GPOP: A scalable cache- and memory-efficient framework for Graph Processing Over Partitions

Kartik Lakhotia, Sourav Pati, Rajgopal Kannan, Viktor Prasanna
Ming Hsieh Department of Electrical Engineering
University of Southern California, Los Angeles, USA
{klakhoti,spati,rajgopak,prasanna}@usc.edu

arXiv:1806.08092v2 [cs.DC] 1 Jan 2019

Abstract

The past decade has seen the development of many shared-memory graph processing frameworks intended to reduce the effort of developing high performance parallel applications. However many of these frameworks, based on Vertex-centric or Edge-centric paradigms suffer from several issues, such as poor cache utilization, irregular memory accesses, heavy use of synchronization primitives or theoretical inefficiency, that deteriorate over-all performance and scalability.

Recently, a cache and memory efficient partition-centric paradigm was proposed for computing PageRank [17]. In this paper, we generalize this approach to develop a novel Graph Processing Over Partitions (GPOP) framework that is cache-efficient, scalable and work-efficient. GPOP exploits the locality of partitioning to dramatically improve the cache performance of a variety of graph algorithms. It achieves high scalability by enabling completely lock and atomic free computation. GPOP's built-in analytical performance model enables it to use a hybrid of source and partition centric communication modes in a way that ensures work-efficiency each iteration while simultaneously boosting high bandwidth sequential memory accesses. Finally, the GPOP framework is designed with programmability in mind. It completely abstracts away underlying parallelism and programming model details from the user, and provides an easy to program set of APIs with the ability to selectively continue the active vertex set across iterations. Such functionality is useful for many graph algorithms but not intrinsically supported by the current frameworks.

We extensively evaluate the performance of GPOP for a variety of graph algorithms, using several large datasets. We observe that GPOP incurs upto 8.6× and 5.2× less L2 cache misses compared to Ligra and GraphMat, respectively. In terms of execution time, GPOP is upto 19× and 6.1× faster than Ligra and GraphMat, respectively.

1 Introduction

Real world problems arising in web and social networks, transportation networks, biological systems etc. are often modeled as graph computation problems. Applications in these domains generate huge amounts of data that require efficient large-scale graph processing. To this purpose, many distributed frameworks have been proposed to process very large graphs on clusters [7, 13, 19, 20]. However, because of the high communication overheads of distributed systems, even single threaded implementations of graph algorithms have been shown to outperform many such frameworks running on several machines [23].

The growth in DDR capacity also allows large graphs to fit in the main memory of a single server. Consequently, many frameworks have been developed for high performance graph analytics on multicore platforms [12, 26–28, 31, 34]. However, multi-threaded graph algorithms may incur race conditions and hence, require expensive synchronization (atomics or locks) primitives that can significantly decrease performance and scalability. Furthermore, graph computations are characterized by large communication volume and irregular access patterns that make it challenging to efficiently utilize the resources even on a single machine [**?**].

Recent advances in hardware technologies offer potentially new avenues to accelerate graph analytics, in particular, new memory technologies, such as Hybrid Memory Cube (HMC) and scratchpad caches. However, many graph analytics frameworks are based on the conventional push-pull Vertex-centric processing paradigm [5, 12, 28, 34], which allows every thread to access and update data of arbitrary vertices in the graph. Without significant pre-processing, this leads to unpredictable and fine-grained random memory accesses, thereby decreasing the utility of the wide memory buses and deterministic caching features offered by these new architectures. Some frameworks and application specific programs[17, 27, 36, 39] have adopted optimized edge-centric programming models that improve access locality and reduce synchronization overhead. However, these programming models require touching all or a large fraction of edges of the graph in each iteration, and are not work optimal for algorithms with dynamic active vertex sets, such as bfs, seeded random walk etc. A work inefficient implementation can significantly underperform an efficient serial algorithm if the useful work done per iteration is very small.

In this paper, we describe GPOP: a Graph Processing Over Partitions framework [1] that enhances cache-efficiency and memory performance to efficiently execute a variety of graph algorithms. GPOP executes an algorithm as bulk synchronous steps of message exchange between vertex subsets called partitions (scatter), and message processing (gather).

[1] https://github.com/souravpati/GPOP

GPOP also utilizes a hierarchical active list structure to ensure that work done in each iteration is proportional to the number of active edges along with a novel dual communication mode technique that enables completely sequential memory accesses while scattering messages from high workload partitions. Thus, GPOP extracts maximum performance from the memory hierarchy while guaranteeing theoretically efficient implementations.

GPOP provides exclusive ownership of a partition to one thread to avoid unnecessary synchronization and achieve good scalability. We utilize this property to develop an easy to program set of APIs, that completely abstract away the underlying parallelism from the user and ensure correctness of the parallel implementation without requiring the user to lock or atomically update the vertex data. This is in stark contrast with many popular frameworks [12, 28, 34] where the burden of ensuring correctness of a parallel program is left to the user. GPOP also facilitates flexible manipulation and *selective continuity* of the active vertex set across iterations, which is not possible with existing frameworks.

The major contributions of our work are as follows:

1. We propose the Graph Processing Over Partitions (GPOP) framework with novel optimizations that (a) improves cache performance, (b) achieves high DRAM bandwidth, (c) minimizes synchronization and, (d) ensures work efficiency of a given algorithm.
2. GPOP provides an easy to program set of APIs allowing selective continuity in frontiers across iterations. This functionality is essential for many algorithms such as Nibble, Heat Kernel PageRank [29] etc., but is not supported intrinsically by the current frameworks.
3. GPOP ensures parallel correctness without the use of locks or atomics and demonstrates good scalability with upto 17.9× speedup on BFS using 36 threads, over a sequential implementation.
4. We evaluate the performance of GPOP by executing a variety of graph algorithms on large real-world and synthetic datasets. We show that GPOP incurs radically less cache misses and achieves upto 19× and 6.1× speedup over Ligra and GraphMat, respectively.

2 Graph Processing Challenges

Typically, graph applications execute iteratively, with a *frontier* of active vertices in each iteration, until a halting criteria is met. The performance of a graph analytics framework depends on the cache efficiency, DRAM communication, use of synchronization primitives and theoretical efficiency of its underlying programming model. We say that a parallel implementation is *theoretically efficient* if the amount of work done in an iteration is proportional to the number of outgoing edges from active vertices i.e active edges (for single iteration execution and some applications, such as BFS, this is equivalent to being work-efficient). Note that theoretical efficiency implies the work done is optimal for the given algorithm, within constant factors. In this section, we illustrate several challenges faced by current frameworks, in regard with the aforementioned aspects.

Vertex-centric (VC) programming is one of the most popular approaches used in graph processing [20]. Shared-memory frameworks [12, 28, 34] typically execute an iteration using either (ref. Algorithm 1)

• *Push* VC – where the set V_a of active vertices (*frontier*) is divided amongst worker threads who modify the attributes of their out-neighbors. In general, synchronization primitives such as atomics are required to ensure correctness of concurrent threaded access to common out-neighbors; this can limit the scalability of the program. While atomics can enable limited number of simple actions (compare and swap, fetch and add), applications performing a complex set of operations may require more expensive *locks* on vertices.

• *Pull* VC – where threads traverse the in-edges of vertices and modify their private attributes, requiring no synchronization. However, since edges from active and inactive neighbors are interleaved, all edges need to be probed resulting in a theoretically inefficient implementation and large runtime even if only few vertices are active.

