

PHAST: Hardware-Accelerated Shortest Path Trees

Daniel Delling, Andrew V. Goldberg, Andreas Nowatzky, and Renato F. Werneck

Microsoft Research Silicon Valley

Mountain View, CA, 94043, USA

Email: {dadellin, goldberg, andnow, renatow}@microsoft.com

Abstract—We present a novel algorithm to solve the non-negative single-source shortest path problem on road networks and other graphs with low highway dimension. After a quick preprocessing phase, we can compute all distances from a given source in the graph with essentially a linear sweep over all vertices. Because this sweep is independent of the source, we are able to reorder vertices in advance to exploit locality. Moreover, our algorithm takes advantage of features of modern CPU architectures, such as SSE and multi-core. Compared to Dijkstra’s algorithm, our method needs fewer operations, has better locality, and is better able to exploit parallelism at multi-core and instruction levels. We gain additional speedup when implementing our algorithm on a GPU, where our algorithm is up to three orders of magnitude faster than Dijkstra’s algorithm on a high-end CPU. This makes applications based on all-pairs shortest-paths practical for continental-sized road networks. Several algorithms, such as computing the graph diameter, exact arc flags, or centrality measures (exact reaches or betweenness), can be greatly accelerated by our method.

Keywords—shortest paths; GPU; route planning; high performance computing

I. INTRODUCTION

In recent years, performance gains in computer systems have come mainly from increased parallelism. As a result, exploiting the full potential of a modern computer has become more difficult. Applications must not only work on multiple cores, but also access memory efficiently, taking into account issues such as data locality. Parallel algorithms are often unavailable or involve a compromise, performing more operations than the best sequential algorithm for the same problem. In this paper we introduce a compromise-free algorithm for the single-source shortest path problem on road networks. Our algorithm performs fewer operations than existing ones, while taking advantage of locality, multi-core and instruction-level parallelism.

The *single-source shortest path problem* is a classical optimization problem. Given a graph $G = (V, A)$, a length $\ell(a)$ assigned to each arc $a \in A$, and a source vertex s , the goal is to find shortest paths from s to all other vertices in the graph. Algorithms for this problem have been studied since the 1950’s. The *non-negative single-source shortest path problem* (NSSP), in which $\ell(a) \geq 0$, is a special case that comes up in several important applications. It can be solved more efficiently than the general case with Dijkstra’s algorithm [1], [2]. When implemented with the appropriate priority queues [3], its running time is within a factor of three of breadth-first search (BFS), a simple linear-time traversal of the graph. This indicates that any significant practical

improvements in performance must take advantage of better locality and parallelism. Both are hard to achieve based on Dijkstra’s algorithm [4], [5].

Motivated by web-based map services and autonomous navigation systems, the problem of finding shortest paths in road networks has received a great deal of attention recently; see e.g. [6], [7] for overviews. However, most research focused on accelerating point-to-point queries, in which both a source s and a target t are known. Up to now, however, Dijkstra’s algorithm was still the fastest known solution to the NSSP problem.

We present a new algorithm for the NSSP problem that works well for certain classes of graphs, including road networks. We call it *parallel hardware-accelerated shortest path trees (PHAST)*. Building on previous work on point-to-point algorithms, PHAST uses *contraction hierarchies* [8] to essentially reduce the NSSP problem to a traversal of a shallow, acyclic graph. This allows us to take advantage of modern computer architectures and get a significant improvement in performance.

The PHAST algorithm requires a preprocessing phase, whose cost needs a moderate number of shortest path computations to be amortized. Moreover, PHAST only works well on certain classes of graphs, namely those with low *highway dimension* [9]. Fortunately, however, road networks are among them. Several important practical applications require multiple shortest path computations on road networks, such as preprocessing for route planning (see, e.g., [10], [11], [12], [13], [14]) or the computation of certain centrality measures, like betweenness [15], [16]. For these applications, PHAST is extremely efficient. On continental-sized road networks, a purely sequential version of our algorithm is two orders of magnitude faster than the best previous solution. Moreover, PHAST scales well on multi-core machines. On a standard 4-core workstation, one can compute all-pairs shortest paths in a few days instead of several months.

Another development in modern computers is the availability of very powerful, highly-parallel, and relatively cheap graphics processing units (GPUs). They have a large number of specialized processors and a highly optimized memory system. Although aimed primarily at computer graphics applications, GPUs have increasingly been used to accelerate general-purpose computations [17]. In this paper, we propose an efficient GPU implementation of PHAST. Note that this is nontrivial, since GPUs are mostly geared towards computation on regular data objects, unlike actual road networks. Still, our

implementation achieves significant speedups even compared to the (highly optimized) CPU implementation of PHAST itself. On a standard workstation equipped with a high-end consumer graphics card, we gain another order of magnitude over CPU-based PHAST. This reduces the computation of all-pairs shortest paths on a continental-sized road network to about half a day, making applications requiring such computations practical.

This paper is organized as follows. Section II reviews Dijkstra’s algorithm and the point-to-point algorithm PHAST builds upon, contraction hierarchies [8]. Section III describes the basic PHAST algorithm. Section IV shows how to improve locality to obtain a faster single-core version of the algorithm. Section V shows how the algorithm can be parallelized in different ways, leading to even greater speedups on multi-core setups. Section VI describes a typical GPU architecture and a GPU implementation of PHAST. Section VII presents how to extend PHAST to compute the auxiliary data needed for some applications. Section VIII reports detailed experimental results. Final remarks are made in Section IX.

II. BACKGROUND

A. Dijkstra’s Algorithm

We now briefly review the NSSP algorithm proposed by Dijkstra [1] and independently by Dantzig [2]. For every vertex v , the algorithm maintains the length $d(v)$ of the shortest path from the source s to v found so far, as well as the predecessor (parent) $p(v)$ of v on the path. Initially $d(s) = 0$, $d(v) = \infty$ for all other vertices, and $p(v) = \text{null}$ for all v . The algorithm maintains a priority queue of *unscanned* vertices with finite d values. At each step, it removes from the queue a vertex v with minimum $d(v)$ value and *scans* it: for every arc $(v, w) \in A$ with $d(v) + \ell(v, w) < d(w)$, it sets $d(w) = d(v) + \ell(v, w)$ and $p(w) = v$. The algorithm terminates when the queue becomes empty.

Efficient implementations of this algorithm rely on fast priority queues. On graphs with n vertices and m arcs, an implementation of the algorithm using binary heaps runs in $O(m \log n)$ time. One can do better, e.g., using k -heaps [18] or Fibonacci heaps [19], the latter giving an $O(m + n \log n)$ bound. If arc lengths are integers in $[0 \dots C]$, bucket-based implementations of Dijkstra’s algorithm work well. The first such implementation, due to Dial [20], gives an $O(m + nC)$ bound. There have been numerous improvements, including some that are very robust in practice. In particular, multi-level buckets [21] and smart queues [3] run in $O(m + n \log C)$ worst-case time.

Smart queues actually run in linear time if arc lengths have a uniform distribution [3]. In fact, experimental results show that, when vertex IDs are randomly permuted, an implementation of NSSP using smart queues is usually within a factor of two of breadth-first search (BFS), and never more than three, even on especially-built bad examples.

For concreteness, throughout this paper we will illustrate the algorithms we discuss with their performance on one well-known benchmark instance representing the road network

of Western Europe [7], with 18 million vertices and 42 million arcs. (More detailed experiments, including additional instances, will be presented in Section VIII.) If vertex IDs are assigned at random, the smart queue algorithm takes 8.0 seconds on an Intel Core-i7 920 clocked at 2.67 GHz, and BFS takes 6.0 seconds. The performance of both algorithms improve if one reorders the vertices such that neighboring vertices tend to have similar IDs, since this reduces the number of cache misses during the computation. Interestingly, as observed in [22], reordering the vertices according to a simple depth first search (DFS), as explained in detail in Section VIII, already gives good results: Dijkstra’s algorithm takes 2.8 seconds, and BFS takes 2.0. We tested several other layouts but were unable to obtain significantly better performance. Therefore, unless otherwise noted, in the remainder of this paper we use the DFS layout when reporting running times.

