

# Accelerating Direction-Optimized Breadth First Search on Hybrid Architectures

Scott Sallinen<sup>(✉)</sup>, Abdullah Gharaibeh, and Matei Ripeanu

University of British Columbia, Vancouver, Canada  
{scotts,abdullah,matei}@ece.ubc.ca

**Abstract.** Large scale-free graphs are famously difficult to process efficiently: the skewed vertex degree distribution makes it difficult to obtain balanced partitioning. Our research instead aims to turn this into an advantage by partitioning the workload to match the strength of the individual computing elements in a Hybrid, GPU-accelerated architecture. As a proof of concept we focus on the direction-optimized breadth first search algorithm. We present the key graph partitioning, workload allocation, and communication strategies required for massive concurrency and good overall performance. We show that exploiting specialization enables gains as high as 2.4x in terms of time-to-solution and 2.0x in terms of energy efficiency by adding 2 GPUs to a 2 CPU-only baseline, for synthetic graphs with up to 16 Billion undirected edges as well as for large real-world graphs. We also show that, for a capped energy envelope, it is more efficient to add a GPU than an additional CPU. Finally, our performance would place us at the top of today's [Green]Graph500 challenges for Scale29 graphs.

## 1 Introduction

Large-scale graph processing plays an important role for a wide range of applications where internal data structures are represented as graphs: from social networks, to protein-protein interactions, to the neuron structure of the human brain. A large class of real-world graphs, the *scale-free graphs*, have a heterogeneous vertex degree distribution: most vertices have a low degree but a few vertices are highly connected [2, 10, 20].

Breadth First Search (BFS) is the building block for many graph algorithms (e.g., Betweenness Centrality, Connected Components) and it has similar structural properties to other algorithms (e.g., Single Source Shortest Paths - SSSP). BFS also exposes the main challenges in graph processing: data-dependent memory accesses, low compute-to-memory access ratio, and low memory access locality. Additionally, for scale-free graphs, the amount of parallelism exposed is highly heterogeneous and data-dependent. For these reasons the Graph500 [7] and GreenGraph500 [8] graph processing competitions have adopted BFS on scale-free graphs as their main benchmark to compare the efficiency of graph processing platforms in terms of time-to-solution and energy.

Recently, Beamer et al. [1] introduced the *direction-optimized* BFS algorithm that takes advantage of the scale-free property (Sect. 2.2). This algorithm combines the classic *top-down* BFS traversal, with inverse *bottom-up* steps and offers a sizable speedup. To date, however, all implementations have focused either on CPU-only platforms [21] or require that the graph fits *entirely* in the accelerator memory [22].

Our past work [5] tests the hypothesis that assembling processing elements with diverse characteristics (i.e., massively-parallel processors optimized for high throughput, and traditional processors optimized for fast sequential processing) is a good match for scale-free graph workloads. While we have proven this hypothesis for a wide set of algorithms (including traditional top-down BFS, Connected Components, SSSP and PageRank), direction-optimized BFS poses new challenges: (i) as it is up to one order of magnitude faster than traditional BFS, it stresses the communication channels between the processing elements of the heterogeneous platform, exposing new bottlenecks; (ii) it requires both pull and push access to vertex state that has to be efficiently exposed by the supporting middleware; (iii) as the processing elements do not share memory, a low-overhead solution must be designed to coordinate them to switch between bottom-up and top-down phases of the algorithm, and finally, (iv) it requires specialized graph partitioning and workload allocation strategies that match the characteristics of the workload to those of the processing elements.

This paper makes the following contributions:

1. It provides further evidence that specialization – i.e., intelligent graph partitioning such that the resulting workload matches a heterogeneous set of processing elements – is key to extracting maximum efficiency when facing a fixed cost or energy constraint. (Sections 3.2, 4.1)
2. It extends TOTEM, our heterogeneous graph processing engine, to support a new class of frontier-based algorithms which require exposing both push and pull access to distributed vertex state. (Section 3)
3. It introduces optimizations key to boost the performance of direction-optimized BFS on a heterogeneous platform: partitioning and workload allocation, communication reduction, and improving access locality. (Sections 3.2, 3.3 and 3.4)

We evaluate these techniques across multiple hardware configurations and multiple large-scale graph workloads. Our evaluation shows an improvement of time-to-solution by up to 2.4x and energy efficiency by up to 2.0x against a CPU-only implementation, and compares favorably against state-of-the-art single node solutions (e.g., GALOIS). (Sections 4.2, 4.3)

## 2 Challenges and Opportunities

The key difficulty when processing scale-free graphs is a result of the heterogeneous vertex connectivity. (For example, over 70% of all vertices in the Twitter follower graph [3] have less than 40 in/out edges. The remaining vertices have increasingly large connectivity: the largest having over 3 million edges.)

