

Social Hash Partitioner: A Scalable Distributed Hypergraph Partitioner

Igor Kabiljo
Facebook

ikabiljo@fb.com

Brian Karrer
Facebook

briankarrer@fb.com

Mayank Pundir
Facebook

mpundir@fb.com

Sergey Pupyrev
Facebook

spupyrev@fb.com

Alon Shalita
Facebook

alon@fb.com

ABSTRACT

We design and implement a distributed algorithm for balanced k -way hypergraph partitioning that minimizes fanout, a fundamental hypergraph quantity also known as the communication volume and $(k - 1)$ -cut metric, by optimizing a novel objective called *probabilistic fanout*. This choice allows a simple local search heuristic to achieve comparable solution quality to the best existing hypergraph partitioners.

Our algorithm is arbitrarily scalable due to a careful design that controls computational complexity, space complexity, and communication. In practice, we commonly process hypergraphs with billions of vertices and hyperedges in a few hours. We explain how the algorithm’s scalability, both in terms of hypergraph size and bucket count, is limited only by the number of machines available. We perform an extensive comparison to existing distributed hypergraph partitioners and find that our approach is able to optimize hypergraphs roughly 100 times bigger on the same set of machines.

We call the resulting tool *Social Hash Partitioner*, and accompanying this paper, we open-source the most scalable version based on recursive bisection.

1. INTRODUCTION

The goal of graph partitioning is to divide the vertices of a graph into a number of equal size components, so as to minimize the number of edges that cross components. It is a classical and well-studied problem with origins in parallel scientific computing and VLSI design placement [8, 9]. Hypergraph partitioning is relatively less studied than graph partitioning. Unlike a graph, in a hypergraph, an edge, referred to as a hyperedge, can connect to any number of vertices, as opposed to just two. The revised goal is to divide the vertices of a hypergraph into a number of equal size components, while minimizing the number of components hyperedges span.

While graph partitioning is utilized in a variety of applications, hypergraph partitioning can be a more accurate model of many real-world problems [10, 11, 14]. For example, it has been successfully applied for optimizing distributed systems [19, 26, 29] as well as distributed scientific computation [10, 14]. For another example, hypergraph partitioning accurately models the problem of minimizing the number of transactions in distributed data placement [11].

Our primary motivation for studying hypergraph partitioning comes from the problem of storage sharding common in distributed databases. Consider a scenario with a large dataset whose data records are distributed across several storage servers. A query to the database may consume several data records. If the data records are located on multiple servers, the query is answered by sending requests to each server. Hence, the assignment of data records to servers determines the number of requests needed to process a query; this number is often called the *fanout* of the query. Queries with low fanout can be answered more quickly, as there is less chance of contacting a slow server [12, 29]. Thus, a common optimization is to choose an assignment of data records that collocates the data required by different queries.

Storage sharding is naturally modeled by the hypergraph partitioning problem; see Figure 1. For convenience, and entirely equivalent to the hypergraph formulation, we follow the notation of [15, 19] and define the problem using a bipartite graph. Let $G = (\mathcal{Q} \cup \mathcal{D}, E)$ be an undirected bipartite graph with disjoint sets of query vertices, \mathcal{Q} , and data vertices, \mathcal{D} . The goal is to partition \mathcal{D} into k parts, that is, find a collection of k disjoint subsets V_1, \dots, V_k covering \mathcal{D} that minimizes an objective function. The resulting subsets, also called *buckets*, should be *balanced*, that is, $|V_i| \leq (1 + \epsilon) \frac{n}{k}$ for all $1 \leq i \leq k$ and some $\epsilon \geq 0$, where $n = |\mathcal{D}|$.

This optimization is exactly equivalent to balanced hypergraph partitioning with the set of vertices \mathcal{D} and hyperedges $\{v_1, \dots, v_r\}$, $v_i \in \mathcal{D}$ for every $q \in \mathcal{Q}$ with $\{q, v_i\} \in E$. The data vertices are shared between the bipartite graph and its hypergraph representation, and each query vertex corresponds to a single hyperedge that spans all data vertices connected to that query vertex in the bipartite graph. Figures 1b and 1c show this equivalent representation.

For a given partitioning $P = \{V_1, \dots, V_k\}$ and a query vertex $q \in \mathcal{Q}$, we define the *fanout* of q as the number of distinct buckets having a data vertex incident to q :

$$\text{fanout}(P, q) = |\{V_i : \exists \{q, v\} \in E, v \in V_i\}|.$$

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org.

Proceedings of the VLDB Endowment, Vol. 10, No. 11
Copyright 2017 VLDB Endowment 2150-8097/17/07.




Figure 1: The storage sharding problem modeled with the (b) bipartite graph and (c) hypergraph partitioning. Given three queries (red), $\{1, 2, 6\}$, $\{1, 2, 3, 4\}$, $\{4, 5, 6\}$, the goal is to split six data vertices (blue) into two buckets, V_1 and V_2 , so that the average query fanout is minimized. For $V_1 = \{1, 2, 3\}$, $V_2 = \{4, 5, 6\}$, fanout of the queries is 2, 2, and 1, respectively.

The quality of partitioning P is the average query fanout:

$$\text{fanout}(P) = \frac{1}{|Q|} \sum_{q \in Q} \text{fanout}(P, q).$$

The fanout minimization problem is, given a graph G , an integer $k > 1$, and a real number $\epsilon \geq 0$, find a partitioning of G into k subsets with the minimum average fanout.

Modulo irrelevant constant factors, the fanout is also called the communication volume [13] and the $(k - 1)$ cut metric. Arguably this is the most widely used objective for hypergraph partitioning [7, 13].¹ Fanout is also closely related to the sum of external degrees, but not identical [22].²

Not surprisingly, the hypergraph partitioning problem is NP-hard [6]. Hence, exact algorithms are only capable of dealing with fairly small problem sizes. Practical approaches all use heuristics. Even then, a significant problem with the majority of proposed algorithms is scalability. Scalability is desirable because modern graphs can be massive; for example, the Facebook Social Graph contains billions of vertices and trillions of edges, consuming many hundred of petabytes of storage space [29].

Scalability can be achieved through distributed computation, where the algorithm is distributed across many workers that compute a solution in parallel. Earlier works on distributed hypergraph partitioning have proposed and implemented such algorithms, but as we argue in Section 2 and experimentally demonstrate in Section 4, *none of the existing tools are capable of solving the problem at large scale*. No existing partitioner was able to partition a hypergraph with a billion hyperedges using four machines, and hypergraphs relevant to storage sharding at Facebook can be two or three orders of magnitude larger.

With this in mind, *we design, implement, and open source a scalable distributed algorithm for fanout minimization*. We denote the resulting implementation *Social Hash Partitioner* (SHP) because it can be used as the hypergraph partitioning component of the Social Hash framework introduced in [29]. The greater framework has several other major components and the specific choice of hypergraph partitioner was only briefly mentioned previously, and this entirely self-contained paper delves into our partitioner in detail. The contributions of the paper are the following:

¹Although communication volume seems a more popular term for the measure in hypergraph partitioning community, we utilize fanout here to be consistent with database sharding terminology.

²The sum of external degrees (SOED) objective is equivalent to unnormalized fanout plus the number of query vertices with fanout greater than one; that is, SOED is the communication volume plus the hyperedge cut.