Algorithm 1 Vertex-centric push and pull methods

$N_i(v)/N_o(v)$: in/out neighbors of vertex v
$val[v]$: attribute of v (eg. distance in SSSP, parent in BFS)

```
1: function PUSH(G(V, E), V_a)
2:     for v ∈ V_a do in parallel
3:         for all u ∈ N_o(v) do
4:             val[u] = updateAtomic(val[v], val[u])
5: function PULL(G(V, E))
6:     for v ∈ V do in parallel
7:         for all u ∈ N_i(v) do
8:             val[v] = update(val[v], val[u])
```

Most frameworks store the graph in a Compressed Sparse format (CSR or CSC) [**?**] which allows efficient sequential access to edges of a given vertex. However, probing neighbor vertex values generates fine-grained random accesses as neighbors are scattered. For large graphs, such accesses increase cache misses and DRAM traffic, becoming the bottleneck in graph processing. Figure 1 is an illustrative example. Here, > 75% of the total DRAM traffic in a PageRank iteration is generated by the random vertex value accesses.

To improve locality of memory accesses, Edge-Centric (EC) programming and blocking techniques [8, 27, 39] are also used. In EC, vertex values are written on edges in a semi-sorted manner (analogous to bucket sort) to increase spatial locality and cache-line utilization. Blocking restricts the range of vertices accessed by each thread and thus, improves

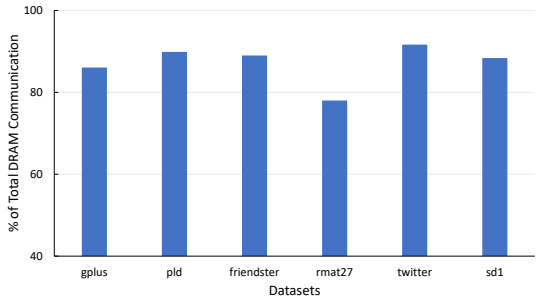

Figure 1. DRAM Traffic generated by random accesses to vertex data in PageRank

cache hit ratio. While such methods improve cache performance, they still incur coarse-grained random accesses limiting the bandwidth of CPU-DRAM communication. Recently, a Partition-Centric Programming Methodology (PCPM) [17] for PageRank computation was proposed that avoids random DRAM accesses and sustains high bandwidth. However, all of these approaches [8, 17, 27, 39] iterate over all vertices and edges in the graph, making them theoretically inefficient and extremely expensive for executing iterations with very few active vertices.

3 Graph Processing Over Partitions

In this paper, we describe GPOP: a Graph processing framework that enhances cache and memory performance, while *guaranteeing* theoretically efficient execution of a variety of graph algorithms. GPOP can be viewed as a generalization of partition based Pagerank [17]. Dividing the vertex set into *cacheable* disjoint partitions, GPOP implements each iteration of an algorithm in 2 phases:

1. Inter-partition communication (*Scatter*) → In this phase, the out edges of active vertices in a partition are streamed in and vertex data is communicated to other partitions in the form of messages (fig. 2).
2. Intra-partition vertex data updates (*Gather*) → In this phase, incoming messages of a partition are streamed in and the vertex data is updated as per a user-defined function. This phase also generates a list of vertices that become active and will communicate in the next iteration.

GPOP processes multiple partitions in parallel during both phases. For Scatter, each partition is provided a distinct memory space (called *bin*) to write its outgoing messages. During Gather, a single thread has exclusive access to all vertices in the partitions it processes. Thus, GPOP can communicate and compute vertex data without using synchronization primitives, such as locks or atomics.

Each thread in GPOP processes one partition at a time. Within each partition (in either phase), vertex data is reused while edges and messages are only accessed once. GPOP ensures high *temporal locality* since a partition's vertex data is designed to fit in the private cache of a core. At the same

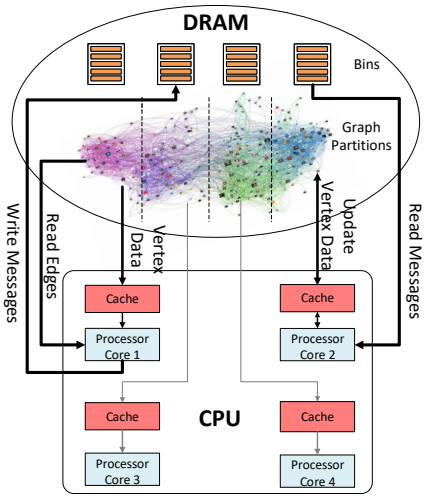

Figure 2. Conceptual demonstration of two phases in GPOP - Core 1 is writing messages in Scatter phase and Core 2 is processing them in Gather phase. Note that in the actual implementation, all threads are in the same phase at any given instant of time.

time, GPOP enjoys high *spatial locality* of message reads and writes by storing them in consecutive locations in bins. All threads in GPOP synchronize at the end of each Scatter and Gather phase to ensure that all partitions have received and processed all incoming messages before they get overwritten in the next iteration. Algorithm 2 describes the execution flow for an iteration of a typical graph application in GPOP.

GPOP chooses between two different scatter modes based on an analytical model that evaluates the tradeoff in work done per active edge versus maximizing main memory performance. In particular, GPOP achieves completely sequential DRAM accesses for scattering partitions with large number of active edges while slightly increasing the work. However, in all cases work done is always within a constant predetermined factor of the number of active edges, thus guaranteeing theoretical efficiency.

3.1 GPOP System Optimizations

GPOP decreases DRAM communication volume, improves cache-efficiency and spatial locality while ensuring theoretical efficiency of the underlying implementation through careful system optimizations, as detailed below.

A. Graph Partitioning: GPOP uses a lightweight index scheme for partitioning vertices into k disjoint equisized sets[2]. Vertex set V^p of partition p consists of all vertices with indices $\in [p\frac{V}{k}, (p+1)\frac{V}{k})$ while (E_a^p) E^p denotes the set of (active) out-edges from V^p. k is chosen to be the smallest

[2]Graph partitioning algorithms such as [6, 18] can optimally pack out-neighbors of vertices into partitions, further reducing the number of (partition-centric) messages exchanged in GPOP. However there is a trade-off between pre-processing time and the speedup obtained. To demonstrate the fundamental performance improvements delivered by GPOP, we downplay partitioning effects by using only the lightweight scheme described.

number satisfying *both* the following performance metrics 1) *Cache-efficiency*: We make partitions cacheable - the data for $q = V/k$ vertices should fit in the largest private cache (L2 cache in Xeon servers). 2) *Load balance*: there is sufficient parallelism to avoid threads being completely idle. For this purpose, we set $k \geq 4t$, where t is the number of threads. Having more partitions than the number of threads also assists dynamic load balancing at runtime, especially in algorithms such as BFS, where workload changes every iteration.

B. Message Generation and Communication: GPOP drastically reduces total communication volume during Scatter by aggregating messages. (Active) vertex v generates ${}^{p}d_v$ messages (where ${}^{p}d_v$ denotes its out-partition-degree) - exactly one message for each of its neighboring *partitions* (containing out-neighbors of v). We use $r^p = (\sum_{i \in p} {}^{p}d_v)/E^p$ to denote the average aggregation factor of partition p. A message consists of source vertex data and list of all out-neighbors in the neighboring partition. For example, in BFS a vertex can send its visited status or its index (to compute parents in BFS tree) as the vertex data. These messages are stored within distinct memory spaces called *bins*, allotted to each partition-pair. Specifically, $bin[p][q]$ is used to store all messages from partition p to partition q within two arrays - *data_bin[]* and *id_bin[]* - for storing vertex data and list of out-neighbors, respectively. Each *bin* is pre-allocated the maximum amount of space required to store all potential messages that can be exchanged between the corresponding partition-pair.