B. Contraction Hierarchies

We now discuss the *contraction hierarchies* (CH) algorithm, proposed by Geisberger et al. [8] to speed up point-to-point shortest path computations on road networks. It has two phases. The preprocessing phase takes only the graph as input, and produces some auxiliary data. The query phase takes the source s and target t as inputs, and uses the auxiliary data to compute the shortest path from s to t .

The preprocessing phase of CH picks a permutation of the vertices and *shortcuts* them in this order. The *shortcut operation* deletes a vertex v from the graph (temporarily) and adds arcs between its neighbors to maintain the shortest path information. More precisely, for any pair $\{u, w\}$ of neighbors of v such that $(u, v) \cdot (v, w)$ is the only shortest path in between u and w in the current graph, we add a *shortcut* (u, w) with $\ell(u, w) = \ell(u, v) + \ell(v, w)$. The output of this routine is the set A^+ of shortcut arcs and the position of each vertex v in the order (denoted by $\text{rank}(v)$). Although any order gives a correct algorithm, query times and the size of A^+ may vary. In practice, the best results are obtained by on-line heuristics that select the next vertex to shortcut based, among other factors, on the number of arcs added and removed from the graph in each step [8].

The query phase of CH runs a bidirectional version of Dijkstra’s algorithm on the graph $G^+ = (V, A \cup A^+)$, with one crucial modification: both searches only look at *upward* arcs, those leading to neighbors with higher rank. More precisely, let $A^\uparrow = \{(v, w) \in A \cup A^+ : \text{rank}(v) < \text{rank}(w)\}$ and $A^\downarrow = \{(v, w) \in A \cup A^+ : \text{rank}(v) > \text{rank}(w)\}$. During queries, the forward search is restricted to $G^\uparrow = (V, A^\uparrow)$, and the reverse search to $G^\downarrow = (V, A^\downarrow)$. Each vertex v maintains estimates $d_s(v)$ and $d_t(v)$ on distances from s (found by the forward search) and to t (found by the reverse search). These values can be infinity. The algorithm keeps track of the vertex u minimizing $\mu = d_s(u) + d_t(u)$, and each search can stop as soon as the minimum value in its priority queue is at least as large as μ .

Consider the maximum-rank vertex u on the shortest s - t path. As shown by Geisberger et al. [8], u minimizes

$d_s(u) + d_t(u)$ and the shortest path from s to t is given by the concatenation of the s - u and u - t paths. Furthermore, the forward search finds the shortest path from s to u (which belongs to G^\uparrow), and the backward search finds the shortest path from u to t (which belongs to G^\downarrow).

This simple algorithm is surprisingly efficient on road networks. On the European instance, random s - t queries visit fewer than 400 vertices (out of 18 million) on average and take a fraction of a millisecond on a standard workstation. Preprocessing takes only about 10 minutes and adds fewer shortcuts than there are original arcs.

Note that the forward search can easily be made target-independent by running Dijkstra’s algorithm in G^\uparrow from s until the priority queue is empty. Even with this loose stopping criterion, the *upward search* only visits about 500 vertices on average. Also note that the distance label $d_s(v)$ of a scanned vertex v does not necessarily represent the actual distance from s to v —it may be only an upper bound. We would have to run a backward search from v to find the actual shortest path from s to v .

From a theoretical point of view, CH works well in networks with low *highway dimension* [9]. Roughly speaking, these are graphs in which one can find a very small set of “important” vertices that hit all “long” shortest paths.

III. BASIC PHAST ALGORITHM

We are now ready to discuss a basic version of PHAST, our new algorithm for the NSSP problem. It has two phases, preprocessing and (multiple) NSSP computations. The algorithm is efficient only if there are sufficiently many NSSP computations to amortize the preprocessing cost.

The preprocessing phase of PHAST just runs the standard CH preprocessing, which gives us a set of shortcuts A^+ and a vertex ordering. This is enough for correctness. We discuss improvements in the next section.

A PHAST query initially sets $d(v) = \infty$ for all $v \neq s$, and $d(s) = 0$. It then executes the actual search in two subphases. First, it performs a simple forward CH search: it runs Dijkstra’s algorithm from s in G^\uparrow , stopping when the priority queue becomes empty. This sets the distance labels $d(v)$ of all vertices visited by the search. The second subphase scans all vertices in G^\downarrow in descending rank order. (For correctness, any reverse topological order will do.) To scan v , we examine each incoming arc $(u, v) \in A^+$; if $d(v) > d(u) + \ell(u, v)$, we set $d(v) = d(u) + \ell(u, v)$.

Theorem 3.1: For all vertices v , PHAST computes correct distance labels $d(v)$ from s .

Proof: We have to prove that, for every vertex v , $d(v)$ eventually represents the distance from s to v in G (or, equivalently, in G^+). Consider one such v in particular, and let w be the maximum-rank vertex on the shortest path from s to v in G^+ . The first phase of PHAST is a forward CH query, and is therefore guaranteed to find the shortest s - w path and to set $d(w)$ to its correct value (as shown in [8]). By construction, G^+ contains a shortest path from w to v in which vertices

appear in descending rank order. The second phase scans the arcs in this order, which means $d(v)$ is computed correctly. ■

In this and the following few sections, our discussion will focus on the computation of distance labels only, and not the actual shortest path trees. Section VII shows how to compute parent pointers and other auxiliary data in a straightforward manner.

On our benchmark instance, PHAST performs a single-source shortest path computation in about 2.0 seconds, which is the same as BFS and lower than the 2.8 seconds needed for Dijkstra’s algorithm.

IV. IMPROVEMENTS

In this section we describe how the performance of PHAST can be significantly improved by taking into account the features of modern computer architectures. We focus on its second phase (the linear sweep), since the time spent on the forward CH search is negligible: less than 0.05 ms on our benchmark instance. In Section IV-A, we show how to improve locality when computing a single tree. Then, Section IV-B discusses how building several trees simultaneously not only improves locality even further, but also enables the use of special instruction sets provided by modern CPUs. Finally, Section IV-C explains how a careful initialization routine can speed up the computation.

A. Reordering Vertices

To explain how one can improve locality and decrease the number of cache misses, we first need to address data representation. For best locality, G^\uparrow and G^\downarrow are represented separately, since each phase of our algorithm works on a different graph.

Vertices have sequential IDs from 0 to $n - 1$. We represent G^\uparrow using a standard cache-efficient representation based on a pair of arrays. One array, *arclist*, is a list of arcs sorted by tail ID, i.e., arc (u, \cdot) appears before (w, \cdot) if $u < w$. This ensures that the outgoing arcs from vertex v are stored consecutively in memory. Each arc (v, w) is represented as a two-field structure containing the ID of the head vertex (w) and the length of the arc. The other array, *first*, is indexed by vertex IDs; *first*[v] denotes the position in *arclist* of the first outgoing arc from v . To traverse the adjacency list, we just follow *arclist* until we hit *first*[$v + 1$]. We keep a sentinel at *first*[n] to avoid special cases.

The representation of G^\downarrow is identical, except for the fact that *arclist* represents incoming instead of outgoing arcs. This means that *arclist* is sorted by head ID, and the structure representing an arc contains the ID of its tail (not head). Distance labels are maintained as a separate array, indexed by vertex IDs.