- We design an iterative algorithm for hypergraph partitioning aimed at optimizing fanout through a classical local search heuristic [23] with two substantive modifications. The first is a novel objective function that generalizes fanout called *probabilistic fanout*; using this objective for fanout minimization improves result quality and algorithm convergence. Our second modification facilitates a distributed implementation of the algorithm.
- We describe SHP and provide a detailed implementation of the algorithm that runs in a parallel manner, carefully explaining how it limits memory usage, computation, and communication. SHP relies on the vertex-centric programming model and scales to hypergraphs with billions of vertices and hyperedges with a reasonable number of machines. In addition to optimizing fanout, we also show that SHP can optimize other hypergraph objectives at scale. We have open sourced the simple code for this implementation.
- We present results of extensive experiments on a collection of real-world and synthetic hypergraphs. These results show that SHP is arbitrarily scalable while producing partitions of comparable quality to existing partitioners. SHP is able to partition hypergraphs 100x larger than existing distributed partitioners, making it the only solution capable of partitioning billion-vertex bipartite graphs using a cluster of machines.

We emphasize that the paper describes an algorithmic solution utilized in our Social Hash framework. We refer the reader to [29] for the details of the framework, additional applications, and real-world experiments.

2. RELATED WORK

There exists a rich literature on graph partitioning from both theoretical and practical points of view. We refer the reader to surveys by Bichot and Siarry [8] and by Buluç et al. [9]. Next we discuss existing approaches for storage sharding that utilize graph partitioning and review theoretical results on the problem. Then we describe existing libraries for hypergraph partitioning, focusing on the publicly available ones. Finally, we analyze limitations of the tools, which motivate the development of our new solution, SHP.

Storage sharding. Data partitioning is a core component of many existing distributed data management systems [11, 19, 24, 26, 29]. The basic idea is to co-locate related data so as to minimize communication overhead. In many systems, the data partitioning problem is reduced to a variant of graph partitioning [11, 19, 24]; a solution is then

found using a publicly available library such as `Metis` [21]. However, hypergraph partitioning is a better model in several scenarios of storage sharding [26, 29]. There exists much fewer tools for partitioning of hypergraphs, as the problem tends to be harder and more computationally expensive than graph partitioning.

Theoretical aspects. The fanout minimization problem is a generalization of balanced k -way graph partitioning (also called minimum bisection when $k = 2$), which is a central problem in design and analysis of approximation algorithms. Andreev and Räcke [6] show that, unless $P=NP$, there is no algorithm with a finite approximation factor for the problem when one requires a perfect balance, that is, $\varepsilon = 0$. Hence, most works focus on the case $\varepsilon > 0$. Leighton et al. [27] and Simon and Teng [30] achieved an $\mathcal{O}(\log k \log n)$ approximation algorithm for $\varepsilon = 1$, that is, when the maximum size of the resulting buckets is $\frac{2n}{k}$. The bound has been improved to $\mathcal{O}(\sqrt{\log k \log n})$ for $\varepsilon = 1$ by Krauthgamer et al. [25] and to $\mathcal{O}(\log n)$ for $\varepsilon > 0$ by Feldmann and Foschini [17].

We stress that being of high theoretical importance, the above approximation algorithms are too slow to be used for large graphs, as they require solving linear programs. Hence, most of the existing methods for graph and hypergraph partitioning are heuristics based on a simple local search optimization [9]. We follow the same direction and utilize a heuristic that can be implemented in an efficient way. It is unlikely that one can provide strong theoretical guarantees on a local search algorithm for our problem. Since fanout minimization is a generalization of minimum bisection, it would imply a breakthrough result.

Existing tools for hypergraph partitioning. As a generalization of graph partitioning, hypergraph partitioning is a more complicated topic and the corresponding algorithms are typically more compute and memory intensive. Here we focus on existing solutions designed for hypergraph partitioning. `PaToH` [10], `hMetis` [20], and `Mondriaan` [32] are software packages providing single-machine algorithms. The tools implement different variants of local refinement algorithms, such as Kernighan-Lin [23] or Feduccia-Mattheyses [18], that incrementally swap vertices among partitions to reduce an optimization objective, until the process reaches a local minimum. Since such local search algorithms can suffer from getting stuck in local minima, a multi-level paradigm is often used. The idea is to create a sequence of “coarsened” hypergraphs that approximate the original hypergraph but have a smaller size. Then the refinement heuristic is applied on the smaller hypergraph, and the process is reverted by an uncoarsening procedure. Note that the above software packages all require random-access to the hypergraph located in memory; as a result these packages can only handle smaller hypergraphs.

`Parkway` [31] and `Zoltan` [14] are distributed hypergraph partitioners that are based on a parallel version of the multi-level technique. Unfortunately, as we argue below and show in our experiments, the provided implementations do not scale well. We also mention a number of more recent tools for hypergraph partitioning. `UMP` [13] is a serial partitioner with a novel refinement heuristic. It aims at minimizing several objective functions simultaneously. `HyperSwap` [33] is a distributed algorithm that partitions hyperedges, rather than vertices. `rFM` [28], allows replication of vertices in ad-

dition to vertex moves. We do not include these algorithms in our experimental study as they are not open-source and they are rather complex to be re-implemented in a fair way.

Limitations of existing solutions. `Parkway` and `Zoltan` are two hypergraph partitioners that are designed to work in a distributed environment. Both of the tools implement a multi-level coarse/refine technique [20]. We analyzed the algorithms and identified the following scalability limitations.

- First, multi-level schemes rely on an intermediate step in which the coarsest graph is partitioned on a single machine before it gets uncoarsened. While the approach is applicable for graph partitioning (when the coarsest graph is typically fairly small), it does not always work for hypergraphs. For large instances, the number of *distinct* hyperedges can be substantial, and even the coarsest hypergraph might not fit the memory of a single machine.
- Second, the refinement phase itself is often equivalent to the local refinement scheme presented in our work, which if not done carefully can lead to scalability issues. For example, `Parkway` is using a single coordinator to approve vertex swaps while retaining balance. This coordinator holds the concrete lists of vertices and their desired movements, which leads to yet another single machine bottleneck.
- Third, the amount of communication messages between different machines/processors is an important aspect for a distributed algorithm. Neither `Zoltan` nor `Parkway` provide strong guarantees on communication complexity. For example, the authors of `Zoltan` present their evaluation for mesh-like graphs (commonly used in scientific computing) and report relatively low communication overhead. Their results might not hold for general non-planar graphs.

In contrast, our algorithm is designed (as explained in the next section) to avoid these single machine bottlenecks and communication overheads.

3. SOCIAL HASH PARTITIONER

Our algorithm for the fanout minimization problem, as mentioned in Section 1, assumes that the input is represented as a bipartite graph $G = (\mathcal{Q} \cup \mathcal{D}, E)$ with vertices representing queries and data objects, and edges representing which data objects are needed to answer the queries. The input also specifies the number of servers that are available to serve the queries, $k > 1$, and the allowed imbalance, $\varepsilon > 0$.

3.1 Algorithm

For ease of presentation, we start with a high-level overview of our algorithm. The basic idea is inspired by the Kernighan-Lin heuristic [23] for graph partitioning; see Algorithm 1. The algorithm begins with an initial random partitioning of data vertices into k buckets. For every vertex, we independently pick a random bucket, which for large graphs guarantees an initial perfect balance. The algorithm then proceeds in multiple steps by performing vertex swaps between the buckets in order to improve an objective function. The process is repeated until a convergence criterion is met (e.g., no swapped vertices) or the maximum number of iterations is reached.

