq \frac{j\text{Speed}^{(d-1)}}{j\text{Load}^{(d-1)}} \end{aligned} \quad (7)$$

Since the PUs are sorted non-increasingly (and with Eq. (6)), we obtain:

$$\frac{c_s(p_{d-1})}{m_{cap}(p_{d-1})} \geq \frac{c_s(p_d)}{m_{cap}(p_d)} > \frac{j\text{Speed}^{(d)}}{j\text{Load}^{(d)}} \quad (8)$$

By Line 11, the computational load left at the end of iteration  $d$  is what was left beforehand minus what we assigned to PU  $p_d$  – similarly for the joint speed:

$$\begin{aligned} j\text{Load}^{(d)} &= j\text{Load}^{(d-1)} - tw(b_{d-1}) \\ j\text{Speed}^{(d)} &= j\text{Speed}^{(d-1)} - c_s(p_{d-1}) \end{aligned}$$

Let us take a closer look to  $j\text{Load}^{(d)}$ . Since PU  $p_{d-1}$  is non-saturated,  $tw(p_{d-1}) = \text{desW}(b_{d-1})$ , so that we get:

$$\begin{aligned} j\text{Load}^{(d)} &= j\text{Load}^{(d-1)} - tw(b_{d-1}) \\ &= j\text{Load}^{(d-1)} - c_s(p_{d-1}) \cdot \frac{j\text{Load}^{(d-1)}}{j\text{Speed}^{(d-1)}} \\ &= j\text{Load}^{(d-1)} \cdot \frac{j\text{Speed}^{(d-1)} - c_s(p_{d-1})}{j\text{Speed}^{(d-1)}} \\ &= j\text{Load}^{(d-1)} \frac{j\text{Speed}^{(d)}}{j\text{Speed}^{(d-1)}} \end{aligned} \quad (9)$$

Combining Eq. (8) and Eq. (9) yields:

$$\begin{aligned} \frac{c_s(p_{d-1})}{m_{cap}(p_{d-1})} & > \frac{j\text{Speed}^{(d)}}{j\text{Load}^{(d)}} = \frac{j\text{Speed}^{(d)}}{j\text{Load}^{(d-1)} \cdot \frac{j\text{Speed}^{(d)}}{j\text{Speed}^{(d-1)}}} \\ &= \frac{j\text{Speed}^{(d)} j\text{Speed}^{(d-1)}}{j\text{Load}^{(d-1)} j\text{Speed}^{(d)}} = \frac{j\text{Speed}^{(d-1)}}{j\text{Load}^{(d-1)}}, \end{aligned} \quad (10)$$

which is the desired contradiction of Eq. (7).  $\square$

## B. Proof of Theorem 1

*Proof.* The running time is dominated by the sorting step, since the loop has  $k$  iterations and all other operations need only constant time.

Clearly, all saturated PUs receive their optimal load, since more load would exceed their memory capacity. Let  $d$  now be the iteration in which the first non-saturated PU appears in our sorted sequence. We need to show that all PUs  $p_i$  with  $i \geq d$  receive a load that is proportionate to  $j\text{Load}^{(d)} \cdot \frac{c_s(p_d)}{j\text{Speed}^{(d)}}$  – as this minimizes Eq. (2).

Let us apply an inductive argument, starting in iteration  $d$ . Clearly, PU  $tw(b_d)$  fulfills this condition. For the inductive step, we assume the claim to be true for all iterations  $i'$  with  $d \leq i' \leq i$ . We now consider iteration  $i+1$ :

$$tw(b_{i+1}) = c_s(p_{i+1}) \cdot \frac{j\text{Load}^{(i)} - tw(b_i)}{j\text{Speed}^{(i)} - c_s(p_i)}$$

In order to prove that the assignment in iteration  $i+1$  is proportionate to the load and speed situation in iteration  $i$  (and thus  $d$  by our assumption), it is sufficient to show that

$$\frac{j\text{Load}^{(i)} - tw(b_i)}{j\text{Speed}^{(i)} - c_s(p_i)} = \frac{j\text{Load}^{(i)}}{j\text{Speed}^{(i)}}.$$

Inserting the value for  $tw(b_i)$  and some rearranging yields:

$$\begin{aligned} & \frac{j\text{Load}^{(i)} - tw(b_i)}{j\text{Speed}^{(i)} - c_s(p_i)} \\ &= \frac{j\text{Load}^{(i)} - \left( \frac{j\text{Load}^{(i)} \cdot c_s(p_i)}{j\text{Speed}^{(i)}} \right)}{j\text{Speed}^{(i)} - c_s(p_i)} \\ &= \frac{j\text{Speed}^{(i)} \left( j\text{Load}^{(i)} - \left( \frac{j\text{Load}^{(i)} \cdot c_s(p_i)}{j\text{Speed}^{(i)}} \right) \right)}{j\text{Speed}^{(i)} \left( j\text{Speed}^{(i)} - c_s(p_i) \right)} \\ &= \frac{j\text{Load}^{(i)} \left( \frac{j\text{Speed}^{(i)} - c_s(p_i)}{j\text{Speed}^{(i)} - c_s(p_i)} \right)}{j\text{Speed}^{(i)}} \\ &= \frac{j\text{Load}^{(i)}}{j\text{Speed}^{(i)}}. \end{aligned}$$

$\square$