This property makes obtaining balanced partitions difficult, as generally the memory footprint of a partition is proportional with the number of edges allocated to it, while the processing time is a complex function that depends on the number of vertices and edges allocated to the partition, and the specific properties of the workload (e.g., compute intensity, access locality).

Past work has generally assumed a homogeneous compute platform and has prioritized balancing partitions in terms of size [16]. This leads, however, to unbalanced partitions in terms of processing time due to the high-connectivity nodes. Recent strategies such as ‘high degree vertex delegation’ [17] continue to assume a homogeneous platform and aim for better load-balancing while dealing with high degree vertices.

## 2.1 Improving Performance with Hardware Accelerators

A GPU-accelerated system offers the opportunity to benefit from heterogeneity: instead of attempting to balance partitions by evenly distributing the workload based on memory footprint, one can choose to ‘embrace’ heterogeneity and partition such that the workload generated by a partition matches best the strengths of the processing element the partition is allocated to – e.g., by creating partitions that expose massive parallelism and allocating them to a GPU [4, 5].

However, efficiently harnessing a GPU-accelerated setup brings new challenges: First, it is difficult to design partition and allocation strategies that harness the platform efficiently. Second, GPUs tend to have over an order of magnitude less memory than the host and cannot process large partitions. (A key constraint – for example, the edge list of a Scale30 graph, a synthetic graph used in the Graph500 benchmark, occupies 256 GB in the memory-efficient Compressed Sparse Row format, yet a Kepler K40 GPU has only 12 GB of memory).

Note that past projects have explored GPU solutions, but either assume that the graph fits the memory of one [9, 22], or multiple GPUs [14]. In both cases, due to the limited memory space available, the scale of the graphs that can be processed is significantly smaller than the large graphs presented in this paper.


In any case, techniques that aim to optimize graph processing on the GPU are complementary to the approach proposed in this work in that they can be applied to the compute kernels to improve the overall performance of the hybrid system. In fact, this work uses the “virtual warp” technique proposed by Hong et al. [9] which aims to reduce divergence among the threads of a warp and hence improve the GPU kernel’s performance.

## 2.2 Improving Performance with Direction-Optimized BFS

Similar to other graph algorithms, level-synchronous Breadth First Search (BFS) exposes the concept of a *frontier*: the set of active vertices, which, for BFS, form the current level. To discover the next level, i.e., the next frontier, the traditional top-down BFS approach explores all edges of the vertices in the current frontier and builds the next frontier as the new vertices that can be reached (i.e., the vertices that have not been visited before). For scale-free graphs, this can cause

high write traffic as many edges out of the current frontier can attempt to add the same vertex into the next frontier.

*Direction-Optimized* BFS [1] is based on the key observation that the next frontier can also be built in a different, *bottom-up* way: by iterating over the vertices that have never been activated and selecting those that have a neighbour in the current frontier. Depending on graph topology and the current state of the algorithm, a bottom-up step can improve performance for two reasons. First, it can result in exploring fewer edges: once it has been determined that a vertex has a neighbour in the frontier it is not necessary to visit its other edges, thus reducing work particularly for high degree vertices. Second, it generates less contention as the thread that updates a vertex state (i.e., marks it as belonging to the new frontier) only reads its neighbour’s state but does not update it.



**Fig. 1.** Processing time per BFS level (left axis) and the average degree of vertices in the frontier (right axis). Graphs: *Left* Synthetic Scale30. *Right* Twitter [3].

During BFS processing of a scale-free graph, the vertices with high connectivity are quickly reached. Next, with these *few high degree* vertices in the frontier, the number of vertices in the next frontier will be large. At this point, it becomes more efficient to employ a bottom-up step: this will reduce the number of edges explored, and will alleviate write pressure corresponding to the many vertices that will be added to the frontier. Figure 1 supports this observation: the average degree of the frontier is large immediately after start (i.e., during the initial top-down steps of direction-optimized BFS). Next, during the bulk of the computation time, the average degree of a processed vertex is lower and continues to decrease; as a result, bottom-up steps are more efficient, but in effect, the many low degree vertices may end up being processed unsuccessfully at multiple BFS levels before they are finally included in the frontier. As the average degree of the frontier continues to decrease, top-down processing becomes again more efficient during the last few steps of the algorithm.