Although the algorithm is similar to the classical one [23], we introduce two critical modifications. The first concerns

Algorithm 1: Fanout Optimization

Input : graph $G = (Q \cup \mathcal{D}, E)$, the number of buckets k , imbalance ratio ε
Output: buckets V_1, V_2, \dots, V_k

for $v \in \mathcal{D}$ **do** /* initial partitioning */
 \lfloor $\text{bucket}[v] \leftarrow \text{random}(1, k)$;
repeat /* local refinement */
 for $v \in \mathcal{D}$ **do**
 for $i = 1$ **to** k **do**
 \lfloor $\text{gain}_i[v] \leftarrow \text{ComputeMoveGain}(v, i)$;
 /* find best bucket */
 $\text{target}[v] \leftarrow \arg \max_i \text{gain}_i[v]$;
 /* update matrix */
 if $\text{gain}_{\text{target}[v]}[v] > 0$ **then**
 \lfloor $S_{\text{bucket}[v], \text{target}[v]} \leftarrow S_{\text{bucket}[v], \text{target}[v]} + 1$;
 /* compute move probabilities */
 for $i, j = 1$ **to** k **do**
 \lfloor $\text{probability}[i, j] \leftarrow \frac{\min(S_{i,j}, S_{j,i})}{S_{i,j}}$;
 /* change buckets */
 for $v \in \mathcal{D}$ **do**
 if $\text{gains}[v] > 0$ **and**
 $\text{random}(0, 1) < \text{probability}[\text{bucket}[v], \text{target}[v]]$
 then
 \lfloor $\text{bucket}[v] \leftarrow \text{target}[v]$;
 until converged or iteration limit exceeded;

the objective function and is intended to improve the quality of the final result. The second is related to how we choose and swap vertices between buckets; this modification is needed to make distributed implementations efficient.

Optimization objective. Empirically, we found that fanout is rather hard to minimize with a local search heuristic. Such a heuristic can easily get stuck in a local minimum for fanout minimization. Figure 2 illustrates this with an example which lacks a single move of a data vertex that improves fanout. All move gains are non-positive, and the local search algorithm stops in the non-optimal state. To alleviate this problem, we propose a modified optimization objective and assume that a query $q \in Q$ only requires an adjacent data vertex $v \in \mathcal{D}$, $\{q, v\} \in E$ for some probability $p \in (0, 1)$ fixed for all queries. This leads to so-called *probabilistic fanout*. The probabilistic fanout of a given query q , denoted by $\text{p-fanout}(q)$, is the expected number of servers that need to be contacted to answer the query given that each adjacent server needs to be contacted with independent probability p .

Formally, let $P = \{V_1, \dots, V_k\}$ be a given partitioning of \mathcal{D} , and let the number of data vertices in bucket V_i adjacent to query vertex q be $n_i(q) = |\{v : v \in V_i \text{ and } \{q, v\} \in E\}|$. Then server i is expected to be queried with probability $1 - (1-p)^{n_i(q)}$. Thus, the p-fanout of q is $\sum_{i=1}^k (1 - (1-p)^{n_i(q)})$, and our probabilistic fanout objective, denoted p-fanout, for Algorithm 1 is, given $p \in (0, 1)$, minimize

$$\frac{1}{|Q|} \sum_{q \in Q} \text{p-fanout}(q) = \frac{1}{|Q|} \sum_{q \in Q} \sum_{i=1}^k (1 - (1-p)^{n_i(q)}).$$




Figure 2: An example in which no single move/swap of data vertices improves query fanout. Probabilistic fanout (for every $0 < p < 1$) can be improved by exchanging buckets of vertices 4 and 5 or by exchanging buckets of vertices 3 and 6. Applying both of the swaps reduces (non-probabilistic) fanout of q_1 and q_3 , which yields an optimal solution.

This revised objective is a smoother version of the (non-probabilistic) fanout. It is simple to observe that $\text{p-fanout}(q)$ is less than or equal to $\text{fanout}(q)$ for all $q \in Q$. If the number of data adjacencies of q in bucket i is large enough, that is, $n_i(q) \gg 1$, then the bucket contributes to the objective a value close to 1. If $n_i(q) = 1$, then the contribution is p , which is smaller for $p < 1$. In contrast, the non-probabilistic fanout contribution is simply 1, the same for all cases with $n_i(q) \geq 1$.

From a theoretical perspective, the way probabilistic fanout smooths fanout is by averaging the fanout objective over an ensemble of random graphs similar to the bipartite graph being partitioned. A bipartite graph from this random graph ensemble is created by independently selecting edges of the original bipartite graph with probability p . Then the probabilistic fanout is precisely the expectation of fanout across this random graph ensemble. In essence, p-fanout minimization is forced to select a partition that performs robustly across a collection of similar hypergraphs, reducing the impact of local minima. With the new objective the state in Figure 2 is no longer a local minimum, as data vertices 1 and 2 could be swapped to improve the p-fanout of both q_1 and q_3 .

An interesting aspect of p-fanout is how it behaves in extreme cases. As we show next, when $p \rightarrow 1$, p-fanout becomes the (non-probabilistic) fanout. When $p \rightarrow 0$, the new measure is equivalent to optimizing a weighted *edge-cut*, an objective suggested in prior literature [4, 5, 10]. In practice it means that p-fanout is a generalization of these measures and Algorithm 1 can be utilized to optimize either by setting small or large values of p .

LEMMA 1. *Minimizing p-fanout in the limit as $p \rightarrow 1$ is equivalent to minimizing fanout.*

PROOF. We write $\text{p-fanout}(q)$ for a query $q \in Q$ as

$$\sum_{i=1}^k (1 - (1-p)^{n_i(q)}) = \sum_{i=1}^k (1 - e^{n_i(q) \log(1-p)}).$$

Now as $p \rightarrow 1$, $\log(1-p)$ goes to negative infinity and the exponential term is zero unless $n_i(q) = 0$ in which case it equals one. Let $\delta(x) = 1$ if x is true and $\delta(x) = 0$, otherwise. In the limit, the above expression equals

$$\sum_{i=1}^k (1 - \delta(n_i(q) = 0)) = \sum_{i=1}^k \delta(n_i(q) > 0),$$

which is fanout of q . \square

Next we show that optimizing p-fanout as $p \rightarrow 0$ is equivalent to a graph partitioning problem on an edge-weighted graph constructed from data vertices. For a pair of data vertices, $u, v \in \mathcal{D}$, let $w(u, v)$ be the number of common queries shared by these data vertices, that is, $w(u, v) = |\{q \in \mathcal{Q} : \{q, u\} \in E \text{ and } \{q, v\} \in E\}|$. Consider a (complete) graph with vertex set \mathcal{D} and let $w(u, v)$ be the weight of an edge between $u, v \in \mathcal{D}$. For a given graph and a partition of its vertices, an *edge-cut* is the sum of edge weights between vertices in different buckets.

LEMMA 2. *Minimizing p-fanout in the limit as $p \rightarrow 0$ is equivalent to graph partitioning amongst the data vertices while minimizing weighted edge-cut, where the edge weight between $u \in \mathcal{D}$ and $v \in \mathcal{D}$ is given by $w(u, v)$.*

PROOF. We begin from the definition of p-fanout and consider the Taylor expansion around $p = 0$:

$$\begin{aligned} \text{p-fanout} &= \sum_{q \in \mathcal{Q}} \sum_{i=1}^k \left(1 - (1-p)^{n_i(q)}\right) \\ &= \sum_{q \in \mathcal{Q}} \sum_{i=1}^k \left(n_i(q)p + n_i(q)(n_i(q)-1)\frac{p^2}{2} + \mathcal{O}(p^3) \right) \\ &= C - \frac{p^2}{2} \sum_{q \in \mathcal{Q}} \sum_{i=1}^k n_i(q)^2 + \mathcal{O}(p^3). \end{aligned}$$

The first term, $C = \sum_{q \in \mathcal{Q}} \sum_{i=1}^k n_i(q)(\frac{p^2}{2} - p)$, is a constant proportional to the number of edges in the graph. Thus, it is irrelevant to the minimization. The last term, $\mathcal{O}(p^3)$, can also be ignored for optimization when $p \rightarrow 0$. We simplify the second term further.

