

Engineering Highway Hierarchies

PETER SANDERS, Karlsruhe Institut für Technologie
DOMINIK SCHULTES, Technische Hochschule Mittelhessen

Highway hierarchies exploit hierarchical properties inherent in real-world road networks to allow fast and exact point-to-point shortest-path queries. A fast preprocessing routine iteratively performs two steps: First, it removes edges that only appear on shortest paths *close* to source or target; second, it identifies low-degree nodes and bypasses them by introducing shortcut edges. The resulting hierarchy of highway networks is then used in a Dijkstra-like bidirectional query algorithm to considerably reduce the search space size without losing exactness. The crucial fact is that ‘far away’ from source and target it is sufficient to consider only high-level edges.

Experiments with road networks for a continent show that using a preprocessing time of around 15 min, one can achieve a query time of around 1ms on a 2.0GHz AMD Opteron.

Highway hierarchies can be combined with goal-directed search, they can be extended to answer many-to-many queries, and they can be used as a basis for other speed-up techniques (e.g., for transit-node routing and highway-node routing).

Categories and Subject Descriptors: G.2.2 [Discrete Mathematics]: Graph Theory—*Graph algorithms*

General Terms: Algorithms, Experimentation

Additional Key Words and Phrases: Algorithm engineering, road network, route planning, speed-up technique

1. INTRODUCTION

Computing fastest routes in road networks from a given source to a given target location is one of the showpieces of real-world applications of algorithmics. Many people frequently use this functionality when planning trips with their cars. There are also many applications like logistic planning or traffic simulation that need to solve a huge number of shortest-path queries. In principle, we could use Dijkstra’s algorithm [Dijkstra 1959], but for large road networks this would be far too slow. Therefore, there is considerable interest in speed-up techniques for route planning. Most approaches, including ours, assume that the road network is *static* (i.e., it does not change).¹ Then, we can allow some preprocessing that generates auxiliary data that can be used to

¹More precisely, it changes so slowly that we can afford to rerun preprocessing from time to time.

accelerate subsequent queries from specific source nodes to specific target nodes. The preprocessing should be sufficiently fast to deal even with very large road networks, the auxiliary data should occupy only a moderate amount of space, and the queries should be as fast as possible.

1.1. Related Work

A detailed overview on shortest-path speed-up techniques can be found in Delling et al. [2009]. This article is an extension of the conference papers Sanders and Schultes [2005, 2006]. Its main role is to provide proofs and to give additional experiments that help understanding how the method works. Much of the material found here can also be found Schultes [2008].

Bidirectional Search. A classical speed-up technique is *bidirectional search*, which simultaneously searches forward from the source and backward from the target until the search frontiers meet. Many more advanced speed-up techniques use bidirectional search as an ingredient.

Goal Direction. Road networks allow effective goal-directed search using A^* search [Hart et al. 1968]: Lower bounds on the remaining distance to the target define a vertex potential that directs search toward the target. This approach was recently shown to be very effective if lower bounds are computed using precomputed shortest-path distances to a carefully selected set of about 20 Landmark nodes [Goldberg and Harrelson 2005; Goldberg and Werneck 2005] using the Triangle inequality (ALT).

The precomputed cluster distances (PCD) technique [Maue et al. 2006, 2009] also uses precomputed distances for goal-directed search, yielding speed-ups comparable to ALT but using less space. The network is partitioned into clusters and the shortest connection between any pair of clusters is precomputed. Then, during a query, upper and lower bounds can be derived that can be used to prune the search.

Another goal-directed approach is to precompute for each edge “signposts” that support the decision whether the target can possibly be reached on a shortest path via this edge. During a query, only promising edges have to be considered. Various instantiations of this general idea have been presented [Schulz et al. 1999; Wagner and Willhalm 2003; Lauther 2004, 2006; Köhler et al. 2005, 2006; Möhring et al. 2005, 2007]. While these methods exhibit good query performance, preprocessing times are quite large.

Separators. Perhaps the most well-known property of road networks is that they are almost planar (i.e, techniques developed for planar graphs will often also work for road networks). Queries accurate within a factor $(1 + \epsilon)$ can be answered in near-constant time using $O((n \log n)/\epsilon)$ space and preprocessing time [Thorup 2001]. Recently, this approach has been efficiently implemented and experimentally evaluated on a road network with 1 million nodes [Muller and Zachariassen 2007]. While the query times are very good (less than $20\mu s$ for $\epsilon = 0.01$), the preprocessing time and space consumption are quite high (2.5 hours and 2GB, respectively). Using $O(n \log^3 n)$ space and preprocessing time, query time $O(\sqrt{n} \log n)$ can be achieved [Fakcharoenphol and Rao 2001] for directed planar graphs without negative cycles.