Given this representation, a simple reordering of the vertices leads to improved memory locality during the second phase of PHAST, which works on G^\downarrow .

To determine a good new order, we first assign *levels* to vertices. Levels can be computed as we shortcut vertices during preprocessing, as follows. Initialize all levels to zero; when

shortcutting a vertex u , we set $L(v) = \max\{L(v), L(u) + 1\}$ for each current neighbor v of u , i.e., for each v such that $(u, v) \in A^\uparrow$ or $(v, u) \in A^\downarrow$. By construction, we have the following lemma.

Lemma 4.1: If $(v, w) \in A^\downarrow$, then $L(v) > L(w)$.

This means that the second phase of PHAST can process vertices in descending order of level: vertices on level i are only visited after all vertices on levels greater than i have been processed. This respects the topological order of G^\downarrow .

Within the same level, we can scan the vertices in any order. In particular, by processing vertices within a level in increasing order of IDs,¹ we maintain some locality and decrease the running time of PHAST from 2.0 to 0.7 seconds.

We can obtain additional speedup by actually reordering the vertices. We assign lower IDs to vertices at higher levels; within each level, we keep the DFS order. Now PHAST will be correct with a simple linear sweep in increasing order of IDs. It can access vertices, arcs, and head distance labels sequentially, with perfect locality. The only non-sequential access is to the distance labels of the arc tails (recall that, when scanning v , we must look at the distance labels of its neighbors). Keeping the DFS relative order within levels helps reduce the number of associated cache misses.

As most data access is now sequential, we get a substantial speedup. With reordering, one Nssp computation is reduced from 0.7 seconds to 172 milliseconds, which is about 16.4 times faster than Dijkstra’s algorithm. We note that the notion of reordering vertices to improve locality has been applied before to hierarchical point-to-point speedup techniques [12], albeit in a more ad hoc manner.

B. Computing Multiple Trees

Reordering ensures that the only possible non-sequential accesses during the second stage of the algorithm happen when reading distance labels of arc tails. More precisely, when processing vertex v , we must look at all incoming arcs (u, v) . The arcs themselves are arranged sequentially in memory, but the IDs of their tail vertices are not sequential.

We can improve locality by running multiple Nssp computations simultaneously. To grow trees from k sources $(s_0, s_1, \dots, s_{k-1})$ at once, we maintain k distance labels for each vertex $(d_0, d_1, \dots, d_{k-1})$. These are maintained as a single array of length kn , laid out so that the k distances associated with v are consecutive in memory.

The query algorithm first performs (sequentially) k forward CH searches, one for each source, and sets the appropriate distance labels of all vertices reached. As we have seen, the second phase of PHAST processes vertices in the same order regardless of source, so we can process all k sources during the same pass. We do so as follows. To process each incoming arc (u, v) into a vertex v , we first retrieve its length and the ID of u . Then, for each tree i (for $0 \leq i < k$), we compute $d_i(u) + \ell(u, v)$ and update $d_i(v)$ if the new value is an improvement.

¹As mentioned in Section II-A, we assign IDs according to a DFS order, which has a fair amount of locality.

For a fixed v , all $d_i(v)$ values are consecutive in memory and are processed sequentially, which leads to better locality and fewer cache misses.

Increasing k leads to better locality, but only up to a point: storing more distance labels tend to evict other, potentially useful, data from the processor caches. Another drawback is increased memory consumption, because we need to keep an array with kn distance labels.

Still, for small values of k we can achieve significant speedups with a relatively small memory overhead. Setting $k = 16$ reduces the average running time per tree from 171.9 to 96.8 milliseconds on our benchmark instance.

SSE Instructions: We use 32-bit distance labels. Current x86-CPU’s have special 128-bit SSE registers that can hold four 32-bit integers and allow basic operations, such as addition and minimum, to be executed in parallel. We can use these registers during our sweep through the vertices to compute k trees simultaneously, k being a multiple of 4. For simplicity, assume $k = 4$. When processing an arc (u, v) , we load all four distance labels of u into an SSE register, and four copies of $\ell(u, v)$ into another. With a single SSE instruction, we compute the packed sum of these registers. Finally, we build the (packed) minimum of the resulting register with the four distance labels of v , loaded into yet another SSE register. Note that computing the minimum of integers is only supported by version 4.1 of SSE or higher.

For $k = 16$, using SSE instructions reduces the average runtime per tree from 96.8 to 37.1 milliseconds, for an additional factor of 2.6 speedup. In total, this algorithm is 76 times faster than Dijkstra’s algorithm on one core.

C. Initialization

PHAST (like Dijkstra’s algorithm) assumes that all distance labels are set to ∞ during initialization. This requires a linear sweep over all distance labels, which takes about 10 milliseconds. This is negligible for Dijkstra’s algorithm, but represents a significant time penalty for PHAST. To avoid this, we mark vertices visited during the CH search with a single bit. During the linear sweep, when scanning a vertex v , we check for this bit: If it is not set, we know $d(v) = \infty$, otherwise we know that v has been scanned during the upward search and has a valid (though not necessarily correct) value. After scanning v we unmark the vertex for the next shortest path tree computation. The results we have reported so far already include this implicit initialization.

V. EXPLOITING PARALLELISM

We now consider how to use parallelism to speed up PHAST on a multi-core CPU. For computations that require shortest path trees from several sources, the obvious approach for parallelization is to assign different sources to each core. Since the computations of the trees are independent from one another, we observe excellent speedups. Running on four cores, without SSE, the average running time per tree ($k = 1$) decreases from 171.9 to 47.1 ms, a speedup of 3.7. (Recall that k indicates the number of sources per linear sweep.) Setting k

to 16 (again without SSE), the running time drops from 113.3 to 28.5 ms per tree, also a speedup of 3.7.

However, we can also parallelize a single tree computation. On our benchmark instance, the number of the vertex levels is around 140, orders of magnitude smaller than the number of vertices. Moreover, low levels contain many more vertices than upper levels. Half of all vertices are in level 0, for example. This allows us to process vertices of the same level in parallel if multiple cores are available. We partition vertices in a level into (roughly) equal-sized blocks and assign each block to a thread (core). When all threads terminate, we start processing the next level. Blocks and their assignment to threads can be computed during preprocessing. Running on 4 cores, we can reduce a single NSSP computation from 171.9 to 49.7 milliseconds on the same machine, a factor of 3.5 speedup. Note that this type of parallelization is the key to our GPU implementation of PHAST, explained in the next section.

VI. GPU IMPLEMENTATION

Our improved implementation of PHAST is limited by the memory bandwidth. One way to overcome this limitation is to use a modern graphics card. The NVIDIA GTX 580 (Fermi) we use in our tests has a higher memory bandwidth (192.4 GB/s) than a high-end Intel Xeon CPU (32 GB/s). Although clock frequencies tend to be lower on GPUs (less than 1 GHz) than CPUs (higher than 3 GHz), the former can compensate by running many threads in parallel. The NVIDIA GTX 580 has 16 independent cores, each capable of executing 32 threads (called a *warp*) in parallel. It follows a Single Instruction Multiple Threads (SIMT) model, which uses predicated execution to preserve the appearance of normal thread execution at the expense of inefficiencies when the control-flow diverges. Moreover, barrel processing is used to hide DRAM latency. For maximal efficiency, all threads of a warp must access memory in certain, hardware-dependent ways. Accessing 32 consecutive integers of an array, for example, is efficient. Another constraint of GPU-based computations is that communication between main and GPU memory is rather slow. Fortunately, off-the-shelf GPUs nowadays have enough on-board RAM (1.5 GB in our case) to hold all the data we need.