$$\begin{aligned} -\frac{p^2}{2} \sum_{q \in \mathcal{Q}} \sum_{i=1}^k n_i(q)^2 &= -\frac{p^2}{2} \sum_{i=1}^k \sum_{q \in \mathcal{Q}} n_i(q)^2 \\ &= -\frac{p^2}{2} \sum_{i=1}^k \left(\sum_{u \in V_i} \sum_{v \in V_i} w(u, v) \right). \end{aligned}$$

Therefore, minimizing probabilistic fanout in the limit as $p \rightarrow 0$ is equivalent to maximizing the sum, taken over all buckets, of edge weights between data vertices within the same bucket, or *maximizing within-bucket edge weights*. Alternatively, this is also equivalent to *minimizing intra-bucket edge weights*, that is, minimizing weighted edge-cut between buckets with edge weights given by $w(u, v)$. \square

This $p \rightarrow 0$ limit is interesting because the resulting optimization is an instance of the clique-net model suggested as a heuristic for hypergraph partitioning [4, 5, 10]. The idea is to convert the hypergraph partitioning problem to the (simpler) graph partitioning problem. To this end, a hypergraph is transformed to an edge-weighted unipartite graph, G^c , on the same set of vertices, by adding a clique amongst all pairs of vertices connected to a hyperedge. Multiple edges between a vertex pair in the resulting graph are combined by summing their respective weights. The buckets produced by a graph partitioning algorithm on the new graph are then used as a solution for hypergraph partitioning.

An obstacle for utilizing the clique-net model is the size of the resulting (unipartite) graph G^c . If there is a hyperedge

connecting $\Omega(n)$ vertices, then G^c contains $\Omega(n^2)$ edges, even if the original hypergraph is sparse. Hence a common strategy is to use some variant of edge sampling to filter out edges with low weight in G^c [4, 5, 10]. Lemma 2 shows that this is unnecessary: One can apply Algorithm 1 with a small value of fanout probability, p , for solving the hypergraph partitioning problem in the clique-net model.

Performing swaps. In order to iteratively swap vertices between buckets, we compute *move gains* for the vertices, that is, the difference of the objective function after moving a vertex from its current bucket to another one. For every vertex $v \in \mathcal{D}$, we compute k values, referred to as $gain_i[v]$, indicating the move gains to every bucket $1 \leq i \leq k$. Then every vertex chooses a *target* bucket that corresponds to the highest move gain. (For minimization, we select the bucket with the lowest move gain, or equivalently the highest negative move gain.) This information is used to calculate the number of vertices in bucket i that chose target bucket j , denoted $S_{i,j}$, for all pairs of buckets $1 \leq i, j \leq k$. Ideally, we would move all these vertices from i to j to maximally improve the objective. However, to preserve balance across buckets, we exchange only $\min(S_{i,j}, S_{j,i})$ pairs of vertices between buckets i and j .

In a distributed environment, we cannot pair vertices together for swaps easily, and so we instead perform the swaps in an approximate manner. Our implementation defines a probability for each vertex in bucket i with target bucket j to be moved as $\frac{\min(S_{i,j}, S_{j,i})}{S_{i,j}}$. All vertices are then simultaneously moved to their target buckets respecting the computed probabilities, for all target buckets. With this choice of probabilities, the expected number of vertices moved from i to j and from j to i is the same; that is, the balance constraint is preserved in expectation. After that swap, we recompute move gains for all data vertices and proceed with the next iteration.

We compute the move gains as follows. Assume that a vertex $v \in \mathcal{D}$ is moved from bucket i to bucket j . Our objective function, p-fanout, might change only for queries adjacent to v and for the terms corresponding to V_i and V_j . Let $\mathcal{N}(v) \subseteq \mathcal{Q}$ be the subset of queries adjacent to v . The move gain is then

$$\begin{aligned} gain_j(v) &= \sum_{q \in \mathcal{N}(v)} \left(2 - (1-p)^{n_i(q)-1} - (1-p)^{n_j(q)+1} \right) \\ &\quad - \sum_{q \in \mathcal{N}(v)} \left(2 - (1-p)^{n_i(q)} - (1-p)^{n_j(q)} \right) \\ &= p \cdot \sum_{q \in \mathcal{N}(v)} \left((1-p)^{n_j(q)} - (1-p)^{n_i(q)-1} \right). \quad (1) \end{aligned}$$

Next we provide details of the implementation.

3.2 Implementation

Our implementation relies on the vertex-centric programming model and runs in the Giraph framework [2]. In Giraph, the input graph is stored as a collection of vertices that maintain some local data (e.g., a list of adjacent vertices). The vertices are distributed to multiple machines in a cluster and communicate with each other via sending messages. A computation in Giraph is split into supersteps that each consist of the following processing steps: (i) a vertex executes a user-defined function based on local vertex data

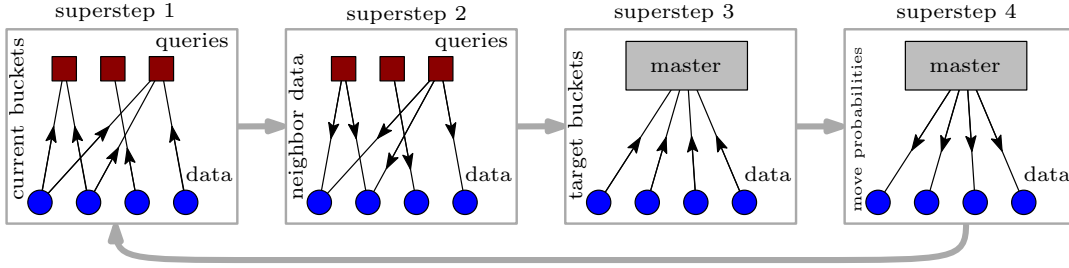


Figure 3: Distributed implementation of the fanout minimization algorithm in the vertex-centric framework with four supersteps and synchronization barriers between them: (1) collecting query neighbor data, (2) computing move gains, (3) proposal of target buckets, (4) sending move probabilities and performing moves.

and on received messages, (ii) the resulting output is sent along outgoing edges. Note that since vertices operate only with local data, such processing can easily be executed in parallel and in a distributed environment. Supersteps end with a synchronization barrier, which guarantees that messages sent in a given superstep are received at the beginning of the next superstep. The whole computation is executed iteratively for a certain number of rounds, or until a convergence property is met.

Algorithm 1 is implemented in the vertex-centric model in the following way; see Figure 3. The first two supersteps compute move gains for all data vertices. As can be seen from Equation 1, a move gain of $v \in \mathcal{D}$ depends on the state of adjacent query vertices. Specifically, we need to know the values $n_i(q)$ for every $q \in \mathcal{Q}$ and all buckets $1 \leq i \leq k$; we call this information the *neighbor data* of query q . The first superstep is used to collect the neighbor data; to this end, every data vertex sends its current bucket to the adjacent queries, which aggregate the received messages into the neighbor data. On the second superstep, the query vertices send their neighbor data back to adjacent data vertices, which use this information to compute their move gains according to Equation 1.

Once data vertices have computed move gains, we choose their target buckets and the next superstep aggregates this information in matrix S , that stores the number of candidate data vertices moving between pairs of buckets. The matrix is collected on a dedicated machine, called *master*, which computes move probabilities for the vertices. On the last superstep, the probabilities are propagated from master to all data vertices and the corresponding moves take effect. This sequence of four supersteps continues until convergence.

3.3 Complexity

Our primary consideration in designing the algorithm is to keep the implementation scalable to large instances. To this end, we limit space, computational, and communication complexity such that each is bounded by $\mathcal{O}(k|E|)$, where k is the number of buckets.