### 3 The Hybrid Algorithm

To harness the opportunities offered by a heterogeneous platform, several issues need to be addressed: (i) partitioning the graph and allocating partitions to processing elements to match their strengths and limits (i.e., massive parallelism yet limited accelerator memory); (ii) efficiently communicating between partitions; and (iii) coordinating the participating processing elements.

Our past work [5] demonstrates that the benefits of using a heterogeneous platform exceed the cost of communication between partitions hosted on different processing elements: the reduction in compute time obtained by adding GPUs is far larger than the added communication time over the PCI bus to synchronize between partitions. However, direction-optimized BFS adds new challenges: first, the bottom-up steps are up to one order of magnitude faster than the equivalent top-down steps thus potentially expose the communication over the PCI bus as the key bottleneck. Second, all processing elements participating in a direction-optimized BFS computation need to coordinate and chose direction at the same time. This coordination can add further communication overheads, as there is no shared state between the processing elements.

### 3.1 Direction-Optimized BFS for Partitioned Graphs

Our BFS Algorithm 1 is based on the Bulk Synchronous Parallel (BSP) model which supports well a system setup where processing elements that do not share memory, and where the graph needs to be partitioned for processing. Each BFS level of the algorithm contains a communication operation: the top-down steps use a Push-based method; the bottom-up steps a Pull-based method (described in Algorithms 2 and 3). The partitions each have an array of frontiers corresponding to other partitions. In addition, vertices have an associated partition ID, allowing the algorithm to determine whether or not the vertex is remote, and which frontier to use.

*Top-down Steps* explore the edges of the vertices in the frontier of the local partition: each such vertex activates – mark as belonging to the next frontier – a vertex that is either local or remote (belonging to another partition). During the *push* phase, the remote information is passed to the corresponding partitions, then all partitions wait for synchronization before starting the next round to ensure they all have pushed all necessary information.

*Bottom-up Steps* start by aggregating, in each partition, the global frontier by *pulling* the required information from all other partitions. Then each partition completes its level by checking its local not-yet-activated vertices against this global frontier. Then partitions synchronize: this ensures that all partitions have completed creating their new local frontier and are ready for the next round.

*Optimizations.* Key performance gains are achieved due to batch communication and message reduction: the push and pull operations only happen once per BSP round, and only to remote neighbours (i.e., only the data relevant to remote partitions). An additional optimization we apply is specific to the case when the user requires computing the BFS traversal tree as in the Graph500 benchmark (and not only labeling nodes with their ‘depth’): To reduce communication overhead, the parent of a vertex is not communicated during the traversal but is collected from the different address spaces in a final aggregation step (only the visited status is updated during traversal).

**Algorithm 1.** Direction-optimized BFS algorithm for partitioned graphs.

---

```

1  func BFS_Kernel (Partition PID, StepType STEP)
2  if (STEP == TOP-DOWN) then
3    parallel foreach (Vertex in Frontier[PID]) do
4      foreach (Nbr in Vertex.Neighbours) do
5        if (!Nbr.isVisited) then
6          NextFrontier[Nbr.partition].Add(Nbr)
7          BFSParentTree[Nbr.partition][Nbr] = Vertex
8          Nbr.isVisited = true
9        end if
10     end for
11   end parallel for
12   PushFrontiers(PID)
13 else if (STEP == BOTTOM-UP)
14   PullFrontiers(PID)
15   parallel foreach (Vertex in Partition[PID]) do
16     if (!Vertex.isVisited) then
17       foreach (Nbr in Vertex.Neighbours) do
18         if (Vertex in Frontier[Nbr.partition]) then
19           NextFrontier[PID].Add(Vertex)
20           BFSParentTree[PID][Vertex] = Nbr
21           Vertex.isVisited = true
22         break for
23       end if
24     end for
25   end if
26 end parallel for
27 end if
28 Synchronize()
29 end func

```

---

**Algorithm 2.** Push Frontiers

---

```

1  func PushFrontiers (Partition PID)
2  foreach (P in Partitions != PID) do
3    NextFrontier[P] ==> Frontier[P]
4    (local) ==> (remote)
5  end for
6  end func