Our GPU-based variant of PHAST, called *GPHAST*, satisfies all the constraints mentioned above. In a nutshell, *GPHAST* outsources the linear sweep to the GPU, while the CPU remains responsible for computing the upward CH trees. During initialization, we copy both G^\downarrow and the array of distance labels to the GPU. To compute a tree from s , we first run the CH search on the CPU and copy the search space (with less than 2 KB) to the GPU. As in the single-tree parallel version of PHAST, we then process each level in parallel. The CPU starts, for each level i , a *kernel* on the GPU, which is a (large) collection of threads that all execute the same code and that are scheduled by the GPU hardware. Each thread computes the distance label of exactly one vertex. With this approach, the overall access to the GPU memory within a warp

is efficient in the sense that DRAM bandwidth utilization is minimized.

Unfortunately, due to the nature of PHAST, we cannot take advantage of the much faster shared memory of the GPU. The reason is that data reuse is very low: each arc is only looked at exactly once, and each distance label is written once and read very few times (no more than twice on average). This severely limits the usefulness of shared memory. Still, it is implicitly utilized by the Fermi GPU because of the on-chip cache, which was configured to use the maximal amount of on-chip shared memory (48 KB out of 64 KB).

On an NVIDIA GTX 580, installed on the machine used in our previous experiments, a single tree can be computed in 5.53 ms. This represents a speedup of 511 over Dijkstra’s algorithm, 31 over the sequential variant of PHAST, and 9 over the 4-core CPU version of PHAST. Note that *GPHAST* uses a single core from the CPU.

We also tested reordering vertices by degree to make each warp work on vertices with the same degree. However, this has a strong negative effect on the locality of the distance labels of the tails of the incoming arcs. Hence, we keep the same ordering of vertices as for our CPU implementation of PHAST.

Multiple Trees: If the GPU has enough memory to hold additional distance labels, *GPHAST* also benefits from computing many trees in parallel. When computing k trees at once, the CPU first computes the k CH upward trees and copies all k search spaces to the GPU. Again, the CPU activates a GPU kernel for each level. Each thread is still responsible for writing exactly one distance label. We assign threads to warps such that threads within a warp work on the same vertices. This allows more threads within a warp to follow the same instruction flow, since they work on the same part of the graph. In particular, if we set $k = 32$, all threads of a warp work on the same vertex.

For $k = 16$, *GPHAST* needs 2.21 ms per shortest path tree. That is about 1280 times faster than the sequential version of Dijkstra’s algorithm, 78 times faster than sequential PHAST, and 8.5 times faster than computing 64 trees on the CPU in parallel (16 sources per sweep, one sweep per core). *GPHAST* can compute all-pairs shortest paths (i.e., n trees) in roughly 11 hours on a standard workstation. On the same machine, n executions of Dijkstra’s algorithm would take about 200 days, even if we compute 4 trees in parallel (one on each core).

VII. COMPUTING AUXILIARY INFORMATION

Our discussion so far has assumed that PHAST computes only distances from a root s to all vertices in the graph. We now discuss how it can be extended to compute the actual shortest path trees (i.e., parent pointers) and show how PHAST can be used for several concrete applications.

A. Building Trees

One can easily change PHAST to compute parent pointers in G^+ . When scanning v during the linear sweep phase, it suffices to remember the arc (u, v) responsible for $d(v)$. Note

that some of the parent pointers in this case will be shortcuts (not original arcs). For many applications, paths in G^+ are sufficient and even desirable [23].

If the actual shortest path tree in G is required, it can be easily obtained with one additional pass through the arc list of G . During this pass, for every original arc $(u, v) \in G$, we check whether the identity $d(v) = d(u) + \ell(u, v)$ holds; if it does, we make u the parent of v . As long as all original arc lengths are strictly positive, this leads to a shortest path tree in the original graph.

In some applications, one might need to compute all distance labels, but the full description of a single s - t path. In such cases, a path in G^+ can be expanded into the corresponding path in G in time proportional to the number of arcs on it [8].

B. Applications

With the tree construction procedure at hand, we can now give practical applications of PHAST: computing exact diameters and centrality measures on continental-sized road networks, as well as faster preprocessing for point-to-point route planning techniques. The applications require extra bookkeeping or additional traversals of the arc list to compute some auxiliary information. As we shall see, the modifications are easy and the computational overhead is relatively small.

a) Diameter: The diameter of a graph G is defined by the longest shortest path in G . Its exact value can be computed by building n shortest path trees. PHAST can easily do it by making each core keep track of the maximum label it encounters. The maximum of these values is the diameter. To use GPHAST, we maintain an additional array of size n to keep track of the maximum value assigned to each vertex over all n shortest path computations. In the end, we do one sweep over all vertices to collect the maximum. This is somewhat memory-consuming, but it keeps the memory accesses within the warps efficient.

b) Arc Flags: A well-known technique to speed up the computation of point-to-point shortest paths is to use *arc flags* [10], [11]. A preprocessing algorithm first computes a partition \mathcal{C} of V into *cells*, then attaches a *label* to each arc a . A label contains, for each cell $C \in \mathcal{C}$, a Boolean flag $F_C(a)$ which is true if there is a shortest path starting with a to at least one vertex in C . During queries, a modified variant of Dijkstra’s algorithm only considers arcs for which the flag of the target cell is set to true. This approach can easily be made bidirectional and is very efficient, with speedups of more than three orders of magnitude over a bidirectional version of Dijkstra’s algorithm [10].

The main drawback of this approach is its preprocessing time. While a good partition can be computed in a few minutes [24], [25], [26], [27], computing the flags requires building a shortest path tree from each boundary vertex, i.e., each vertex with an incident arc from another cell. In a typical setup, one has to compute about 40 000 shortest path trees resulting in preprocessing times of about 10.5 hours (on four cores). Instead of running Dijkstra’s algorithm, however, we

can run GPHAST with tree reconstruction, reducing the time to set flags to less than 3 minutes.

c) Centrality Measures: PHAST can also be used to compute the exact *reach* [13] of a vertex v . The reach of v is defined as the maximum, over all shortest s - t paths containing v , of $\min\{dist(s, v), dist(v, t)\}$. This notion is very useful to accelerate the computation of point-to-point shortest paths. The best known method to calculate exact reaches for all vertices within a graph requires computing all n shortest path trees, which we can do much faster with PHAST. Fast heuristics [12] compute reach upper bounds only and are fairly complicated.

Another frequently used centrality measure based on shortest paths is *betweenness* [15], [16] which is defined as $c_B(v) = \sum_{s \neq v \neq t \in V} \sigma_{st}(v) / \sigma_{st}$, where σ_{st} is the number of shortest paths between two vertices s and t , and $\sigma_{st}(v)$ is the number of shortest s - t paths on which v lies on. Computing exact betweenness relies on n shortest path tree computations [28]. Again, replacing Dijkstra’s algorithm by PHAST makes exact betweenness tractable for continent-size road networks. Moreover, it could also be helpful for accelerating known approximation techniques [28], [29].

Note that these and other applications often require traversing the resulting shortest path tree in bottom-up fashion. This can be easily done in a cache-efficient way by scanning vertices in level order.

VIII. EXPERIMENTAL RESULTS

A. Experimental Setup

We implemented our CPU-based PHAST with all optimizations from Section IV in C++ and compiled it with Microsoft Visual C++ 2010. For parallelization, we use OpenMP. CH queries use a binary heap as priority queue; we tested other data structures, but their impact on performance is negligible because the queue size is small.

As already mentioned, we run most of our evaluation on an Intel Core-i7 920 running Windows Server 2008R2. It has four cores clocked at 2.67 GHz and 12 GB of DDR3-1066 RAM. Our standard benchmark instance is the European road network, with 18 million vertices and 42 million arcs, made available by PTV AG [30] for the 9th DIMACS Implementation Challenge [7]. The length of each arc represents the travel time between its endpoints.