Space complexity. We ensure every vertex v of G consumes only $\mathcal{O}(|\mathcal{N}(v)| + k)$ memory, where $|\mathcal{N}(v)|$ is the number of neighbors of vertex v , so the total memory consumption is $\mathcal{O}(|E| + k|V|)$. Every vertex keeps its adjacency list, which is within the claimed limits. Additionally, every query vertex $q \in \mathcal{Q}$ stores its neighbor data containing $|\mathit{fanout}(q)| = \mathcal{O}(|\mathcal{N}(q)|)$ entries. Every data vertex $v \in \mathcal{D}$ stores move gains from its current bucket to all other buck-

ets, which includes up to $\mathcal{O}(k)$ values per vertex. The master machine stores the information of move proposals for every pair of buckets, that is, its memory consumption is $\mathcal{O}(k^2) = \mathcal{O}(k|V|)$, which is again within the claimed memory bound. Notice that since Giraph distributes vertices among machines in a Giraph cluster randomly, the scheme does not have a single memory bottleneck. All the machines are equally loaded, and in order to have enough memory for large graphs at a fixed number of buckets k , it is sufficient to increase the cluster size, while keeping the implementation unchanged.

Computational complexity. Our algorithm is within a bound of $\mathcal{O}(k|E|)$, assuming a constant number of refinement iterations. The computationally resource intensive steps are calculating the query neighbor data, which is performed in $\mathcal{O}(|E|)$ time, and processing this information by data vertices. The latter step is bounded by $\mathcal{O}(k|\mathcal{N}(v)|)$ steps for every $v \in \mathcal{D}$, as this is the amount of information being sent to v in superstep 2. Finally, computing and processing matrix S requires $\mathcal{O}(|V| + k^2) = \mathcal{O}(k|V|)$ time.

Communication complexity. Another important aspect of a distributed algorithm is its communication complexity, that is, the amount of messages sent between machines during its execution. The supersteps 1, 3, and 4 are relatively “lightweight” and require only $|E|$, $|V|$, and $|V|$ messages of constant size, respectively. The “heavy” one is superstep 2 in which every query vertex sends its neighbor data to all its neighbors. We can upper bound the amount of sent information by $k|E|$, as neighbor data for every $q \in \mathcal{Q}$ contains k entries. In practice, however, this is bounded by $\sum_{q \in \mathcal{Q}} \mathit{fanout}(q) \cdot |\mathcal{N}(q)|$, as the zero entries of neighbor data (having $n_i(q) = 0$) need not be sent as messages. Hence, a reasonable estimation of the amount of sent messages is $\mathit{fanout} \cdot |E|$ per iteration of Algorithm 1, where fanout is the average fanout of queries on the current iteration.

We stress here that Giraph has several built-in optimizations that can further reduce the amount of sent and received messages. For example, if exactly the same message is sent between a pair of machines several times (which might happen, for example, on superstep 2), then the messages are combined into a single one. Similarly, if a message is sent between two vertices residing on the same machine, then the message can be replaced with a read from the local memory of the machine.

Another straightforward optimization is to maintain some state of the vertices and only recompute the state between

iterations of the algorithm when necessary. A natural candidate is the query neighbor data, which is recomputed only when an adjacent data vertex is moved; if the data vertex stays in the same bucket on an iteration, then it does not send messages on superstep 1 for the next iteration.

Recursive partitioning. The discussion above suggests that Algorithm 1 is practical for small values of k , e.g., when $k = \mathcal{O}(1)$. In this case the complexity of the algorithm is linear in the size of the input hypergraph. However in some applications, substantially more buckets are needed. In the extreme with $k = \Omega(|V|)$, the implementation requires quadratic time and memory, and the run-time for large instances could be unreasonable even on a large Giraph cluster.

A common solution to this problem in the hypergraph partitioning literature is to observe that the partitioning can be constructed recursively. In recursive partitioning, the algorithm splits data vertices into $r > 1$ parts V_1, \dots, V_r . This splitting algorithm is recursively applied to all the graphs induced by vertices $\mathcal{Q} \cup V_1, \dots, \mathcal{Q} \cup V_r$ independently. The process continues until we achieve k buckets, which requires $\lceil \log_r k \rceil$ levels of recursion. A typical strategy is to employ $r = 2$; that is, recursive bisection [15, 30].

Notice that Algorithm 1 can be utilized for recursive partitioning with just a single modification. At every recursion step, data vertices are constrained as to which buckets they are allowed to be moved to. For example, at the first level of recursive bisection, all data vertices are split into V_1 and V_2 . At the second level, the vertices are split into four buckets V_3, V_4, V_5, V_6 so that the vertices $v \in V_1$ are allowed to move between buckets V_3 and V_4 , and the vertices $v \in V_2$ are allowed to move between V_5 and V_6 . In general, the vertices of V_i for $1 \leq i \leq k/2$ are split into V_{2i+1} and V_{2i+2} .

An immediate implication of the constraint is that every data vertex only needs to compute r move gains on each iteration. Similarly, a query vertex needs to send only a subset of its neighbor data to data vertices that contains at most r values. Therefore, the memory requirement as well as the computational complexity of recursive partitioning is $\mathcal{O}(r|E|)$ per iteration, while the amount of messages sent on each iteration does not exceed $\mathcal{O}(r|E|)$, which is a significant reduction over direct (non-recursive) partitioning when $r \ll k$. This improvement sometimes comes with the price of reduced quality; see Section 4 for a discussion. Accompanying the paper, we open-source a version that performs recursive bisection, as it is the most scalable.

3.4 Advanced implementation

While the basic algorithm described above performs well, this subsection describes additional improvements that we have included in our implementation, motivated by our practical experience.

First, randomly selecting vertices to swap between a pair of buckets i and j may not select those with the highest move gains to swap. In the ideal serial implementation, we would have two queues of gains, one corresponding to vertices in bucket i that want to move to j , and the other for vertices in bucket j that want to move to i , sorted by move gain from highest to lowest. We would then pair vertices off for swapping from highest to lowest.

This is difficult to implement exactly in a distributed environment. Instead of maintaining two queues for each pair

of buckets, we maintain two histograms that contain the number of vertices with move gains in exponentially sized bins. We then match bins in the two histograms for maximal swapping with probability one, and then probabilistically pair the remaining vertices in the final matched bins. In superstep 4, the master distributes move probabilities for each bin, most of which are either one or zero. This change allows our implementation to focus on moving the most important gains first. A further benefit is that we can allow non-positive move gains for target buckets. A pair of positive and negative histogram bins can swap if the sum of the gains is expected to be positive, which frees up additional movement in the local search.

Additionally, we utilize the imbalance allowed by ε to consider imbalanced swaps. For recursive partitioning, we typically do not want to allow ε imbalance for the early recursive splits since that will substantially constrain movement at later steps of the recursion. Instead, using ε multiplied by the ratio of the current number of recursive splits to the final number of recursive splits works reasonably well in practice.

4. EXPERIMENTS

Here we describe our experiments that are designed to answer the following questions:

- What is the impact of reducing fanout on query processing in a sharded database (Section 4.2.1)?
- How well does our algorithm perform in terms of quality and scalability compared to existing hypergraph partitioners (Sections 4.2.2 and 4.2.3)?
- How do various parameters of our algorithm contribute to its performance and final result (Section 4.2.4)?

4.1 Datasets

We use a collection of hypergraphs derived from large social networks and web graphs; see Table 1. We transform the input hypergraph into a bipartite graph, $G = (\mathcal{Q} \cup \mathcal{D}, E)$, as described in Section 1.

In addition, we use five large synthetically generated graphs that have similar characteristics as the Facebook friendship graph [16]. These generated graphs are a natural source of hypergraphs; in our storage sharding application, to render a profile-page of a Facebook user, one might want to fetch information about a user’s friends. Hence, every user of a social network serves both as query and as data in a bipartite graph.

In all experiments, isolated queries and queries of degree one (single-vertex hyperedges) are removed, since they do not contribute to the objective, having fanout equal to one in every partition.