```

---

**Algorithm 3.** Pull Frontiers

---

```

1  func PullFrontiers (Partition PID)
2  foreach (P in Partitions != PID) do
3    Frontier[P] <== NextFrontier[P]
4    (local) <== (remote)
5  end for
6  end func

```

---

### 3.2 Partition Specialization

A performance-critical decision is graph partitioning; in particular, we need to identify the part of the graph that should be placed on the space-constrained GPUs such that overall performance is maximized. We first observe that even though the bottom-up steps can significantly improve performance for some BFS levels (due to the reduction in total edge checks), these bottom-up steps take the longest out of the overall execution (Fig. 1). Thus accelerating these steps is essential for overall processing performance, and our focus.

We partition such that the low-degree vertices are assigned to the GPUs. The intuition behind this decision is threefold: first, processing the many low-degree vertices in parallel fits the GPU compute model (i.e., many small computations with insignificant load imbalance); second, the low-degree vertices occupy a small amount of memory (as they are not attached to many edges), a critical issue to the space constrained GPUs; and third, and most importantly, processing the low-degree vertices during the bottom-up steps is the main bottleneck as we have empirically verified. As we argue in the next section, this partitioning solution adds one additional advantage: it makes it easier to decide when to switch to bottom-up processing without communicating between partitions.

### 3.3 Switching Processing Direction for a Partitioned Graph

The direction-optimized algorithm requires coordinating all processing elements when the processing switches from top-down to bottom-up (after processing the

first few BFS levels) then switching back to top-down processing. These decisions are generally taken based on global information [1, 22] (e.g., the anticipated size of the next frontier) yet obtaining a more precise estimate is costly on a platform that does not offer shared memory.

*Top-down to Bottom-up Switch.* We estimate the next frontier based on a static percent of the edges out of the current frontier. This worked well in most executions on the scale-free graphs we have experimented with. However, when using this technique on our partitioned setup, it would normally be necessary to synchronize frontier information across each partition. However, as shown in Fig. 1, the frontier is rapidly built by the few high degree vertices, while the low degree vertices have virtually no impact on the decision to switch as they are discovered later. For this reason, the coordinator for switching can be the partition responsible for the high degree vertices: the CPU. This method is less costly than communicating among partitions to precisely anticipate frontier size, while retaining nearly identical accuracy in predicting the optimal switch point.

*Bottom-up to top-down switch.* The performance gains tends to be low from switching back to top-down processing as the final BFS levels require little time anyways. For this reason, partitions return to top-down after a fixed number of steps, so that all partitions return at the same time without state communicating or voting.

### 3.4 Optimizations to Improve Access Locality

After partitioning, a vertex is identified by two elements: a global ID which corresponds to its place in the original graph, and a local ID, which corresponds to its place in the partition. This indexing provides flexibility that can be exploited as follows. First, since the partition retains the global ID, permutation of local IDs enables optimizations: we can reorder vertices in memory to improve local partition access locality [18]. Second, the adjacency lists can be ordered in decreasing order of vertex connectivity, so that the highest degree vertex in the adjacency list comes first. This optimization shortens the bottom up-steps as the higher degree vertices are most likely to belong in the frontier, thus scanning the neighbour list has a higher chance to stop earlier (also noted by [21]).

Finally, note that the optimizations discussed in this section are applicable to both CPU-only and GPU-accelerated setups. Indeed, such optimizations makes it even more challenging to show the benefits of a heterogeneous platform as they significantly improve the performance of the CPU-only baseline, and hence they expose the communication and coordination overheads. However, as we show in the next section, our optimizations related to reducing communication overheads successfully eliminate it as a potential bottleneck.

## 4 Experimental Results

**Software Platform.** TOTEM [5], the framework we use to support our exploration, hides the complexity of developing graph algorithms from the programmer by providing abstractions for communication, the ability to specify graph

partitioning strategies, as well as common optimizations such as bitmap frontier representations and vertex and adjacency list ordering. We implemented the direction-optimized BFS algorithm on top of TOTEM, as well as the optimizations discussed – it is important to note that both the GPU-accelerated and the CPU-only experiments use the same CPU kernel (i.e., they both apply the optimizations discussed in Sect. 3.4).