PHAST builds upon contraction hierarchies. We implemented a parallelized version of the CH preprocessing routine [8]. For improved efficiency, we use a slightly different priority function for ordering vertices. The priority of a vertex u is given by $2ED(u) + CN(u) + H(u) + 5L(u)$, where $ED(u)$ is the difference between the number of arcs added and removed (if u were contracted), $CN(u)$ is the number of contracted neighbors, $H(u)$ is the total number of arcs represented by all shortcuts added, and $L(u)$ is the level u would be assigned to. In this term, we bound $H(u)$ such that every incident arc of u can contribute at most 3. This ensures that this term is only important during the beginning of the contraction process. For faster preprocessing, we bound

witness searches by 5 hops up to an average degree (of the uncontracted graph) of 5 and by 10 hops up to degree 10. Beyond that, we use no limit. Finally, our code is parallelized: after contracting a vertex, we update the priorities of all neighbors simultaneously. On 4 cores, this gives a speedup of 2.5 over a sequential implementation.

With our implementation, CH preprocessing takes about five minutes (on four cores) and generates upward and downward graphs with 33.8 million arcs each. The number of levels is 140, with half the vertices assigned to the lowest level, as shown in Figure 1. Note that the lowest 20 levels contain all but 100 000 vertices, while all but 1 000 vertices are assigned to the lowest 66 levels. We stress that the priority term has limited influence on the performance of PHAST. It works well with any function that produces a “good” contraction hierarchy (leading to fast point-to-point queries), such as those tested by Geisberger et al. [8].

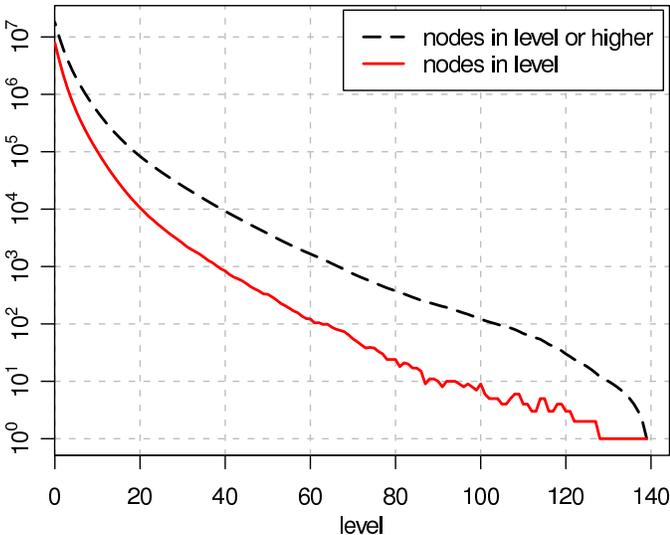


Fig. 1. Vertices per level

B. Single Tree

We now evaluate the performance of Dijkstra’s algorithm and PHAST when computing a single tree. We tested different priority queues (for Dijkstra’s algorithm) and different graph layouts (for both algorithms). We start with a *random* layout, in which vertex IDs are assigned randomly, to see what happens with poor locality. We also consider the *original* layout (as given in the input graph as downloaded); it has some spatial locality. Finally, we consider a *DFS* layout, with IDs given by the order vertices are discovered during a DFS from a random vertex. The resulting figures are given in Table I. As a reference we also include the running time of a simple BFS.

We observe that both the layout and the priority queue have an impact on Dijkstra’s algorithm. It is four times slower when using the binary heap and the random layout than when using a bucket-based data structure and the DFS layout. For single-core applications, the smart queue implementation [3] (based on multi-level buckets) is robust and memory-efficient.

TABLE I
PERFORMANCE OF VARIOUS ALGORITHMS ON THE EUROPEAN ROAD NETWORK. THREE GRAPH LAYOUTS ARE CONSIDERED: RANDOM, AS GIVEN IN THE INPUT, AND DFS-BASED.

algorithm	details	time per tree [ms]		
		random	input	DFS
Dijkstra	binary heap	11159	5859	5180
	Dial	7767	3538	2908
	smart queue	7991	3556	2826
BFS	—	6060	2445	2068
PHAST	original ordering	1286	710	678
	reordered by level	406	179	172
	reordered + four cores	144	53	50

In our setup, however, Dial’s implementation [20], based on single-level buckets, is comparable on a single core and scales better on multiple cores. For the remainder of this paper, all numbers given for Dijkstra’s algorithm executions refer to Dial’s implementation with the DFS layout.

The impact of the layout on PHAST is even more significant. By starting from the DFS ordering and then ordering by level, the average execution time for one source improves from 1286ms to 172ms, a speedup of 7.5. For all combinations tested, PHAST is always faster than Dijkstra’s algorithm. The speedup is about a factor of 16.5 for sequential PHAST, while this number increases to 57 if we use four cores to scan the vertices within one level in parallel, as explained in Section V.

To evaluate the overhead of PHAST, we also ran a lower bound test. To determine the memory bandwidth of the system, we sequentially and independently read from all arrays (*first*, *arclist*, and the distance array) and then write a value to each entry of the distance array. On our test machine, this takes 65.6ms; PHAST is only 2.6 times slower than this.

Note that this lower bound merely iterates through the arc list in a single loop. In contrast, most algorithms (including PHAST) loop through the vertices, and for each vertex loop through its (few) incident arcs. Although both variants visit the same arcs in the same order, the second method has an inner loop with a very small (but varying) number of iterations, thus making it harder to be sped up by the branch predictor. Indeed, it takes 153ms to traverse the graph exactly as PHAST does, but storing at $d(v)$ the sum of the lengths of all arcs into v . This is only 19ms less than PHAST, which suggests that reducing the number of cache misses (from reading $d(u)$) even further by additional reordering is unlikely to improve the performance of PHAST significantly.

C. Multiple Trees

Next, we evaluate the performance of PHAST when computing many trees simultaneously. We vary both k (the number of trees per linear sweep) and the number of cores we use. We also evaluate the impact of SSE instructions. The results are given in Table II.

Without SSE, we observe almost perfect speedup when using four cores instead of one, for all values of k . With SSE, however, we obtain smaller speedups when running on multiple cores. The more cores we use, and the higher k we

TABLE II

AVERAGE RUNNING TIMES PER TREE WHEN COMPUTING MULTIPLE TREES IN PARALLEL. WE CONSIDER THE IMPACT OF USING SSE INSTRUCTIONS, VARYING THE NUMBER OF CORES, AND INCREASING THE NUMBER OF SOURCES PER SWEEP (k). THE NUMBERS IN PARENTHESES REFER TO THE EXECUTION TIMES WHEN SSE IS ACTIVATED; FOR THE REMAINING ENTRIES WE DID NOT USE SSE.

sources/ sweep	time per tree [ms]		
	1 core	2 cores	4 cores
1	171.9	86.7	47.1
4	121.8 (67.6)	61.5 (35.5)	32.5 (24.4)
8	105.5 (51.2)	53.5 (28.0)	28.3 (20.8)
16	96.8 (37.1)	49.4 (22.1)	25.9 (18.8)

pick, the more data we have to process in one sweep. Still, using all optimizations (SSE, multi-core) helps: the algorithm is more than 9 times faster with $k = 16$ on four cores than with $k = 1$ on one core.

For many cores and high values of k , memory bandwidth becomes the main bottleneck for PHAST. This is confirmed by executing our lower bound test on all four cores in parallel. For $k = 16$ and four cores, it takes 12.8ms per “tree” to traverse all arrays in optimal (sequential) order. This is more than two thirds of the 18.8ms needed by PHAST. This indicates that PHAST is approaching a barrier. This observation was the main motivation for our implementation of PHAST on a GPU, whose performance we discuss next.