4.2 Evaluation

We evaluate two versions of our algorithm, direct partitioning into k buckets, SHP- k , and recursive bisection with $\log_2 k$ levels, SHP-2. The algorithms are implemented in Java and SHP-2 is available at [1]. Both versions can be run in a single-machine environment using one or several threads executed in parallel, or in a distributed environment using a Giraph cluster. Throughout the section, we use imbalance ratio $\varepsilon = 0.05$.

Table 1: Properties of hypergraphs used in our experiments


hypergraph	source	$ Q $	$ D $	$ E $
email-Enron	[3]	25,481	36,692	356,451
soc-Epinions	[3]	31,149	75,879	479,645
web-Stanford	[3]	253,097	281,903	2,283,863
web-BerkStan	[3]	609,527	685,230	7,529,636
soc-Pokec	[3]	1,277,002	1,632,803	30,466,873
soc-LJ	[3]	3,392,317	4,847,571	68,077,638
FB-10M	[16]	32,296	32,770	10,099,740
FB-50M	[16]	152,263	154,551	49,998,426
FB-2B	[16]	6,063,442	6,153,846	2×10^9
FB-5B	[16]	15,150,402	15,376,099	5×10^9
FB-10B	[16]	30,302,615	40,361,708	10×10^9

4.2.1 Storage Sharding

Here we argue and experimentally demonstrate that fanout is a suitable objective function for our primary application, storage sharding. We refer the reader to the description of the Social Hash framework [29] for more detailed evaluation of the system and other applications of SHP.

In many applications, queries issue requests to multiple storage servers, and they do so in parallel. As such, the latency of a multi-get query is determined by the slowest request. By reducing fanout, the probability of encountering a request that is unexpectedly slower than the others is reduced, thus reducing the latency of the query. This is the fundamental argument for using fanout as the objective function for the assignment problem in the context of storage sharding. We ran a simple experiment to confirm our understanding of the relationship between fanout and latency. We issued trivial remote requests and measured (i) the latency of a single request (fanout = 1) and (ii) the latency of several requests sent in parallel (fanout > 1) (that is, the maximum over the latencies of single requests). Figure 4a shows the results of this experiment and illustrates the dependency between various percentiles of multi-get query latency and fanout of the query. The observed latencies match our expectations and indicate that reducing fanout is important for database sharding; for example, one can almost half the average latency by reducing fanout from 40 to 10.

There are several possible caveats to our analysis of the relationship between fanout and latency in the simplistic experiment. For example, reducing fanout generally increases the size of the largest request to a server, which could increase latency. With this in mind, we conduct a more realistic experiment with 40 servers storing a subset of the Facebook friendship graph. For the experiment, the data is stored in a memory-based, key-value store, and there is one data record per user. In order to shard the data, we minimize fanout using our SHP algorithm applied for the graph. We sample a live traffic pattern, and issued the same set of queries, while measuring fanout and latency of each query. The dependency between fanout and latency are shown in Figure 4b. Here the queries needed to issue requests to only 9.9 servers on average. Notice that we do not include measurements for fanout > 35, as there are very few such queries; this also explain several “bumps” in latency for queries with large fanout. The results demonstrate that decreasing fanout from 40 (corresponding to a “random” sharding) to 10 (“social” sharding) yields a $2x$ lower average latency for the queries, which agrees with the results of the earlier experiment.

**Figure 4:** Distribution of latency for multi-get queries with various fanout, where t is the average latency of a single call.

Finally, we mention that after we deployed storage sharding optimized with SHP to one of the graph databases at Facebook, containing thousands of storage servers, we found that measured latencies of queries decreased by over 50% on average, and CPU utilization also decreased by over 50%; see [29] for more details.

4.2.2 Quality

Next we compare the quality of our algorithm as measured by the optimized fanout with existing hypergraph partitioners. We identified the following tools for hypergraph partitioning whose implementations are publicly available: hMetis [20], Mondriaan [32], Parkway [31], PaToH [10], and Zoltan [14]. These are tools that can process hypergraphs and can optimize fanout (or the closely related sum of external degrees) as the objective function.

For a fair comparison, we set allowed imbalance $\varepsilon = 0.05$ and used default optimization flags for all partitioners. We also require that partitions can be computed with 10 hours without errors. We computed the fanout of partitions produced with these algorithms for various numbers of buckets, k , and found that hMetis and PaToH generally produced higher fanout than the other partitioners on our hypergraphs. So for clarity, we focus on results from SHP-2 and SHP- k , along with Mondriaan, Parkway, and Zoltan.

Figure 5 compares the fanout of partitions produced by these algorithms for various bucket count against the minimum fanout partition produced across all algorithms. No partitioner is consistently the best across all hypergraphs and bucket count. However, Zoltan and Mondriaan generally produce high quality partitions in all circumstances.


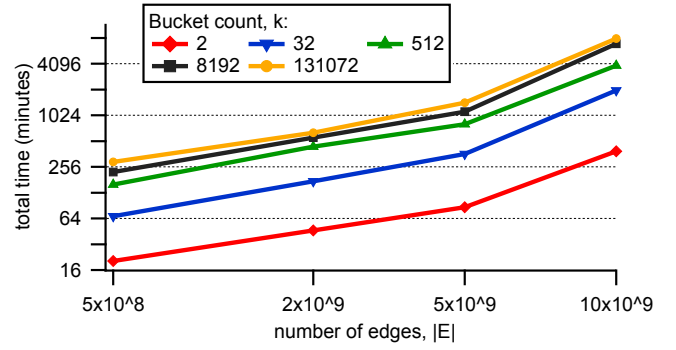



Figure 5: Fanout optimization of different single-machine partitioners across hypergraphs from Table 1 for $k \in \{2, 8, 32, 128, 512\}$.

SHP-2, SHP- k and to a lesser extent Parkway are more inconsistent. For example, both versions of SHP have around a 10 – 30% increased fanout, depending on the bucket count, over the minimum fanout on the web hypergraphs. On the other hand, SHP generally performs well on FB-10M and FB-50M, likely the hypergraphs closest to our storage sharding application. We also note a trade-off between quality and speed for recursive bisection compared to direct k -way optimization. The fanout from SHP-2 is typically, but not always, 5 – 10% larger than SHP- k .



(a) Total time of SHP-2 across four largest hypergraphs from Table 1 for various number of buckets and 4 machines.



(b) Run-time and total time of SHP-2 on FB-10B for various number buckets and machines in a cluster.

Figure 6: Run-time and total time (run-time of a single machine multiplied by the number of machines) of SHP-2 across largest hypergraphs from Table 1 in a distributed setting.

Because no partitioner consistently provides the lowest fanout, we conclude that using all partitioners and taking the lowest fanout across all partitions is an attractive option if possible. However, as we shall see in the next section, SHP-2 may be the only option for truly large hypergraphs due to its superior scalability.

4.2.3 Scalability

In this section, we evaluate SHP’s performance in a distributed setting using the four largest hypergraphs in Table 1. We use 4 machines for our experiments each having the same configuration: Intel(R) Xeon(R) CPU E5-2660 @ 2.20GHz with 144GB RAM.

First, we numerically verify the computational complexity estimates from Section 3.3. Figure 6a shows SHP-2’s total wall time as a function of the number of hyperedges across the four largest hypergraphs. Notice that the y-axis is on a log-scale in the figure. The data verifies that SHP-2’s computational complexity is $\mathcal{O}(\log k |E|)$, as predicted.

Next we analyze scalability of our approach with various number of worker machines in a cluster. Figure 6b(left) illustrates the run-time of SHP-2 using 4, 8, and 16 machines; Figure 6b(right) illustrates the total wall time using the same configuration. While there is a reduction in the run-time, the speedup is not proportional to the ratio of added machines. We explain this by an increased communication between the machines, which contributes to the performance of our algorithm.