Algorithm 2 An Iteration of a Graph algorithm in GPOP

```
    ^pN_o(v) → set of neighboring partitions of v
1:  for p ∈ sPartList do in parallel                    ▷ Scatter
2:      for all v ∈ V_a^p do
3:          for all q ∈ N_o^p[v] do
4:              insert {val[v], V^q ∩ N_o(v)} in bin[p][q]
5:      for all v ∈ V_a^p do                            ▷ reInit Frontier
6:          val[v] = update1(val[v])
7:          if !cond1(val[v]) then  V_a^p = V_a^p \ v
8:  __synchronize()__
9:  for p ∈ gPartList do in parallel                    ▷ Gather
10:     for all msg ∈ bin[:][p] do
11:         for all v ∈ msg.id do
12:             val[v] = acc(msg.val, val[v])
13:             V_a^p = V_a^p ∪ v
14:     for all v ∈ V_a^p do                            ▷ filter Frontier
15:         val[v] = update2(val[v])
16:         if !cond2(val[v]) then  V_a^p = V_a^p \ v
17: __synchronize()__
```

Each partition p is scattered by one thread at a time with exclusive access to the corresponding row of bins $bin[p][:]$, using one of two communication modes:

Source-centric (SC) mode: When the number of active edges are small, the thread reads and processes only active vertices in p (lines 2-4, Alg. 2). 1) (Aggregated) messages from a given active vertex are written (in destination partition order) into the corresponding $bin[p][:]$ before processing the next vertex. 2) Successive messages to any partition q from active vertices in p are written to contiguous addresses in $bin[p][q]$, labeled as *insertion points* of bins. The small number of partitions ensures that *all* insertion points used by a thread are likely to fit in the cache, thus enabling efficient data packing in the cache lines with increased spatial locality.

The SC mode is optimal in terms of work done ($\approx r^p E_a^p$ messages in expectation), however, a thread will switch partitions (bins) being written into which can negatively impact the sustained memory bandwidth and overall performance. This motivates an alternate communication mode for GPOP.

Partition-centric (PC) mode: All vertices in V^p are scattered. To ensure communication-efficiency, (aggregated) messages from V^p are scattered in partition-centric order, i.e messages are *not* interleaved between partitions (lines 2-4, Alg. 3). Here we use the Partition-Node bipartite Graph (PNG) data layout (proposed in [17]) to Scatter. Using this layout, a thread generates and writes all messages to $bin[p][q]$ before moving to $bin[p][q+1]$, thereby ensuring completely sequential DRAM accesses and high sustained memory bandwidth.

PC mode requires all out-edges of a partition to be traversed. Consequently, the order in which messages are written always stays the same and adjacency information written once can be *reused* across multiple iterations. Messages to $bin[p][q]$ only need to contain vertex data from p, written in the same locations of $data_bin[]$ every iteration, while neighbor identifiers are present in a separate $id_bin[]$ written only once during pre-processing, thereby dramatically reducing DRAM communication volume.

Algorithm 3 GPOP Scatter in Partition-Centric mode

```
    N_i^p(q) → vertices in p that have out-neighbors in q
1:  for p ∈ sPartList do in parallel                    ▷ Scatter
2:      for all q ∈ P do
3:          for all v ∈ N_i^p(q) do
4:              insert val[v] in bin[p][q]
```

The PC and SC modes trade off DRAM bandwidth efficiency for slight increase in work (redundant messages). To achieve the best of both worlds, GPOP analytically evaluates the processing cost and chooses the lower cost mode *individually* for each partition *each* iteration (Sec.3.2).

Finally, during Gather, the thread for partition q *exclusively* reads all messages from $bin[:][q]$. These reads generate completely sequential memory accesses (in both modes). All messages in $bin[:][q]$ update vertex data of partition q (which is cached) providing good temporal locality. Thus, GPOP enables *atomic* and *lock-free* communication and computation

in both phases with improved cache performance. Many applications also require *conditional* reinitialization of active frontier in each iteration (parallel Nibble, Heat Kernel PageRank [29] etc.). GPOP intrinsically supports this functionality. After processing messages in either phase, user functions update (cached) data of active vertices and re-evaluate their active status (lines 5-7, 14-16 in alg. 2).

C. 2-level Active List: If Gather is unaware of the bins that are empty (not written into at all in the preceding Scatter phase), it will probe all the bins performing at least $\theta(k^2)$ work in the process. Theoretically, probing all partitions can make the implementation inefficient since $k = \theta(V)$ (partition size is capped by cache size). Practically, even if k is small (in thousands for very large graphs), k^2 can be large (in millions). This can heavily deteriorate performance, especially when the useful work done per iteration is very small, for instance, in the Nibble algorithm [30].

We propose a hierarchical list data structure to avoid these performance inefficiencies. At the top level, we implement a global list, called *gPartList*, that enumerates all partitions (columns of bins) which have received at least one incoming message in the current iteration. At the second level, we implement a local list for each partition, called *binPartList*, that enumerates all the bins of a given partition which have received at least one incoming message. Thus, when a partition p sends a message to partition q during scatter phase, it inserts (1) the message into $bin[p][q]$, (2) q into *gPartList* and (3) p into $binPartList[q]$.

During Gather, the *gPartList* serves as a queue of *tasks* to be executed in parallel. The thread allocated to partition q processes all $bin[:][q]$ listed in $binPartList[q]$. Finally, during Gather, GPOP also 1) Populates another list (*sPartList*) enumerating all partitions with at least one active vertex and 2) For each partition p, creates a frontier V_a^p of vertices that will be active in the next iteration.

3.2 Performance Modeling

Next we describe GPOP's analytical performance model for choosing the optimal communication mode for each partition. Let s_i denote the size of a vertex index and $s_v = sizeof(val(v))$ denote the size of a message data element (vertex value), respectively (both assumed 4B for the applications in this paper). Then the total communication volume generated by partition p in SC mode is given by:

$$(V_a^p + E_a^p)s_i + 2(r^p E_a^p s_v + E_a^p s_i)$$

The first term represents Scatter phase traffic for reading V_a^p offsets (CSR) and E_a^p edges. These generate messages containing an expected $r^p E_a^p$ data values and E_a^p destination IDs (second term) - which are exactly read by the succeeding Gather phase (factor of 2). Note that it is possible to exactly count messages in SC but we use r^p as a faster approximation.

The PC mode reads $r^p E^p$ edges and k offsets from the PNG layout and writes $r^p E^p$ messages containing only vertex data. The Gather phase, however, reads both the vertex data and the destination IDs from these messages. Total DRAM communication for PC mode is given by:

$$r^p E^p s_i + k s_i + 2rE^p s_v + E^p s_i = E^p((r^p + 1)s_i + 2rs_v) + k s_i$$

Since graph algorithms are mostly memory-bound, we compute the execution time as a ratio of communication volume and bandwidth (BW_{SC} or BW_{PC}). GPOP chooses PC mode for scattering a partition only if:

$$\frac{E^p((r^p + 1)s_i + 2rs_v) + k s_i}{BW_{PC}} \leq \frac{2r^p E_a^p s_v + 3E_a^p s_i + V_a^p s_i}{BW_{SC}} \quad (1)$$

The ratio $\frac{BW_{PC}}{BW_{SC}}$ is a user configurable parameter which is set to 2 by default. Note that although PC mode generates more messages, it is selected only if it does work within a constant factor of optimal . Also note that V_a^p (E_a^p) for a partition can increase or decrease *each iteration* (depending on the algorithm). By carefully evaluating the communication mode *individually* for each partition, each iteration, GPOP generates work efficient implementations of graph algorithms that achieve high bandwidth for DRAM communication.