D. GPHAST

To evaluate GPHAST, we implemented our algorithm from Section VI using CUDA SDK 3.2 and compiled it with Microsoft Visual C++ 2008. (CUDA SDK 3.2 is not compatible with Microsoft Visual C++ 2010 at the time of writing.) We conducted our experiments on an NVIDIA GTX 580 installed in our benchmark machine. The GPU is clocked at 772 MHz and has 1.5 GB of DDR5 RAM. Table III reports the performance when computing up to 16 trees simultaneously.

TABLE III

PERFORMANCE AND GPU MEMORY UTILIZATION OF GPHAST IN MILLISECONDS PER TREE, DEPENDING ON k , THE NUMBER OF TREES PER SWEEP.

trees / sweep	memory [MB]	time [ms]
1	395	5.53
2	464	3.93
4	605	3.02
8	886	2.52
16	1448	2.21

As the table shows, the performance of GPHAST is excellent. A single tree can be built in only 5.53ms. When computing 16 trees in parallel, the running time per tree is reduced to a mere 2.21ms. This is a speedup of more than three orders of magnitude over plain Dijkstra’s algorithm. On average, we only need 123 picoseconds per distance label, which is roughly a third of a CPU clock cycle. On four cores, CH preprocessing takes 302 seconds, but this cost is amortized away after only 319 trees are computed if one uses GPHAST instead of Dijkstra’s algorithm (also on four cores).

E. Hardware Impact

In this section we study the performance of PHAST on different computer architectures. Although GPHAST is clearly faster than PHAST, GPU applications are still very limited, and general-purpose servers usually do not have high-performance GPUs. Table IV gives an overview of the five machines we tested. It should be noted that M2-1 and M2-4 are older machines (about 5 and 3 years old, respectively), whereas M1-4 (our default machine) is a recent commodity workstation. M2-6 and M4-12 are modern servers costing an order of magnitude more than M1-4. Note that M4-12 has many more cores than M2-6, but has worse sequential performance due to a lower clock rate. With the exception of M1-4, all machines have more than one NUMA node (local memory bank). For these machines, access to local memory is faster than to memory assigned to a different NUMA node. M4-12 has more NUMA nodes (eight) than CPUs (four). M4-12 and M2-6 run Windows Server 2008R2, M2-4 runs Windows Server 2008, and M2-1 runs Windows Server 2003. Note that we use SSE instructions only on M1-4 and M2-6 because the remaining machines do not support SSE 4.2 (we need the minimum operator not supported by previous versions).

TABLE IV

SPECIFICATIONS OF THE MACHINES TESTED. COLUMN $|P|$ INDICATES THE NUMBER OF CPUs, WHEREAS $|c|$ REFERS TO THE TOTAL NUMBER OF PHYSICAL CORES IN THE MACHINE. COLUMN $|B|$ REFERS TO HOW MANY LOCAL MEMORY BANKS THE SYSTEM HAS (NUMA NODES). THE GIVEN MEMORY BANDWIDTH REFERS TO THE (THEORETICAL) SPEED WITH WHICH A CORE CAN ACCESS ITS LOCAL MEMORY.

name	brand	type	CPU		memory				
			clock [GHz]	$ P $ $ c $	type	size [GB]	clock [MHz]	bandw. [GB/s]	$ B $
M2-1	AMD	Opteron 250	2.40	2 2	DDR	16	133	6.4	2
M2-4	AMD	Opteron 2350	2.00	2 8	DDR2	64	266	12.8	2
M4-12	AMD	Opteron 6168	1.90	4 48	DDR3	128	667	42.7	8
M1-4	Intel	Core-i7 920	2.67	1 4	DDR3	12	533	25.6	1
M2-6	Intel	Xeon X5680	3.33	2 12	DDR3	96	667	32.0	2

We tested the sequential and parallel performance of Dijkstra’s algorithm and PHAST on these machines. By default, the operating system moves threads from core to core during execution, which can have a significant adverse effect on memory-bound applications such as PHAST. Hence, we also ran our experiments with each thread pinned to a specific core. This ensures that the relevant distance arrays are always stored in the local memory banks, and results in improved locality at all levels of memory hierarchy. For the same reason, on multi-socket systems, we also copy the graph to each local memory bank explicitly (when running in pinned mode). Table V shows the results.

Running single-threaded, PHAST outperforms Dijkstra’s algorithm by a factor of approximately 19, regardless of the machine. This factor increases slightly (to 21) when we compute one tree per core. The reason for this is that cores share the memory controller(s) of a CPU. Because PHAST has fewer cache misses, it benefits more than Dijkstra’s algorithm from the availability of multiple cores.

TABLE V

IMPACT OF DIFFERENT COMPUTER ARCHITECTURES ON DIJKSTRA’S ALGORITHM AND PHAST. WHEN RUNNING MULTIPLE THREADS, WE EXAMINE THE EFFECT OF PINNING EACH THREAD TO A SPECIFIC CORE OR KEEPING IT UNPINNED (FREE). IN EACH CASE, WE SHOW THE AVERAGE RUNNING TIME PER TREE IN MILLISECONDS.

machine	Dijkstra [ms]			PHAST [ms]				
	single thread	1 tree/core		single thread	1 tree/core		16 trees/core	
		free	pinned		free	pinned	free	pinned
M2-1	6073.5	3967.3	3499.3	315.4	184.4	158.3	99.5	85.0
M2-4	6497.5	1499.0	1232.2	330.5	104.0	56.6	49.0	31.4
M4-12	5183.1	417.3	168.5	272.7	49.1	8.0	18.4	4.0
M1-4	2907.2	951.6	947.7	171.9	48.1	47.1	19.0	18.8
M2-6	2321.7	413.7	288.8	134.9	21.2	14.5	14.5	7.2

The impact of pinning threads and copying the graph to local memory banks is significant. Without pinning, no algorithm performs well when run in parallel on a machine with more than one NUMA node. On M4-12, which has four CPUs (and eight NUMA nodes), we observe speedups of less than 6 (the number of cores of a single NUMA node) when using all 48 cores, confirming that non-local memory access is inefficient. However, if the data is properly placed in memory, the algorithms scale much better. On M4-12, using 48 cores instead of a single one makes PHAST 34 times faster. Unsurprisingly, pinning is not very helpful on M1-4, which has a single CPU.

Computing 16 trees per core within each sweep gives us another factor of 2, independent of the machine. When using all cores, PHAST is consistently about 40 times faster than Dijkstra’s algorithm.

F. Comparison: Dijkstra, PHAST, and GPHAST

Next, we compare Dijkstra’s algorithm with PHAST and GPHAST. Table VI reports the best running times (per tree) for all algorithms, as well as how long it would take to solve the all-pairs shortest-paths problem. We also report how much energy these computations require (for GPU computations, this includes the entire system, including the CPU). In this experiment, we also include the predecessor of the NVIDIA GTX 580, the GTX 480. Also based on the Fermi architecture, the GTX 480 has fewer cores (15 instead of 16) and lower clock rates, for both the cores (701 MHz instead of 772 MHz) and the memory (1848 MHz instead of 2004 MHz). The remaining specifications are the same for both cards.