Now we compare SHP’s performance to the two existing distributed partitioning packages, Parkway and Zoltan. The




Figure 7: Run-time in minutes of distributed hypergraph partitioners across hypergraphs from Table 1 for $k \in \{32, 512, 8192\}$. All tests run on 4 machines. Partitioners that fail process an instance or if their run-time exceeds 10 hours are shown with the maximum value.

comparison with *Zoltan* is particularly relevant since it provides generally lower fanout in the quality comparison.

Figure 7 shows the run-time of these partitioners in minutes across the hypergraphs and various bucket count. If the result was not computed within 10 hours without errors, we display the maximum value in the figure. *Parkway* only successfully ran on one of these graphs within the time allotted, because it runs out of memory on the other hypergraphs in the 4-machine setting. Similarly, *Zoltan* also failed to partition hypergraphs larger than *soc-LJ*. On the other hand, *SHP-k* ran on *FB-10B* for 32 buckets, and only *SHP-2* was able to successfully run on all tests.

Further, note that the x-axis is on a log-scale in Figure 7. So *SHP* can not only run on larger graphs with more buckets than *Zoltan* and *Parkway* on these 4 machines, the run-time is generally substantially faster. *SHP-2* finished partitioning every hypergraph in less than 5 hours, and for the hypergraphs on which *SHP-k* succeeded, it ran less than 8 hours.

While not observable in Figure 7, *Zoltan*'s run-time was largely independent of the bucket count, such that for 8192 buckets on *FB-50M* it was faster than *SHP-k*. This is a relatively rare case; typically *SHP-k*, despite having a run-time that scales linearly with bucket count, was faster in our experiments. While in all examples *Zoltan* was much slower than *SHP-2*, for a division of a small hypergraph into a very large number of buckets, *Zoltan* could conceivably be faster, as *SHP-2*'s run-time scales logarithmically with bucket count.




Figure 8: Fanout optimization with *SHP-2* on *soc-Pokec* as a function of p for different bucket counts.




Figure 9: Progress of fanout optimization with *SHP-k* for $p = 0.5$ and $p = 1.0$ on *soc-LJ* using $k = 8$.

4.2.4 Parameters of SHP

There are two parameters affecting Algorithm 1: the fanout probability and the number of refinement iterations. To investigate the effect of these parameters, we apply *SHP-2* for various values of $0 < p < 1$; see Figure 8 illustrating the resulting percentage reduction in (non-probabilistic) fanout on *soc-Pokec*. Values between $0.4 \leq p \leq 0.6$ tend to produce the lowest fanout, with $p = 0.5$ being a reasonable default choice for all tested values of bucket count k . The point $p = 1$ in the figure corresponds to optimizing fanout directly with *SHP-2*, and yields worse results than $p = 0.5$.

One explanation, as mentioned in Section 3.1, is that the local search is more likely to land in a local minimum with $p = 1$. This is illustrated in Figure 9 for *SHP-k*, where the number of moved vertices per iteration on *soc-LJ* is significantly lower for $p = 1$ than for $p = 0.5$ at earlier iterations. The number of moved vertices for $p = 0.5$ is below 0.1% after iteration 35; this number is below 0.01% after iteration 49. This behavior on *soc-Pokec* and *soc-LJ* for *SHP* was observed typical across many hypergraphs, and motivates our default parameters. We set $p = 0.5$ as the default

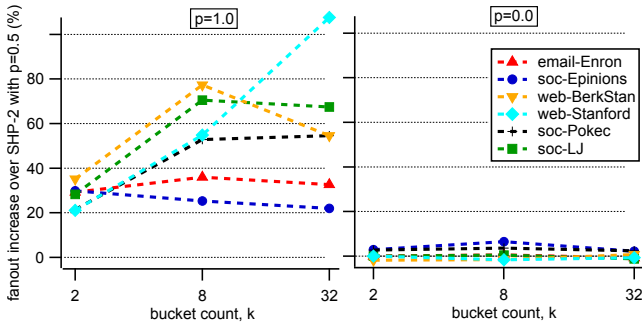


Figure 10: Probabilistic fanout with $p = 0.5$ vs direct fanout optimization ($p = 1.0$) and the clique-net objective ($p = 0.0$) for $k \in \{2, 8, 32\}$.

for p-fanout, 60 for the maximum number of refinement iterations of SHP-k, and 20 iterations per bisection for SHP-2.

Figure 10 quantifies the impact of using probabilistic fanout for SHP-2 across a variety of hypergraphs for 2, 8, and 32 buckets. The left-hand plot displays the substantial percentage increases in fanout caused by using direct fanout optimization over the $p = 0.5$ probabilistic fanout optimization. Similar behavior is seen with SHP-k, and these results demonstrate the importance of using probabilistic fanout for SHP. On average across these hypergraphs, direct fanout optimization would produce fanout values 45% larger than probabilistic fanout with $p = 0.5$.

The right-hand plot of Figure 10 compares p-fanout with the clique-net model defined in Section 3.1, where the clique-net objective is optimized using SHP-2. The comparison to $p = 0.5$ reveals that clique-net optimization is often worse, but typically similar, depending on the graph and the number of buckets. In practice, we suggest optimizing for both p-fanout and clique-net, as which surrogate objective performs better for fanout minimization depends on the graph.

5. DISCUSSION

Storage sharding for production systems has many additional practical challenges [29]. Two requirements that arise from these challenges are (i) *incremental updates* of an existing partition and (ii) *balance across multiple dimensions*.

(i) Incremental updates can be needed to avoid overhead from substantially changing a partition. Our algorithm simply adapts to incremental updates by initializing with a previous partition and running a local search. If a limited search moves too many data vertices, we can modify the move gain calculation to punish movement from the existing partition or artificially lower the movement probabilities returned via master in superstep four.

(ii) Basic k -way hypergraph partitioning balances the number of data vertices per bucket. Trivially, we can consider weighted data vertices, but additionally, a data vertex might have multiple associated dimensions (e.g., CPU cost, memory, disk resources etc.) that each require balance. In practice, we have found that requiring strict balance on many dimensions substantially harms solution quality. Instead, we favor a simple heuristic that produces $c \cdot k$ buckets for some $c > 1$ that have loose balance requirements on all but one dimension, and merges them into k buckets to satisfy load balance across all dimensions.

We stress that the storage sharding problem might have additional requirements that are not captured by our model. For example, one could introduce some replication by allowing data records to be present on multiple servers, as it is done in [11, 26]. However, that would bring in new operational complications (e.g., synchronization between the servers when a data record is modified), which are not always possible in an existing framework. Another potential extension of our model is to consider a better optimization goal that more accurately captures the relation between query latency and distribution of data records among servers. We recently observed that the query fanout is not the only objective affecting its latency; the size of a request to a server also plays a role. For example, a query with fanout = 2 that needs 100 data records can be answered faster if the two servers contain an even number of records, 50 and 50, in comparison with a distribution of 99 and 1. We leave a deeper investigation of this and other potential extensions of the model as an interesting future direction.

6. CONCLUSION AND FUTURE WORK

In this paper, we presented Social Hash Partitioner, SHP, a distributed hypergraph partitioner that can optimize p-fanout, as well as the clique-net objective among others, through local search, and scales to far larger hypergraphs than existing hypergraph partitioning packages. Because the careful design of SHP limits space, computational, and communication complexity, the applicability of our implementation to a hypergraph is only constrained by the number of Giraph machines available. We regularly apply SHP to hypergraphs that contain billions of vertices and hundreds of millions to even billions of hyperedges. These hypergraphs correspond to bipartite graphs with billions of vertices and trillions of edges.