4 GPOP Programming Interface

We carefully assess the program structure in GPOP to construct the right set of APIs that (a) hide all the details of parallel programming model from the user and allow an application to be described algorithmically, (b) do not restrict the functionality of the framework and provide enough expressive power to implement a variety of applications.

In each iteration comprising a Scatter and a Gather phase, the GPOP back-end calls four main user-defined functions that steer application development:

1. **`scatterFunc(node)`**: Called for all (active) vertices during (SC) Scatter; should return an (algorithm-specific) value to propagate (write in bins) to neighbors of `node` (Eg. alg. 4).
2. **`initFunc(node)`**: Applied to all active vertices at the end of Scatter (Alg. 2: lines 5-7). It allows conditional continuation of a vertex in active frontier i.e. a vertex *remains* active for the *next* iteration if this function returns `true`. A specific example is Alg. 4: lines 4-6, where, if $PR[node] \geq 2\epsilon * deg[node]$, then the vertex stays active, irrespective of whether it gets updated or not during Gather. The programmer can also use `initFunc()` to *update* vertex data before Gather begins. In Alg. 4, for example, only half the *PR* of a vertex is retained after scattering amongst its out-neighbors.

To the best of our knowledge, existing frameworks put the onus of implementing such selective continuity in the frontier, on the programmer, thus increasing program complexity. In contrast, GPOP intrinsically provides such functionality that appreciably reduces programmer effort, both in terms of length and complexity of the code.

3. **`gatherFunc(val,node)`**: For all messages read during Gather phase, this function is called with the data value and destination vertex of the message as input arguments. The user can accordingly update the vertex data without

worrying about atomicity or locking. For example, in parallel Nibble (Alg. 4), the `gatherFunc()` adds the input value to the current probability of the node. The concerned vertex is marked active if the function returns `true`.

4. **`filterFunc(node)`**: This function applies to all active vertices in the preliminary frontier built after message processing and can be used to filter out (deactivate) some vertices. For example, in parallel Nibble, `filterFunc()` performs the threshold check on all active vertices. The semantics of this function are similar to `initFunc`. It can also be used to update the final aggregated vertex values computed by `gatherFunc()`. For instance, in PageRank (Alg. 6), we use this function to apply the damping factor to each vertex.

For the case of weighted graphs, the user can define how the weight should be applied to the source vertex' value, using the `applyWeight()` function.

5 Applications

In this paper, we evaluate the performance of GPOP on five key applications. Algorithms 4-7 illustrate the respective user codes in GPOP. Program initialization and iterative calls to Scatter-Gather execution are similar for all applications and therefore only illustrated for Alg. 4 (lines 13-17).

Nibble (Alg. 4) - Computes the probability distribution of a seeded random walk. Nibble is often used for strongly local clustering [29, 30], where it is *mandatory* for the parallel implementation to be work-optimal.

Algorithm 4 Parallel Nibble implementation in GPOP

```
1:  struct UF{
2:  procedure scatterFunc(node)
3:      return PR[node]/(2 * deg[node]);
4:  procedure initFunc(node)
5:      PR[node] = PR[node]/2;                          //update1//
6:      return (PR[node] ≥ ϵ * deg[node]);              //cond1//
7:  procedure gatherFunc(val, node)
8:      PR[node]+ = val;
9:      return true;
10: procedure filterFunc(node)
11:     return (PR[node] ≥ ϵ * deg[node]);              //cond2//
12: };
13: function nibble(G, v_s)                             ▷ v_s → start node
14:     loadFrontier(G, v_s);
15:     PR[] = 0; PR[v_s] = 1;
16:     while (G.frontierSize > 0 and MAX_ITER) do
17:         scatter_and_gather(&G, UF);
```

Breadth-First Search (BFS) (Alg. 5) - Used for rooted graph traversal or search (second kernel in Graph500 [25]). The `scatterFunc` returns the ID of the active node; the `initFunc` resets the frontier ; and the `gatherFunc` updates the parent of a vertex if it was not already visited.

Algorithm 5 Breadth First Search in GPOP

```
    v_s → root node
    parent[] = {−1}; parent[v_s] = v_s;
1:  procedure scatterFunc(node)
2:      return node;
3:  procedure initFunc(node)
4:      return false;
5:  procedure gatherFunc(val, node)
6:      if parent[node] < 0 then
7:          parent[node] = val;
8:          return true;
9:      return false;
10: procedure filterFunc(node)
11:     return true;
```

PageRank (Alg. 6) - A node ranking algorithm; representative of the crucial SpMV multiplication kernel which is used in many scientific and engineering applications [2, 32**?**]. `scatterFunc` returns the degree normalized rank of the input node ; `gatherFunc` simply accumulates the incoming value into the PageRank of node. Since all vertices are always active, `initFunc` and `filterFunc` are simply used to reinitialize PageRank values and apply damping factor, respectively.

Algorithm 6 PageRank computation in GPOP

```
    pageRank[] = {1/|V|}
1:  procedure scatterFunc(node)
2:      return pageRank[node]/deg[node];
3:  procedure initFunc(node)
4:      pageRank[node] = 0;
5:      return true;
6:  procedure gatherFunc(val, node)
7:      pageRank[node] += val;
8:      return true;
9:  procedure filterFunc(node)
10:     pageRank[node] = (1−d)/|V| + d * pageRank[node] ;
11:     return true;
```

Single Source Shortest Path (SSSP) (Alg. 7) - Finds the shortest distance to all vertices in a graph from a source vertex (third kernel in Graph500). The `scatterFunc` returns the distance of the active node; `initFunc` resets the frontier and `gatherFunc` updates the distance of the node if a shorter path is found. The additional `applyWeight()` function specifies that the sum of edge weight (`wt`) and source node distance (`val`) should be propagated to the neighbors.

Label Propagation - Iteratively, (a seed set of) vertices propagate labels and update their neighbors. Either all labeled vertices or vertices whose labels changed in current iteration may become active in the next iteration. Label propagation can be used for learning, community detection, connected

Algorithm 7 SSSP computation in GPOP

```
distance[] = {∞}; distance[v_s] = 0;
1:  procedure scatterFunc(node)
2:      return distance[node];
3:  procedure initFunc(node)
4:      return false;
5:  procedure gatherFunc(val, node)
6:      if val < distance[node] then
7:          distance[node] = val;
8:          return true;
9:      return false;
10: procedure filterFunc(node)
11:     return true;
12: procedure applyWeight(val, wt)
13:     return (val + wt);
```

components, location detection etc [10, 38]. In this paper, we evaluate the performance of connected components algorithm for which the `gatherFunc()` selects the minimum label and returns `true` if the label of the node changes. For brevity, we do not show the pseudocode for this application.

The programs (alg.4-7) show that diverse applications can be implemented in GPOP with few lines of code and minuscule changes in the user-defined functions.

6 Evaluation

6.1 Experimental Setup

We conduct experiments on the following 2 machines

1. *M*1: a dual-socket Broadwell server with two 18-core Intel Xeon E5-2695 v4 processors running Ubuntu 16.04.
2. *M*2: a dual-socket Ivybridge server with two 8-core Intel Xeon E5-2650 v2 processors running Ubuntu 14.04.

The important properties of these machines are listed in table 1. Memory bandwidth is measured using STREAM benchmark [22]. We compile GPOP using G++ 6.3.0 and parallelize it using OpenMP version 4.5. Partition size in GPOP is set to 256 KB (same as L2 cache). Unless explicitly stated otherwise, the results depict performance using 36 threads on *M*1 and 16 threads on *M*2.