On the most powerful machine we tested, M4-12, the CPU-based variant is almost as fast as GPHAST, but the energy consumption under full work load is much higher for M4-12 (747 watts) than for M1-4 with a GPU installed (GTX 480: 390 watts, GTX 580: 375 watts). Together with the fact that GPHAST (on either GPU) still is faster than PHAST on M4-12, the energy consumption per tree is about 2.8 to 3.6 times worse for M4-12. M1-4 without a GPU (completely removed from the system) consumes 163 watts and is about as energy-efficient as M4-12. Interestingly, M2-6 (332 watts) does a better job than the other machines in terms of energy per tree. Still, GPHAST on a GTX 580 is almost three times more efficient than M2-6. We observe that GPHAST is 20%

TABLE VI

DIJKSTRA’S ALGORITHM, PHAST, AND GPHAST COMPARISON. COLUMN “MEMORY USED” INDICATES HOW MUCH MAIN MEMORY IS OCCUPIED DURING THE CONSTRUCTION OF THE TREES (FOR GPHAST, WE ALSO USE 1.5 GB OF GPU MEMORY).

algorithm	device	memory used [GB]	per tree		n trees	
			time [ms]	energy [J]	time [d:h:m]	energy [MJ]
			Dijkstra	M1-4	2.8	947.72
	M2-6	8.2	288.81	95.88	60:04:51	1725.93
	M4-12	31.8	168.49	125.86	35:02:55	2265.52
PHAST	M1-4	5.8	18.81	3.07	3:22:06	55.19
	M2-6	15.6	7.20	2.39	1:12:13	43.03
	M4-12	61.8	4.03	3.01	0:20:09	54.19
GPHAST	GTX 480	2.2	2.69	1.05	0:13:27	18.88
	GTX 580	2.2	2.21	0.83	0:11:03	14.92

faster on the GTX 580 than on the GTX 480, with slightly lower energy consumption.

A GTX 580 graphics card costs half as much as the M1-4 machine on which it is installed, and the machine supports two cards. With two cards, GPHAST would be twice as fast, computing all-pairs shortest paths in roughly 5.5 hours (1.12 ms per tree), at a fifth of the cost of M4-12 or M2-6. In fact, one could even buy some very cheap machines equipped with 2 GPUs each. Since the linear sweep is by far the bottleneck of GPHAST, we can safely assume that the all-pairs shortest-paths computation scales perfectly with the number of GPUs.

G. Other Inputs

Up to now, we have only tested one input, the European road network with travel times. Here, we evaluate the performance of our algorithm if applied to travel distances instead of travel times. CH preprocessing takes about 41 minutes on this input, generating upwards and downwards graphs with 410 levels and 38.8 millions arcs each. Moreover, we evaluate the road network of the US (generated from TIGER/Line data [31]), also made available for the 9th DIMACS Implementation Challenge [7]. It has 24 million vertices and 58.3 million arcs. For travel times (distances), the CH preprocessing takes 10 (28) minutes, and the search graphs have 50.6 (53.7) million arcs and 101 (285) levels. Table VII presents the results.

All algorithms are slower on the US graph, which has about 6 million more vertices than Europe. More interestingly,

TABLE VII
PERFORMANCE OF DIJKSTRA’S ALGORITHM, PHAST, AND GPHAST ON OTHER INPUTS.

algorithm	device	Europe		USA	
		time	distance	time	distance
Dijkstra	M1-4	947.72	609.19	1269.1	947.75
	M2-6	288.81	177.58	380.40	280.17
	M4-12	168.49	108.58	229.00	167.77
PHAST	M1-4	18.81	22.25	27.11	28.81
	M2-6	7.20	8.27	10.42	10.71
	M4-12	4.03	5.03	6.18	6.58
GPHAST	GTX 480	2.69	4.54	4.07	5.41
	GTX 580	2.21	3.88	3.41	4.65

switching from travel times to distances has a positive effect on Dijkstra’s algorithm (there are fewer *decrease-key* operations), but makes PHAST slower (it has more arcs to scan). However, the differences are relatively small. PHAST is always much faster than Dijkstra’s algorithm, and GPHAST yields the best performance on all inputs.

H. Arc-Flags

Our last set of experiments deals with the computation of arc flags (as described in Section VII). The purpose of this test is to show that additional information (besides the distance labels) can indeed be computed efficiently. As input, we again use the road network of Western Europe with travel times. First, we use SCOTCH [26] to create a partition of the graph into 128 cells and 20240 boundary vertices in total. This takes less than a minute. Next, we remove the so-called *1-shell* (attached trees) of the graph, which has roughly 6 million vertices. Optimal flags for arcs within these trees can be set easily [14], [10]. This step takes 2 seconds. We then start the computation of the remaining arc flags. We compute for each boundary vertex two shortest path trees with GPHAST (one forward, one backward) and set the corresponding flags accordingly. We do this by copying G to the GPU, together with an array with 32-bit integers representing 32 flags for each arc. Since we need to compute 128 flags per arc, we copy this array to the main memory after computing 32 flags and reinitialize it. We do so due to memory constraints on the GPU; we set $k = 8$ for the same reason. Overall, with this approach the GPU memory is almost fully loaded.

The last step, tree construction and setting of arc flags, takes 158 seconds. This is 3.91 ms per boundary vertex (and direction) on average, of which 1.94 ms are spent computing the 12 million distance labels. Hence, reconstructing the parent pointers and setting the flags takes exactly as much time as computing the tree. This is expected, since we have to look at almost the same amount of data as during tree construction. On four cores we reduce the overall time for computing flags from 10.5 hours with Dijkstra’s algorithm to less than 9 minutes (including partitioning and the CH preprocessing).

IX. CONCLUDING REMARKS

We presented PHAST, a new algorithm for computing shortest path trees in graphs with low highway dimension, such as road networks. Not only is its sequential version faster than the best existing sequential algorithm, but it also exploits parallel features of modern computer architectures, such as SSE instructions, multiple cores, and GPUs. The GPU implementation can compute a shortest path tree about three orders of magnitude faster than Dijkstra’s algorithm. This makes many applications on road networks, such as the exact computation of centrality measures, practical.

A previously studied parallel algorithm for the NSSP problem is Δ -stepping [32]. It performs more sequential operations than Dijkstra’s algorithm, its parallel implementation requires fine-grained synchronization, and the amount of parallelism in Δ -stepping is less than that in PHAST. Thus, for a large

number of NSSP computations, PHAST is a better choice. For a small number of computations (e.g., a single computation) PHAST is not competitive because of the preprocessing. However, it is not clear if on road networks Δ -stepping is superior to Dijkstra’s algorithm in practice. The only study of Δ -stepping on road networks [4] has been done on a Cray MTA-2, which has an unusual architecture and requires a large amount of parallelism for good performance. The authors conclude that continent-size road networks do not have sufficient parallelism. It would be interesting to see how the Δ -stepping algorithm performs on a more conventional multi-core system or a small cluster. Another interesting project is an MTA-2 implementation of PHAST.

Future work includes studying the performance of PHAST to other networks. A possible approach would be to stop the construction of the contraction hierarchy as soon as the average degree of the remaining vertices exceeds some value. PHAST then has to explore more vertices during the CH upwards search and would only sweep over those vertices that were contracted during preprocessing.

Finally, it would be interesting to study which kinds of map services are enabled by the ability to compute full shortest path trees in real time, and by the fact that preprocessing based on all-pairs shortest paths is now feasible.

Acknowledgements

We thank Diego Nehab for interesting discussions and his valuable input.