Despite the improved scalability and simplicity of the algorithm, our experiments demonstrate that SHP achieves comparable solutions to both single-machine and distributed hypergraph partitioners. We note these results and conclusions might be different for other input hypergraphs (e.g. matrices from scientific computing, planar networks or meshes, etc.), and in cases where the hypergraph is small enough and solution quality is essential, running all available tools is recommended. While SHP occasionally produced the best solution in our experiments, other packages, especially *Zoltan* and *Mondriaan*, often claimed the lowest fanout.

Although our motivation for designing SHP is practical, it would be interesting to study fanout minimization from a theoretical point of view. For example, the classical balanced partitioning problem on unipartite graphs admits an $\mathcal{O}(\sqrt{\log k \log n})$ -approximation for $\varepsilon = 2$ [25] but it is unclear if a similar bound holds for our problem. Alternatively, it would be interesting to know whether there is an optimal algorithm for some classes of hypergraphs, or an algorithm that provably finds a correct solution for certain random hypergraphs (e.g., generated with a planted partition model). Finally, it would be interesting to understand when and how minimizing p-fanout speeds up algorithm convergence and improves solution quality over direct fanout minimization.

Acknowledgments. We thank H. Killapi for contributing to an earlier graph partitioner that evolved into SHP. We also thank M. Stumm for providing comments on this work.

7. ADDITIONAL AUTHORS

Alessandro Presta (Google, apresta@google.com) and Yaroslav Akhremtsev (Karlsruhe Institute of Technology, yaroslav.akhremtsev@kit.edu) – work done at Facebook.

8. REFERENCES

- [1] Social Hash Partitioner. <https://issues.apache.org/jira/browse/GIRAPH-1131>.
- [2] Apache Giraph. <http://giraph.apache.org/>.
- [3] Stanford large network dataset collection. <https://snap.stanford.edu/data>.
- [4] C. J. Alpert, L. W. Hagen, and A. B. Kahng. A hybrid multilevel/genetic approach for circuit partitioning. In *Asia Pacific Conference on Circuits and Systems*, pages 298–301. IEEE, 1996.
- [5] C. J. Alpert and A. B. Kahng. Recent directions in netlist partitioning: A survey. *Integration, the VLSI journal*, 19(1–2):1–81, 1995.
- [6] K. Andreev and H. Räcke. Balanced graph partitioning. *Theory of Computing Systems*, 39(6):929–939, 2006.
- [7] D. A. Bader, H. Meyerhenke, P. Sanders, and D. Wagner. Graph partitioning and graph clustering, 10th DIMACS implementation challenge workshop. *Contemporary Mathematics*, 588, 2013.
- [8] C.-E. Bichot and P. Siarry. *Graph partitioning*. John Wiley & Sons, 2013.
- [9] A. Buluç, H. Meyerhenke, I. Safro, P. Sanders, and C. Schulz. Recent advances in graph partitioning. In *Algorithm Engineering*, pages 117–158. Springer, 2016.
- [10] Ü. V. Çatalyürek and C. Aykanat. Hypergraph-partitioning-based decomposition for parallel sparse-matrix vector multiplication. *IEEE Transactions on Parallel and Distributed Systems*, 10(7):673–693, 1999.
- [11] C. Curino, E. Jones, Y. Zhang, and S. Madden. Schism: a workload-driven approach to database replication and partitioning. *VLDB Endowment*, 3(1-2):48–57, 2010.
- [12] J. Dean and L. A. Barroso. The tail at scale. *Communications of the ACM*, 56:74–80, 2013.
- [13] M. Deveci, K. Kaya, B. Uçar, and Ü. V. Çatalyürek. Hypergraph partitioning for multiple communication cost metrics: Model and methods. *Journal of Parallel and Distributed Computing*, 77:69–83, 2015.
- [14] K. D. Devine, E. G. Boman, R. T. Heaphy, R. H. Bisseling, and Ü. V. Çatalyürek. Parallel hypergraph partitioning for scientific computing. In *International Parallel & Distributed Processing Symposium*, pages 10–pp. IEEE, 2006.
- [15] L. Dhulipala, I. Kabiljo, B. Karrer, G. Ottaviano, S. Pupyrev, and A. Shalita. Compressing graphs and indexes with recursive graph bisection. In *International Conference on Knowledge Discovery and Data Mining*, pages 1535–1544. ACM, 2016.
- [16] S. Edunov, D. Logothetis, C. Wang, A. Ching, and M. Kabiljo. Darwini: Generating realistic large-scale social graphs. *arXiv:1610.00664*, 2016.
- [17] A. E. Feldmann and L. Foschini. Balanced partitions of trees and applications. *Algorithmica*, 71(2):354–376, 2015.
- [18] C. M. Fiduccia and R. M. Mattheyses. A linear-time heuristic for improving network partitions. In *19th Conference on Design Automation*, pages 175–181. IEEE, 1982.
- [19] L. Golab, M. Hadjieletheriou, H. Karloff, and B. Saha. Distributed data placement to minimize communication costs via graph partitioning. In *International Conference on Scientific and Statistical Database Management*, pages 20:1–20:12. ACM, 2014.
- [20] G. Karypis, R. Aggarwal, V. Kumar, and S. Shekhar. Multilevel hypergraph partitioning: applications in VLSI domain. *IEEE Transactions on Very Large Scale Integration Systems*, 7(1):69–79, 1999.
- [21] G. Karypis and V. Kumar. A fast and highly quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on Scientific Computing*, 20(1):359–392, 1999.
- [22] G. Karypis and V. Kumar. Multilevel k-way hypergraph partitioning. *VLSI design*, 11(3):285–300, 2000.
- [23] B. W. Kernighan and S. Lin. An efficient heuristic procedure for partitioning graphs. *Bell system technical journal*, 49(2):291–307, 1970.
- [24] T. Kiefer. *Allocation Strategies for Data-Oriented Architectures*. PhD thesis, Dresden, Technische Universität Dresden, 2016.
- [25] R. Krauthgamer, J. S. Naor, and R. Schwartz. Partitioning graphs into balanced components. In *Symposium on Discrete Algorithms*, pages 942–949. SIAM, 2009.
- [26] K. A. Kumar, A. Quamar, A. Deshpande, and S. Khuller. SWORD: workload-aware data placement and replica selection for cloud data management systems. *The VLDB Journal*, 23(6):845–870, 2014.
- [27] T. Leighton, F. Makedon, and S. Tragoudas. Approximation algorithms for VLSI partition problems. In *Circuits and Systems*, pages 2865–2868. IEEE, 1990.
- [28] R. O. Selvitopi, A. Turk, and C. Aykanat. Replicated partitioning for undirected hypergraphs. *Journal of Parallel and Distributed Computing*, 72(4):547–563, 2012.
- [29] A. Shalita, B. Karrer, I. Kabiljo, A. Sharma, A. Presta, A. Adcock, H. Killapi, and M. Stumm. Social Hash: An assignment framework for optimizing distributed systems operations on social networks. In *NSDI*, pages 455–468, 2016.
- [30] H. D. Simon and S.-H. Teng. How good is recursive bisection? *Journal on Scientific Computing*, 18(5):1436–1445, 1997.
- [31] A. Trifunović and W. J. Knottenbelt. Parallel multilevel algorithms for hypergraph partitioning. *Journal of Parallel and Distributed Computing*, 68(5):563–581, 2008.
- [32] B. Vastenhouw and R. H. Bisseling. A two-dimensional data distribution method for parallel sparse matrix-vector multiplication. *SIAM Review*, 47(1):67–95, 2005.
- [33] W. Yang, G. Wang, L. Ma, and S. Wu. A distributed algorithm for balanced hypergraph partitioning. In *Advances in Services Computing*, pages 477–490. Springer, 2016.