Table 1. System Characteristics

		*M*1	*M*2
Socket	L3 cache	45 MB	25 MB
Core	L1d cache	32 KB	32 KB
	L2 cache	256 KB	256 KB
Main Memory	size	1 TB	128 GB
	Copy BW	55.1 GBps	43.7 GBps
	Add BW	64.2 GBps	48.8 GBps

We observe that the scaling and performance of GPOP relative to other frameworks is better on *M*2 as compared to *M*1. This is due to the higher $\frac{BW}{\#cores}$ ratio of *M*2 that boosts GPOP, and the smaller *L*3 cache in *M*2 that hurts random access performance of the baseline frameworks. For brevity, we only report the results for *M*1 in this paper, with the exception of L2 cache misses. Due to insufficient permissions on *M*1, the L2 cache statistics could only be collected on *M*2. Cache performance is measured using Intel Performance Counter Monitor [33].

We use large real world and synthetic graph datasets to evaluate the performance of GPOP. Table 2 lists the real-world datasets used in the experiments. *Soclj*, *Google+*, *Twitter* and *Friendster* are social networks ; *Pld* and *Sd1* are hyperlink graphs extracted by web crawlers. The synthetic graphs used in this paper are generated by RMAT graph generator with default settings (scale-free graphs) and degree 16.

Table 2. Graph Datasets

Dataset	Description	# Vertices	# Edges
soclj [15]	LiveJournal	4.84 M	68.99 M
gplus [11]	Google+	28.94 M	462.99 M
sd1 [24]	Subdomain graph	94.95 M	1937.49 M
pld [24]	Pay-Level-Domain	42.89 M	623.06 M
friendster [15]	social network	68.35 M	2568.14 M
twitter [15, 16]	follower network	61.58 M	1468.36 M
rmat<*n*> [9]	synthetic graph	2^n M	$16 * 2^n$ M

6.2 Results

6.2.1 Comparison with Baselines

We choose to compare the performance of GPOP against 2 most popular frameworks - Ligra [28] and GraphMat [31] , which are representatives of Vertex-centric push-pull methodology and Scatter-Gather based Sparse Matrix computation, respectively. GraphMat only supports graphs with $< 2^{31}$ edges and cannot process the *Friendster* graph.

To explore the overall benefit of dual communication modes in GPOP (section 3), we also measure the runtime with only SC mode scatter (GPOP_SC). For the parallel Nibble algorithm, the GraphMat implementation is not available and we only compare against Ligra implementation given in [29].

Execution Time: Figure 3 provides the normalized runtime of the five graph algorithms discussed in section 5 implemented using GPOP and baseline frameworks. GPOP consistently outperforms GraphMat for all the algorithms by $2\times - 6.1\times$ speedup on large datasets. Except BFS, GPOP executes all other algorithms faster than Ligra achieving upto $19\times$ speedup for PageRank.

GPOP executes PageRank and Label Propagation dramatically faster than both Ligra and GraphMat. The speedup is more substantial for large graphs such as *Friendster* where the cache and memory bandwidth optimizations of GPOP become extremely crucial for performance. Further, in PageRank, all vertices are always active allowing GPOP to use the high performance PC mode throughout the course of the

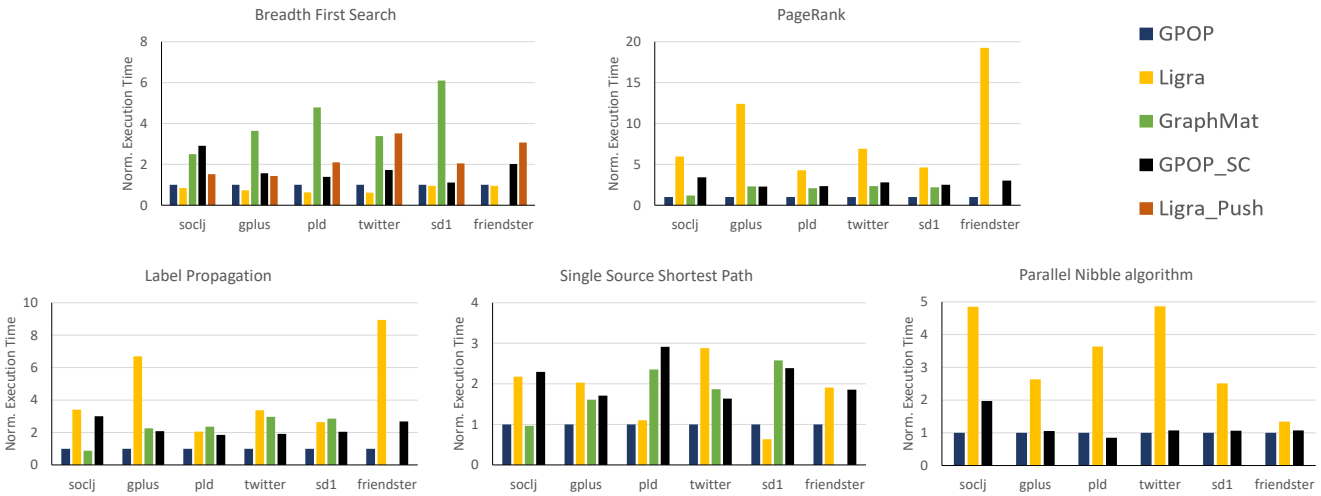

Figure 3. Comparison of Execution time (normalized with GPOP runtime) for various Graph algorithms

algorithm. For both PageRank and Label Propagation, GPOP is $1.8\times - 3.4\times$ faster than GPOP_SC.

In case of BFS, the pull direction optimization in Ligra enables early termination of loops (iterating on adjacency lists), drastically reducing the number of edges traversed. GPOP and GraphMat do not support pull direction methodology and traverse all active edges. However, GPOP is still $0.61\times - 0.95\times$ as fast as Ligra. For comparison, we also show that BFS in Ligra without direction optimization (Ligra_Push) is upto 3.06× slower than GPOP. GraphMat iterations are not theoretically efficient and do $O(V)$ work in traversing the frontier. In comparison, BFS in GPOP is work-efficient providing 4× and 6× speedup over GraphMat for *Sd1* and *Pld* graphs, respectively.

For SSSP, GPOP achieves upto 2.9× and 2.6× speedup over Ligra and GraphMat, respectively. Ligra converges faster (2.7× less edges traversed in *sd1*) because of *asynchronous* updates i.e. any change in a vertex' distance is immediately visible to its neighbors in the same iteration. Conversely, GPOP and GraphMat implement Jacobi-style Bellman Ford [**?**] in which distance updates are only visible in the next iteration. Therefore, Ligra execution time is comparable to GPOP for *Sd1* and *Pld* graphs, respectively. *Asynchronous* updates can also be enabled in GPOP by scattering the pointer to vertex values instead of value itself. The Gather phase will chase the pointers to obtain the value of source vertex. This is a trade-off between cache efficiency and quick convergence and the programmer has to choose wisely.

In parallel Nibble, the algorithm explores very few vertices in local neighborhood of the seed set, enabling good cache performance even in Ligra. The frontiers are also small and GPOP is unable to utilize the PC mode for high memory performance. Consequently, unlike BFS and SSSP, both GPOP and GPOP_SC provide similar performance for Nibble algorithm. Since all threads are updating a small set of vertices, GPOP benefits by the atomic free updates. Furthermore, the Ligra implementation [29] uses hash-based sparse data structures and suffers from hash collisions. Instead, our implementation uses an array to store probabilities but still achieves theoretical efficiency by amortizing the cost of reading the graph and initialization, for multiple runs of Nibble.