**Hardware Platform.** The experiments were executed on a single machine with two Intel Xeon E5-2670v2 processors with 10 cores at 2.5 GHz and 512 GB of shared memory. The machine hosts two NVidia Kepler K40s with 2880 cores at 0.75 GHz and 12 GB of memory each. The peak memory bandwidth of the host is 59.7 GB/s, while on the GPU is 288 GB/s.

**Methodology.** We employ the experimental methodology defined by Graph500 and GreenGraph500. These require computing the BFS parent of each vertex (as opposed to only its level). While TOTEM uses the CSR format and represents each undirected edge as two directed edges, we do report performance in *undirected traversed edges per second (TEPS)*, as required by Graph500. Reported results are harmonic means over 100 executions. We measure power at the outlet using a WattsUP meter that samples at 1 Hz. To get representative energy consumption, we run each experiment for 10 min (e.g., repeating searches).

**Workloads.** Unless otherwise mentioned, the synthetic graphs used are Scale30 [1B V, 16B E], built with the Graph500 reference code generator and parameters. The real-world graphs used are undirected versions of Twitter [3] [52M V, 1.9B E], Wikipedia [11] [27M V, 601M E], and LiveJournal [12] [4M V, 69M E].


#### 4.1 The Impact of Specialized Partitioning

Figure 2 (*left*) presents the processing rate for a Scale30 graph for configurations with one or two CPUs, and one or two GPUs. There are two takeaways: first, GPUs provide acceleration in all cases; and, relevantly for budget/energy-limited platforms, it is more efficient to add an additional GPU than an additional CPU.


Second, and most importantly, the plot highlights the benefits of workload specialization: with random partitioning adding GPUs only offers a speedup proportional with memory footprint of the offloaded partition. *The intelligent partitioning scheme we introduce offers a super-linear speedup: despite offloading only 8 % of the graph, 2 GPUs improve performance by 2.4x.*

Figure 2 (*right*) evaluates performance over multiple Graph500 sizes and shows that the GPU-accelerated setup with workload specialization consistently offers large gains. Larger-scale graphs tend to have a smaller TEPS rate due to lower data locality. We note that, despite the ability to allocate a larger part of the smaller graphs to the GPUs the gains level off for scale-free graphs: allocating more low-degree vertices becomes exponentially costly as the vertices have higher connectivity. The largest graph offers more potential for improvement if GPUs had more memory: ‘only’ 88 % of non-singleton vertices are allocated to the GPUs. This increases to 97 % for Scale29, and 99 % for Scale28; at which point there is not much room for performance gains from GPUs with larger memory.





**Fig. 2.** *Left:* Direction-optimized BFS processing rate for specialized and random partitioning on hardware configurations with variable number of CPU Sockets (S) and GPUs (G) for a Scale30 graph. *Right:* Processing rates for synthetic graphs with size varying over one order of magnitude: Scale27 to Scale30. The curve labeled 4S presents the performance by Beamer [1] on 4-Socket machine.



**Fig. 3.** BFS run time (ms) for a Scale30 graph broken down into components: initializing BFS status data, computation, and push- and pull-communication.


**Analyzing the Overheads.** Figure 3 highlights that performance is dominated by computation time: initialization, aggregation, and communication (presented separately for push/pull) are only a small fraction of the total runtime.

Figure 4 (*left*) breaks down the total runtime by BFS level for classic top-down BFS and direction-optimized BFS on a traditional (two CPU sockets – labeled 2S) and hybrid (two CPUs and two GPUs – labeled 2S2G) platform. The plot highlights two key points. First, it confirms the benefits of direction-optimization and it shows that these benefits are concentrated on faster processing of bottom-up levels 4 and 5. Second, it highlights the further gains offered by the hybrid platform, and pinpoints the gains to much faster level 4 processing.

As a result of the BSP model, the computation time is determined by the bottleneck processor in each step. Figure 4 (*right*) presents the processing time per-level for each processing element: although occasionally (for levels 5 and beyond) the bottleneck is with the GPUs, the computation time for the initial bottom-up level (level 3) by the CPU dwarfs the rest of the execution time, leaving the other load-balancing inefficiencies nearly irrelevant.

## 4.2 Comparison with Past Work Using Real-World Graphs

We use real-world graphs to compare performance to that of the state-of-the-art graph processing framework GALOIS, whose direction-optimized BFS implementation compares favorably [15] to that of Ligra [19], PowerGraph [6], and



**Fig. 4.** *Left:* Per-level runtime (ms) for top-down (classic) and direction-optimized (D/O) BFS for a 2 CPU platform (2S), versus 2 CPUs and 2 GPUs (2S2G). *Right:* Per-level execution time for CPUs/GPUs of the 2S2G platform for the direction-optimized execution in the left plot. *Workload:* Scale30 graph.