REFERENCES

- [1] E. W. Dijkstra, “A Note on Two Problems in Connexion with Graphs,” *Numerische Mathematik*, vol. 1, pp. 269–271, 1959.
- [2] G. B. Dantzig, *Linear Programming and Extensions*. Princeton University Press, 1962.
- [3] A. V. Goldberg, “A Practical Shortest Path Algorithm with Linear Expected Time,” *SIAM Journal on Computing*, vol. 37, pp. 1637–1655, 2008.
- [4] K. Madduri, D. A. Bader, J. W. Berry, and J. R. Crobak, “Parallel Shortest Path Algorithms for Solving Large-Scale Instances,” in *The Shortest Path Problem: Ninth DIMACS Implementation Challenge*, ser. DIMACS Book, C. Demetrescu, A. V. Goldberg, and D. S. Johnson, Eds. American Mathematical Society, 2009, vol. 74, pp. 249–290.
- [5] U. Meyer and P. Sanders, “ Δ -Stepping : A Parallel Single Source Shortest Path Algorithm,” in *Proceedings of the 6th Annual European Symposium on Algorithms (ESA’98)*, ser. Lecture Notes in Computer Science, vol. 1461, 1998, pp. 393–404.
- [6] D. Delling, P. Sanders, D. Schultes, and D. Wagner, “Engineering Route Planning Algorithms,” in *Algorithmics of Large and Complex Networks*, ser. Lecture Notes in Computer Science, J. Lerner, D. Wagner, and K. Zweig, Eds. Springer, 2009, vol. 5515, pp. 117–139.
- [7] C. Demetrescu, A. V. Goldberg, and D. S. Johnson, Eds., *The Shortest Path Problem: Ninth DIMACS Implementation Challenge*, ser. DIMACS Book. American Mathematical Society, 2009, vol. 74.
- [8] R. Geisberger, P. Sanders, D. Schultes, and D. Delling, “Contraction Hierarchies: Faster and Simpler Hierarchical Routing in Road Networks,” in *Proceedings of the 7th International Workshop on Experimental Algorithms (WEA’08)*, ser. Lecture Notes in Computer Science, C. C. McGeoch, Ed., vol. 5038. Springer, June 2008, pp. 319–333.
- [9] I. Abraham, A. Fiat, A. V. Goldberg, and R. F. Werneck, “Highway Dimension, Shortest Paths, and Provably Efficient Algorithms,” in *Proceedings of the 21st Annual ACM–SIAM Symposium on Discrete Algorithms (SODA’10)*, 2010, pp. 782–793.

- [10] M. Hilger, E. Köhler, R. H. Möhring, and H. Schilling, "Fast Point-to-Point Shortest Path Computations with Arc-Flags," in *The Shortest Path Problem: Ninth DIMACS Implementation Challenge*, ser. DIMACS Book, C. Demetrescu, A. V. Goldberg, and D. S. Johnson, Eds. American Mathematical Society, 2009, vol. 74, pp. 41–72.
- [11] U. Lauther, "An Experimental Evaluation of Point-To-Point Shortest Path Calculation on Roadnetworks with Precalculated Edge-Flags," in *The Shortest Path Problem: Ninth DIMACS Implementation Challenge*, ser. DIMACS Book, C. Demetrescu, A. V. Goldberg, and D. S. Johnson, Eds. American Mathematical Society, 2009, vol. 74, pp. 19–40.
- [12] A. V. Goldberg, H. Kaplan, and R. F. Werneck, "Reach for A*: Shortest Path Algorithms with Preprocessing," in *The Shortest Path Problem: Ninth DIMACS Implementation Challenge*, ser. DIMACS Book, C. Demetrescu, A. V. Goldberg, and D. S. Johnson, Eds. American Mathematical Society, 2009, vol. 74, pp. 93–139.
- [13] R. J. Gutman, "Reach-Based Routing: A New Approach to Shortest Path Algorithms Optimized for Road Networks," in *Proceedings of the 6th Workshop on Algorithm Engineering and Experiments (ALENEX'04)*. SIAM, 2004, pp. 100–111.
- [14] R. Bauer and D. Delling, "SHARC: Fast and Robust Unidirectional Routing," *ACM Journal of Experimental Algorithmics*, vol. 14, no. 2.4, pp. 1–29, August 2009, special Section on Selected Papers from ALENEX 2008.
- [15] J. M. Anthonisse, "The rush in a directed graph," Stichting Mathematisch Centrum, 2e Boerhaavestraat 49 Amsterdam, Tech. Rep. BN 9/71, October 1971.
- [16] L. C. Freeman, "A Set of Measures of Centrality Based Upon Betweenness," *Sociometry*, vol. 40, pp. 35–41, 1977.
- [17] J. D. Owens, D. Luebke, N. Govindaraju, M. Harris, J. Krüger, A. Lefohn, and T. J. Purcell, "A Survey of General-Purpose Computation on Graphics Hardware," *Computer Graphics Forum*, vol. 26, no. 1, pp. 80–113, 2007.
- [18] D. B. Johnson, "Efficient Algorithms for Shortest Paths in Sparse Networks," *Journal of the ACM*, vol. 24, no. 1, pp. 1–13, January 1977.
- [19] R. E. Tarjan and M. L. Fredman, "Fibonacci Heaps and Their Uses in Improved Network Optimization Algorithms," *Journal of the ACM*, vol. 34, no. 3, pp. 596–615, July 1987.
- [20] R. B. Dial, "Algorithm 360: Shortest-Path Forest with Topological Ordering [H]," *Communications of the ACM*, vol. 12, no. 11, pp. 632–633, 1969.
- [21] E. V. Denardo and B. L. Fox, "Shortest-Route Methods: 1. Reaching, Pruning, and Buckets," *Operations Research*, vol. 27, no. 1, pp. 161–186, 1979.
- [22] P. Sanders, D. Schultes, and C. Vetter, "Mobile Route Planning," in *Proceedings of the 16th Annual European Symposium on Algorithms (ESA'08)*, ser. Lecture Notes in Computer Science, vol. 5193. Springer, September 2008, pp. 732–743.
- [23] R. Bauer, D. Delling, P. Sanders, D. Schieferdecker, D. Schultes, and D. Wagner, "Combining Hierarchical and Goal-Directed Speed-Up Techniques for Dijkstra's Algorithm," *ACM Journal of Experimental Algorithmics*, vol. 15, no. 2.3, pp. 1–31, January 2010, special Section devoted to WEA'08.
- [24] G. Karypis and G. Kumar, "A Fast and Highly Quality Multilevel Scheme for Partitioning Irregular Graphs," *SIAM Journal on Scientific Computing*, vol. 20, no. 1, pp. 359–392, 1999.
- [25] V. Osipov and P. Sanders, " n -Level Graph Partitioning," in *Proceedings of the 18th Annual European Symposium on Algorithms (ESA'10)*, ser. Lecture Notes in Computer Science. Springer, 2010, pp. 278–289.
- [26] F. Pellegrini and J. Roman, "Scotch: A software package for static mapping by dual recursive bipartitioning of process and architecture graphs," in *High-Performance Computing and Networking*, ser. Lecture Notes in Computer Science. Springer, 1996, pp. 493–498. [Online]. Available: <http://www.springerlink.com/content/fg24302095611p34/>
- [27] D. Delling, A. V. Goldberg, I. Razenshteyn, and R. F. Werneck, "Graph Partitioning with Natural Cuts," in *25th International Parallel and Distributed Processing Symposium (IPDPS'11)*. IEEE Computer Society, 2011.
- [28] U. Brandes, "A Faster Algorithm for Betweenness Centrality," *Journal of Mathematical Sociology*, vol. 25, no. 2, pp. 163–177, 2001.
- [29] R. Geisberger, P. Sanders, and D. Schultes, "Better Approximation of Betweenness Centrality," in *Proceedings of the 10th Workshop on Algorithm Engineering and Experiments (ALENEX'08)*, I. Munro and D. Wagner, Eds. SIAM, April 2008, pp. 90–100.
- [30] PTV AG - Planung Transport Verkehr, "<http://www.ptv.de>."
- [31] US Census Bureau, "UA Census 2000 TIGER/Line files." 2002.
- [32] U. Meyer and P. Sanders, " Δ -Stepping: A Parallelizable Shortest Path Algorithm," *Journal of Algorithms*, vol. 49, no. 1, pp. 114–152, 2003.