Cache Performance: We report the absolute L2 cache misses of GPOP and the baselines, for PageRank, Label Propagation and SSSP (table 3). Because of the cache optimizations explained in section 3, GPOP incurs dramatically lesser cache misses than GraphMat and Ligra for almost all datasets and algorithms. For PageRank, GPOP exhibits an average reduction of 8.6× and 5.8× in L2 misses, compared to Ligra and GraphMat, respectively. In case of Label Propagation, number of L2 misses in GPOP are 2.8× and 1.5× fewer on average than Ligra and GraphMat, respectively. GraphMat has better cache efficiency than Ligra because of its highly optimized SpMV engine.

As explained earlier, SSSP in Ligra is able to converge faster with less edge traversal in comparison to GPOP and GraphMat. Even with more work and memory references, GPOP still outperforms Ligra by upto 1.6× reduction in cache misses. GraphMat traverses the same number of edges as GPOP and incurs upto 2.2× more cache misses. The performance improvements for *soclj* are modest because of the small size of the graph.

6.2.2 Scalability

We evaluate the scalability of GPOP using PageRank and BFS algorithms. We measure the runtime of these algorithms on synthetic graphs with size ranging from $\{rmat22, |E| = 64M\}$ to $\{rmat27, |E| = 2048M\}$.

Strong scaling (figure 4): GPOP demonstrates very good scalability for BFS, achieving upto 17.9× speedup over a sequential implementation, using 36 threads. Note that the

Table 3. Comparison of L2 Cache Misses incurred by GPOP

Application	Datasets	GPOP	Ligra	GraphMat
PageRank	gplus	**1.3 B**	9.9 B	4.2 B
	sd1	**4.52 B**	52.7 B	19.4 B
	twitter	**3.43 B**	34.2 B	16.3 B
	pld	**1.66 B**	16.7 B	6.9 B
	friendster	**7.57 B**	38.7 B	-
	soclj	**0.17 B**	1.17 B	2.1 B
Label Propagation	gplus	**1.84 B**	3.85 B	2.29 B
	sd1	**6.74 B**	24.8 B	10.5 B
	twitter	**2.74 B**	8.2 B	4.82 B
	pld	**2.33 B**	5.65 B	3.16 B
	friendster	**12.2 B**	31.55 B	-
	soclj	0.75 B	0.36 B	**0.1 B**
SSSP	gplus	0.81 B	**0.78 B**	1.4 B
	sd1	**4.16 B**	5.5 B	7.86 B
	twitter	**1.56 B**	2.34 B	3.36 B
	pld	**1.17 B**	1.31 B	2.47 B
	friendster	**3.7 B**	6.1 B	-
	soclj	0.113 B	0.096 B	**0.08 B**

speedup is higher for large datasets that perform more work per iteration reducing the relative overhead of parallelization and synchronization.

In case of PageRank, GPOP achieves upto 10.5× speedup with 36 threads and scales poorly after 16 threads. This is because PageRank always uses PC mode scatter and nearly saturates the bandwidth with ≈ 20 threads. Therefore, increasing number of threads to 36 while keeping memory bandwidth constant does not improve overall performance.

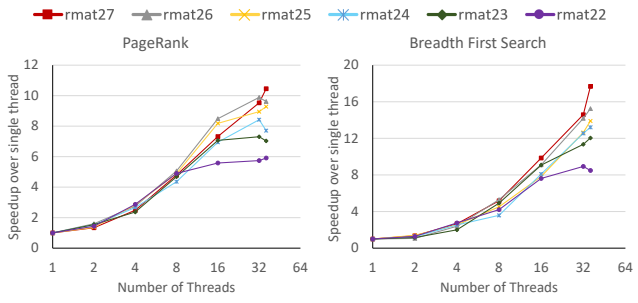

Figure 4. Strong scaling of GPOP for PageRank and BFS

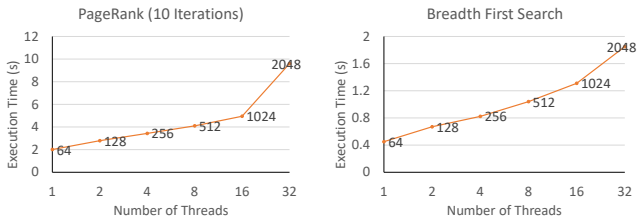

Figure 5. Weak scaling of GPOP for PageRank and BFS. Labels on the graph indicate #edges(in Millions)

Weak scaling (figure 5): The execution time of BFS increases only 4× with a 32× increase in the problem size. When executing small graphs on few threads, GPOP creates very few partitions. This results in very efficient message aggregation and reduced DRAM communication volume that improves the performance of small graphs. Due to the NUMA effect, we also observe a steeper increase in execution time when number of threads increases from 16 to 32.

PageRank executes completely in PC mode and the execution time increases only by 2.5× with a 16 fold increase in the problem size from *rmat22* to *rmat26*. However, due to memory bandwidth limits explained earlier, we see a sharp increase of 1.9× from *rmat26* to *rmat27* dataset.

6.2.3 Communication Mode Comparison

As explained in section 3, GPOP uses dual Partition-centric and Source-centric communication modes during Scatter phase. Here, we compare the performance of these modes individually (GPOP_SC and GPOP_PC) and in conjunction (GPOP).

Figure 6 depicts the execution time of each iteration for BFS, Label Propagation and SSSP (Bellman-Ford), for two of the largest datasets used in this paper - *friendster* and *rmat*27. In BFS and SSSP, frontier is dense in the middle iterations and very sparse in the initial and last few iterations. In Label Propagation, all vertices are active initially and the frontier size decreases as algorithm proceeds.

Execution time of GPOP_PC stays constant for almost all but a few iterations. This is because in these iterations, the 2-level hierarchical list prevents scattering of partitions with no active vertices. Runtime of GPOP_SC varies drastically with frontier size and is larger than GPOP_PC for iterations with large number of active vertices. GPOP achieves the best of both worlds by analytically predicting and choosing the mode with better performance. In almost all iterations, GPOP executes faster than both GPOP_SC and GPOP_PC which empirically validates our analytical model (section 3.2).

7 Related Work

Pregel [20] was one of the earliest (distributed-memory) graph processing frameworks that proposed the Vertex-centric (VC) programming model with the Gather-Apply-Scatter (GAS) abstraction. Many shared memory graph analytic frameworks [12, 28, 34] also adopt the VC model with *push* or *pull* direction execution (algorithm 1). The direction optimizing BFS[**?**] performs a frontier dependent switch between the two directions. Ligra [28] is a framework that generalizes this idea to other graph algorithms. However, it requires atomic instructions and the push-pull methodology is not cache and memory efficient, as discussed in section 2.