GraphLab [13]. (We run GALOIS on our experimental machine. We had extensive discussions with GALOIS authors to make sure comparisons are fair.)

Table 1 shows the following: first, as in Fig. 4 (*left*) direction-optimized BFS largely outperforms top-down BFS. Second, our CPU-only versions of top-down and direction-optimized BFS perform largely similar to their GALOIS counterpart: this provides evidence that the baselines we used earlier for comparison to showcase the gains offered by our solution are fair. Furthermore, since in our hybrid algorithm the CPU and GPU kernels are executing concurrently, and the CPU is the bottleneck processor, improving our CPU kernel improves our overall execution rate, thus we have made all efforts to have efficient CPU-only kernels.

The hybrid direction-optimized version provides a performance boost of 2.0x for Twitter compared to the best CPU-only version. The larger diameter and less scale-free nature of the last two graphs reduce the impact of the direction-optimized approach, and expose more of the hybrid implementation overheads. Additionally, these smaller graphs expose less opportunity for the massive parallelism GPUs could harness. Nevertheless the hybrid implementation still offers a 1.3x speedup for LiveJournal and Wikipedia.

The table also highlights that the hybridization and the algorithm-level optimizations are synergetic, and together, they offer a significant boost in performance over generic and even optimized CPU versions. These results suggest that other scale-free real-world graphs will benefit from the techniques we propose.

### 4.3 The Energy Case

For Scale30 graphs, at 10.86 MTEPS/W, our CPU only implementation respectably ranked #10 in the November 2014 Big Data category of the GreenGraph500 list [8]. The hybrid configuration achieved over 2x better energy efficiency, ranking #6 with 22.36 MTEPS/W. (We note that, our hybrid configuration ranked behind 5 similar submissions by the GraphCrest group [21], that all use more energy-efficient hardware: more and newer CPUs). For Scale29 graphs, on a recently acquired platform (2x Intel E5-2695, DDR4 memory, same GPUs) with

**Table 1.** TOTEM and GALOIS (v2.2.1) performance in billion TEPS (higher is better), across real-world graphs. TOTEM executions use the same CPU kernel. The Naive kernel shown doesn’t apply optimizations discussed in Sect. 3.4.

	Algorithm	Naive-2S	GALOIS-2S	TOTEM-2S	TOTEM-2S2G
Twitter	Top-Down	0.23	0.50	1.39	2.05
	Direction-Optimized		1.96	2.84	<b>5.78</b>
Wikipedia	Top-Down	0.84	0.42	1.14	1.29
	Direction-Optimized		1.12	1.49	<b>2.01</b>
LiveJournal	Top-Down	0.54	0.99	1.26	1.57
	Direction-Optimized		1.23	1.96	<b>2.59</b>

17.3 GTPEPS and 30.1 MTEPS/Watt we would rank at the top of today’s Graph500 and, respectively, GreenGraph500.

The reason behind the energy gains the hybrid platform offers is that the GPU enables faster race-to-idle for the whole system (including energy expensive RAM), which means that the system draws high power for a significantly shorter period. Moreover, the most important factor that contributes to the energy gains is that the GPU, the processor with the higher Thermal Design Power (TDP), races-to-idle much faster than the CPUs (as shown in Fig. 4). Finally we note that the property we observed for performance holds for energy efficiency: it is always better to add a GPU than a second CPU. For example, if we extrapolate the linear performance improvement from 1 CPU to 2 CPUs as in Fig. 2 to 4 homogeneous CPUs, *and conservatively assume these two new CPUs have no additional energy cost*, a system consisting of 4 of our CPUs would be approximately 16 MTEPS/W, still less efficient than our 2 CPU 2 GPU system.

## 5 Summary

This work presents the design, implementation and evaluation of a state-of-the-art BFS algorithm (Beamer et al.’s direction-optimized algorithm [1]) on top of a hybrid, GPU-accelerated platform. We present a number of critical optimizations that take advantage of both the characteristics of the hardware platform we target and common properties of many real-world datasets. We show that while the GPU has limited memory space, large-scale graphs can still benefit from GPU acceleration by carefully partitioning the graph such that the GPU is assigned the part of the workload that otherwise critically limits the overall performance. Moreover, we show that by applying simple yet effective optimizations, such gains are achieved even for discrete GPUs connected to the system via high-latency PCI bus. This offers a strong indication that these gains will hold for high-speed GPU platforms, such as AMD Fusion or NVLink.