Several works [3, 8, 27, 31, 36] borrow the 2-phased GAS abstraction (also used by GPOP) and avoid use of synchronization primitives on shared-memory platforms. GraphMat

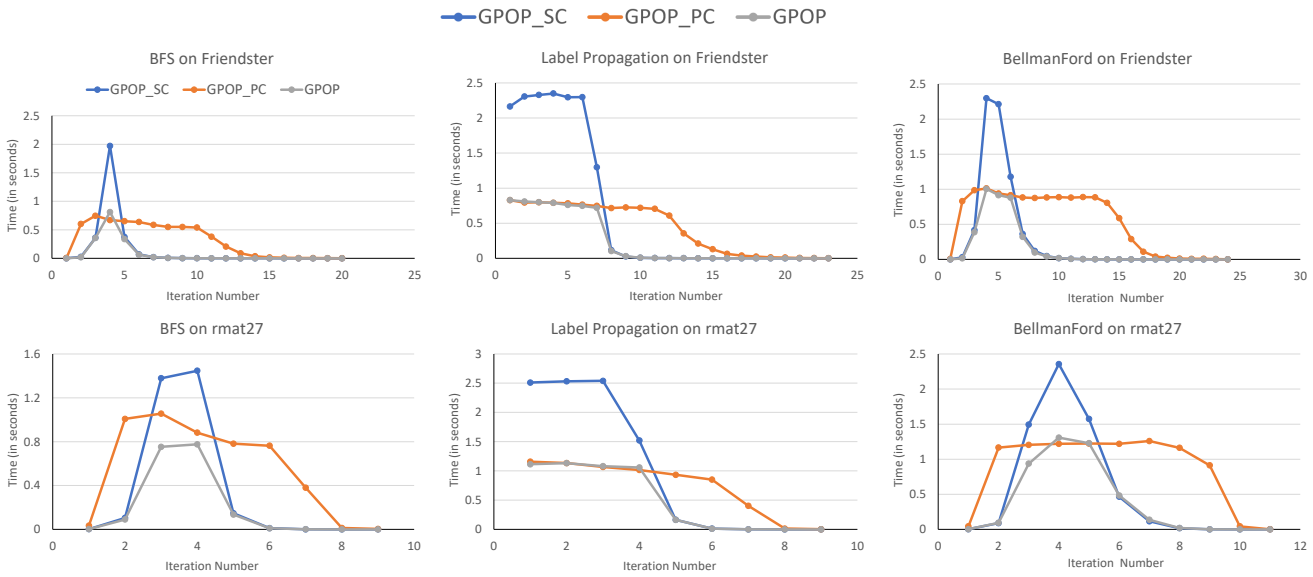

Figure 6. Iteration-wise execution time of GPOP with different communication modes on *friendster* and *rmat27* graphs

[31] maps graph algorithms to efficient Sparse Matrix (SpM) kernels but its bit-vector representation of frontier requires $O(V)$ work in each iteration. X-Stream uses Edge-centric programming and performs $O(E)$ traversal in each iteration. Modified EC or blocking techniques [8, 36] have also been utilized to improve locality and cache performance. The partitioning based PageRank [17] sequentializes memory accesses to enhance sustained bandwidth. However, these works [8, 17, 36] are only applicable to Sparse Matrix Dense Vector multiplication or PageRank computation, where all vertices are active.

The 2-phased SpMSpV algorithm of [3] buckets the columns corresponding to non-zero entries in the SpV and merges them to compute the product. It keeps a constant number of buckets to ensure work efficiency. However, for large graphs, the number of vertices in a bucket can outgrow the cache capacity, resulting in high cache miss ratio. Unlike [3], GPOP caps the partition size to ensure cacheability of vertices and still achieves work-efficiency using the *2-level Active List*.

A key difference between other shared-memory frameworks using GAS [3, 8, 27, 31, 36] and GPOP is communication volume and memory efficiency. The former typically generate messages for neighboring *vertices* and propagate them in a Source-centric fashion. In contrast, GPOP *aggregates* messages to neighboring *partitions* and optimally switches between dual communication modes, thus reducing DRAM traffic while simultaneously increasing bandwidth.

Some frameworks trade off programmability and preprocessing overheads for performance. Polymer [34] is a NUMA-aware framework that reduces random accesses to remote nodes, generated by its underlying Ligra engine. However, these benefits come at the cost of heavy preprocessing and storage requirements. Further, NUMA-aware partitioning has been shown to worsen the performance for some algorithms such as BFS [21]. Grazelle [12] deploys NUMA optimizations along with vectorized loop processing to improve performance of *pull* direction VC. While we do not focus on these aspects in this paper, NUMA awareness can be instilled in GPOP because of the explicit communication and high access locality, and sequential accesses to bins can be vectorized. The programming interface of Polymer and Grazelle is also quite complex and requires hundreds of lines of code to be written for simple applications, such as BFS.

Domain Specific Languages (DSL) are another line of research in high performance graph analytics. Green-Marl [14] is a popular DSL that allows users to write graph algorithms intuitively and relies on its compiler to extract parallelism and performance. GraphIt [35] allows user to explore the trade off space of various optimizations (push vs pull, dense vs sparse frontier, NUMA partitioning etc). Perhaps, cache and memory optimizations in GPOP can also become a part of such languages in the future.

8 Conclusion

In this paper, we presented the GPOP framework for graph processing, which achieves high cache and memory performance while guaranteeing theoretical efficiency. It also allows completely lock and atomic free computation, enabling good scalability.

We carefully developed the user interface of GPOP keeping ease of programming in mind and showed how various graph algorithms can be implemented with miniscule changes in user code.

We compared GPOP against Ligra (Vertex-centric push-pull programming) and GraphMat (Sparse Matrix operations with Vertex-centric GAS abstraction) by running various

graph algorithms on large datasets. GPOP incurs dramatically less L2 cache misses than the baseline frameworks and is upto 19× and 6.1× faster than Ligra and GraphMat, respectively. We found that DRAM bandwidth for sequential accesses is a limiting factor for GPOP's performance. Hence, execution can be further accelerated by using high bandwidth memory platforms, such as Hybrid Memory Cube.

References

[1] Takuya Akiba, Yoichi Iwata, Ken-ichi Kawarabayashi, and Yuki Kawata. 2014. Fast shortest-path distance queries on road networks by pruned highway labeling. In *2014 Proceedings of the Sixteenth Workshop on Algorithm Engineering and Experiments (ALENEX)*. SIAM, 147–154.

[2] Krste Asanovic, Ras Bodik, Bryan Christopher Catanzaro, Joseph James Gebis, Parry Husbands, Kurt Keutzer, David A Patterson, William Lester Plishker, John Shalf, Samuel Webb Williams, and others. 2006. *The landscape of parallel computing research: A view from berkeley*. Technical Report. Technical Report UCB/EECS-2006-183, EECS Department, University of California, Berkeley.

[3] Ariful Azad and Aydin Buluç. 2017. A work-efficient parallel sparse matrix-sparse vector multiplication algorithm. In *Parallel and Distributed Processing Symposium (IPDPS), 2017 IEEE International*. IEEE, 688–697.

[4] Scott Beamer, Krste Asanović, and David Patterson. 2013. Direction-optimizing breadth-first search. *Scientific Programming* 21, 3-4 (2013), 137–148.

[5] Maciej Besta, MichałPodstawski, Linus Groner, Edgar Solomonik, and Torsten Hoefler. 2017. To Push or To Pull: On Reducing Communication and Synchronization in Graph Computations. In *Proceedings of the 26th International Symposium on High-Performance Parallel and Distributed Computing*. ACM, 93–104.

[6] Florian Bourse, Marc Lelarge, and Milan Vojnovic. 2014. Balanced graph edge partition. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 1456–1465.

[7] Aydın Buluç and John R Gilbert. 2011. The Combinatorial BLAS: Design, implementation, and applications. *The International Journal of High Performance Computing Applications* 25, 4 (2011), 496–509.

[8] Daniele Buono, Fabrizio Petrini, Fabio Checconi, Xing Liu, Xinyu Que, Chris Long, and Tai-Ching Tuan. 2016. Optimizing sparse matrix-vector multiplication for large-scale data analytics. In *Proceedings of the 2016 International Conference on Supercomputing*. ACM, 37.

[9] Deepayan Chakrabarti, Yiping Zhan, and Christos Faloutsos. 2004. R-MAT: A recursive model for graph mining. In *Proceedings of the 2004 SIAM International Conference on Data Mining*. SIAM, 442–446.

[10] Ryan Compton, David Jurgens, and David Allen. 2014. Geotagging one hundred million twitter accounts with total variation minimization. In *Big Data (Big Data), 2014 IEEE International Conference on*. IEEE, 393–401.

[11] Neil Zhenqiang Gong, Wenchang Xu, Ling Huang, Prateek Mittal, Emil Stefanov, Vyas Sekar, and Dawn Song. 2012. Evolution of social-attribute networks: measurements, modeling, and implications using google+. In *Proceedings of the 2012 Internet Measurement Conference*. ACM, 131–144.

[12] Samuel Grossman, Heiner Litz, and Christos Kozyrakis. 2018. Making pull-based graph processing performant. In *Proceedings of the 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. ACM, 246–260.

[13] Minyang Han and Khuzaima Daudjee. 2015. Giraph unchained: barrierless asynchronous parallel execution in pregel-like graph processing systems. *Proceedings of the VLDB Endowment* 8, 9 (2015), 950–961.

[14] Sungpack Hong, Hassan Chafi, Edic Sedlar, and Kunle Olukotun. 2012. Green-Marl: a DSL for easy and efficient graph analysis. In *ACM SIGARCH Computer Architecture News*. ACM, 349–362.

[15] Jérôme Kunegis. 2013. Konect: the koblenz network collection. In *Proceedings of the 22nd International Conference on World Wide Web*. ACM, 1343–1350.

[16] Haewoon Kwak, Changhyun Lee, Hosung Park, and Sue Moon. 2010. What is Twitter, a social network or a news media?. In *Proceedings of the 19th international conference on World wide web*. ACM, 591–600.

[17] Kartik Lakhotia, Rajgopal Kannan, and Viktor Prasanna. 2017. Accelerating PageRank using Partition-Centric Processing. *arXiv preprint arXiv:1709.07122* (2017).

[18] Lingda Li, Robel Geda, Ari B Hayes, Yanhao Chen, Pranav Chaudhari, Eddy Z Zhang, and Mario Szegedy. 2017. A Simple Yet Effective Balanced Edge Partition Model for Parallel Computing. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 1, 1 (2017), 14.

[19] Yucheng Low, Danny Bickson, Joseph Gonzalez, Carlos Guestrin, Aapo Kyrola, and Joseph M Hellerstein. 2012. Distributed GraphLab: a framework for machine learning and data mining in the cloud. *Proceedings of the VLDB Endowment* 5, 8 (2012), 716–727.

[20] Grzegorz Malewicz, Matthew H Austern, Aart JC Bik, James C Dehnert, Ilan Horn, Naty Leiser, and Grzegorz Czajkowski. 2010. Pregel: a system for large-scale graph processing. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*. ACM, 135–146.

[21] Jasmina Malicevic, Baptiste Lepers, and Willy Zwaenepoel. 2017. Everything you always wanted to know about multicore graph processing but were afraid to ask. In *2017 USENIX Annual Technical Conference (USENIX ATC 17)*. USENIX Association, 631–643.

[22] John D McCalpin. 1995. STREAM benchmark. *Link: www. cs. virginia. edu/stream/ref. html# what* 22 (1995).

[23] Frank McSherry, Michael Isard, and Derek G. Murray. 2015. Scalability! But at What Cost?. In *Proceedings of the 15th USENIX Conference on Hot Topics in Operating Systems (HOTOS'15)*. USENIX Association, 14–14. http://dl.acm.org/citation.cfm?id=2831090.2831104

[24] Robert Meusel, Sebastiano Vigna, Oliver Lehmberg, and Christian Bizer. 2015. The graph structure in the web: Analyzed on different aggregation levels. *The Journal of Web Science* 1, 1 (2015), 33–47.

[25] Richard C Murphy, Kyle B Wheeler, Brian W Barrett, and James A Ang. 2010. Introducing the graph 500. *Cray Users Group (CUG)* 19 (2010), 45–74.

[26] Donald Nguyen, Andrew Lenharth, and Keshav Pingali. 2013. A lightweight infrastructure for graph analytics. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*. ACM, 456–471.

[27] Amitabha Roy, Ivo Mihailovic, and Willy Zwaenepoel. 2013. X-stream: Edge-centric graph processing using streaming partitions. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*. ACM, 472–488.

[28] Julian Shun and Guy E Blelloch. 2013. Ligra: a lightweight graph processing framework for shared memory. In *ACM Sigplan Notices*, Vol. 48. ACM, 135–146.

[29] Julian Shun, Farbod Roosta-Khorasani, Kimon Fountoulakis, and Michael W Mahoney. 2016. Parallel local graph clustering. *Proceedings of the VLDB Endowment* 9, 12 (2016), 1041–1052.

[30] Daniel A Spielman and Shang-Hua Teng. 2013. A local clustering algorithm for massive graphs and its application to nearly linear time graph partitioning. *SIAM J. Comput.* 42, 1 (2013), 1–26.

[31] Narayanan Sundaram, Nadathur Satish, Md Mostofa Ali Patwary, Subramanya R Dulloor, Michael J Anderson, Satya Gautam Vadlamudi, Dipankar Das, and Pradeep Dubey. 2015. Graphmat: High performance graph analytics made productive. *Proceedings of the VLDB Endowment* 8, 11 (2015), 1214–1225.

[32] Richard Vuduc, James W Demmel, and Katherine A Yelick. 2005. OSKI: A library of automatically tuned sparse matrix kernels. In *Journal of Physics: Conference Series*. IOP Publishing, 521.

[33] Thomas Willhalm, Roman Dementiev, and Patrick Fay. 2016. Intel Performance Counter Monitor-A better way to measure CPU utilization. 2012. *URL: http://software.intel.com/en-us/articles/intel-performance-counter-monitor-a-better-way-to-measure-cpuutilization* (2016).

[34] Kaiyuan Zhang, Rong Chen, and Haibo Chen. 2015. NUMA-aware graph-structured analytics. *ACM SIGPLAN Notices* 50, 8 (2015), 183–193.

[35] Yunming Zhang, Mengjiao Yang, Riyadh Baghdadi, Shoaib Kamil, Julian Shun, and Saman Amarasinghe. 2018. GraphIt-A High-Performance DSL for Graph Analytics. *arXiv preprint arXiv:1805.00923* (2018).

[36] Shijie Zhou, Kartik Lakhotia, Shreyas G Singapura, Hanqing Zeng, Rajgopal Kannan, Viktor K Prasanna, James Fox, Euna Kim, Oded Green, and David A Bader. 2017. Design and implementation of parallel PageRank on multicore platforms. In *High Performance Extreme Computing Conference (HPEC), 2017 IEEE*. IEEE, 1–6.

[37] Shijie Zhou and Viktor K Prasanna. 2017. Accelerating Graph Analytics on CPU-FPGA Heterogeneous Platform. In *2017 29th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)*. IEEE, 137–144.

[38] Xiaojin Zhu and Zoubin Ghahramani. Learning from labeled and unlabeled data with label propagation. (????).

[39] Xiaowei Zhu, Wentao Han, and Wenguang Chen. 2015. GridGraph: Large-Scale Graph Processing on a Single Machine Using 2-Level Hierarchical Partitioning. In *2015 USENIX Annual Technical Conference (USENIX ATC 15)*. USENIX Association, 375–386. https://www.usenix.org/conference/atc15/technical-session/presentation/zhu