Making progress on techniques able to harness heterogeneous platforms is essential in the context of current hardware trends: as the cost of energy continues to increase relative to the cost of silicon, future systems will host a wealth of

different processing units. In this context, partitioning the workload and assigning the partitions to the processing element where they can be executed most efficiently in terms of power or time becomes a key issue.

**Acknowledgement.** This work was supported in part by the Institute for Computing, Information and Cognitive Systems (ICICS) at UBC.

## References

1. Beamer, S., Patterson, D.A.: Searching for a parent instead of fighting over children: A fast breadth-first search implementation for graph500 (2011)
2. Bullmore, E., Sporns, O.: Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat. Rev. Neurosci.* **10**(3), 186–198 (2009)
3. Cha, M., Haddadi, H., Benevenuto, F., Gummadi, P.K.: Measuring user influence in twitter: The million follower fallacy. (2010)
4. Cumming, B., Fourestey, G., Fuhrer, O., Gysi, T.: Application centric energy-efficiency study of distributed multi-core and hybrid cpu-gpu systems. In: SC (2014)
5. Gharaibeh, A., Reza, T., Santos-Neto, E., Sallinen, S., Ripeanu, M.: Efficient large-scale graph processing on hybrid cpu and gpu systems. [arXiv:1312](https://arxiv.org/abs/1312.2041) (2014)
6. Gonzalez, J.E., Low, Y., Gu, H., Bickson, D., Guestrin, C.: Powergraph: distributed graph-parallel computation on natural graphs. In: OSDI (2012)
7. Graph500. <http://www.graph500.org/>
8. Green Graph500. <http://green.graph500.org/>
9. Hong, S., Kim, S.K., Oguntebi, T., Olukotun, K.: Accelerating cuda graph algorithms at maximum warp (2011)
10. Jeong, H., Mason, S.P., Barabási, A.L., Oltvai, Z.N.: Lethality and centrality in protein networks. *Nature* **411**(6833), 41–42 (2001)
11. Kunegis, J.: The koblenz network collection. In: World Wide Web Companion (2013)
12. Leskovec, J., Krevl, A.: SNAP Datasets: Stanford large network dataset collection, June 2014. <http://snap.stanford.edu/data>
13. Low, Y., Gonzalez, J.E., Kyrola, A., Bickson, D., Guestrin, C.E., Hellerstein, J.: Graphlab: a new framework for parallel machine learning (2014). [arXiv:1408.2041](https://arxiv.org/abs/1408.2041)
14. Merrill, D., Garland, M., Grimshaw, A.: Scalable gpu graph traversal. In: ACM SIGPLAN Notices, vol. 47, pp. 117–128. ACM (2012)
15. Nguyen, D., Lenharth, A., Pingali, K.: A lightweight infrastructure for graph analytics. In: SOSP (2013)
16. Pearce, R., Gokhale, M., Amato, N.M.: Scaling techniques for massive scale-free graphs in distributed (external) memory. In: IPDPS (2013)
17. Pearce, R., Gokhale, M., Amato, N.M.: Faster parallel traversal of scale free graphs at extreme scale with vertex delegates. In: SC (2014)
18. Sallinen, S., Borges, D., Gharaibeh, A., Ripeanu, M.: Exploring hybrid hardware and data placement strategies for the graph 500 challenge (2014)
19. Shun, J., Blesl, G.E.: Ligra: a lightweight graph processing framework for shared memory. In: ACM SIGPLAN Notices, vol. 48, pp. 135–146. ACM (2013)
20. Wang, X.F., Chen, G.: Complex networks: small-world, scale-free and beyond. *IEEE Circuits Syst. Mag.* **3**(1), 6–20 (2003)

21. Yasui, Y., Fujisawa, K., Sato, Y.: Fast and energy-efficient breadth-first search on a single NUMA system. In: Kunkel, J.M., Ludwig, T., Meuer, H.W. (eds.) ISC 2014. LNCS, vol. 8488, pp. 365–381. Springer, Heidelberg (2014)
22. You, Y., Bader, D., Dehnavi, M.M.: Designing a heuristic cross-architecture combination for breadth-first search. In: ICPP (2014)