Another previous approach is the *separator-based multilevel method* [Schulz et al. 1999, 2002]. The idea is to use a set of nodes V_1 whose removal partitions the graph $G = G_0$ into small components. Then, consider the *overlay graph* $G_1 = (V_1, E_1)$ where edges in E_1 are *shortcuts* corresponding to shortest paths in G that do not have inner nodes that belong to V_1 . Routing can now be restricted to G_1 and the components containing s and t , respectively. This process can be iterated yielding a multilevel method. A limitation of this approach is that the graphs at higher levels become

much more dense than the input graphs, thus limiting the benefits gained from the hierarchy. Also, computing small separators can become quite costly for large graphs.

Reach-Based Routing/REAL. Let $R(v) := \max_{s,t \in V} R_{st}(v)$ denote the *reach* of node v , where $R_{st}(v) := \min(d(s, v), d(v, t))$. Gutman [2004] observed that a shortest-path search can be stopped at nodes with a reach too small to get to source or target from there. Goldberg et al. [2006, 2007] have considerably strengthened this approach by introducing various improvements, in particular a combination with ALT, yielding the REAL algorithm. Its query performance is similar to our highway hierarchies, while the preprocessing times are usually worse; a comparison can be found in Section 6.7.

Heuristics. In the last decades, commercial navigation systems were developed that had to handle ever more detailed descriptions of road networks on rather low-powered processors. Vendors resorted to heuristics still used today that do not give any performance guarantees: A^* search with estimates on the distance to the target rather than lower bounds or heuristic hierarchical approaches [Ishikawa et al. 1991; Jagadeesh et al. 2002].

1.2. Our Contributions

Our *exact* highway hierarchies (first published in Sanders and Schultes [2005, 2006]) are inspired by *heuristic* hierarchical approaches. It is a bidirectional speed-up technique. While the search is inside some local area around source or target, all roads of the network are considered. Outside these areas, however, the search is restricted to “important” roads. This general idea can be iterated and applied to a hierarchy consisting of several levels. The crucial point is the definition of “important streets.” In previous heuristic variants, this definition is based on a classification of the streets according to their type (motorway, national road, regional road, etc.). Such a classification requires manual tuning of the data and a delicate trade-off between speed and suboptimality of the computed routes. In our exact variant, however, nodes and edges are classified fully automatically in a preprocessing step in such a way that all shortest paths are preserved. By this means, we gain not only exactness, but also greater speed, since we can build high-performance hierarchies consisting of many levels without worrying about the quality of the results.

In the preprocessing phase, we alternate between two procedures: edge reduction and node reduction. *Edge reduction* removes nonhighway edges, that is, edges that only appear on shortest paths close to source or target. More specifically, every node v has a neighborhood radius $r(v)$ we are free to choose. An edge (u, v) is a highway edge if it belongs to some shortest path from a node s to a node t such that (u, v) is neither fully contained in the neighborhood of s nor in the neighborhood of t , that is, $d(s, v) > r(s)$ and $d(u, t) > r(t)$. In all our experiments, neighborhood radii are chosen such that each neighborhood contains a certain number H of nodes. H is a tuning parameter that can be used to control the rate at which the network shrinks.

Node reduction (also called *contraction*) removes low-degree nodes by bypassing them with newly introduced shortcut edges. In particular, all nodes of degree 1 and 2 are removed by this process.

The query algorithm is very similar to bidirectional Dijkstra search with the difference that certain edges need not be expanded when the search is sufficiently far from source or target. Highway hierarchies are the first speed-up technique that was able to handle the largest available road networks giving query times measured in milliseconds. There are two main reasons for this success: Under the previously described reduction routines, the road network shrinks in a geometric fashion from level to level and remains sparse and near planar, that is, levels of the highway hierarchy are in

some sense *self-similar*. The other key property is that preprocessing can be done using limited local searches starting from each node. Preprocessing is also the most nontrivial aspect of highway hierarchies. In particular, long edges (e.g., long-distance ferry connections) make simple-minded approaches far too slow. Instead, we use fast heuristics that compute a superset of the set of highway edges.

Some further optimisations allow to drop the average query times below 1ms on a 2.0GHz machine—even for a road network with more than 30 million nodes. One of these optimizations is an all-pairs distance table that we precompute for the topmost level L so that forward and backward search can be stopped as soon as all entrance points to level L have been found. Then, the remaining gap can be bridged by performing a moderate number of simple table look-ups.

We cannot give a general worst-case bound better than Dijkstra’s. So far, this drawback applies to all other exact speed-up techniques, where an implementation is available, as well. However, in contrast to most of them, we can provide *per-instance worst-case guarantees*, that is, for a given graph, we can determine an upper bound for the search space size of any possible point-to-point query performing only a linear number of unidirectional highway queries. These upper bounds are only a factor around three away from the average search space sizes.

1.3. Subsequent Work

Various other speed-up techniques in some way build on highway hierarchies. Goldberg et al. adopted the introduction of shortcuts in order to improve both preprocessing and query times of the REAL algorithm. There is a many-to-many variant [Knopp et al. 2007] and a combination with ALT [Delling et al. 2006]. The same applies to highway-node routing [Schultes and Sanders 2007], a more recent approach that can be used to handle dynamic scenarios (e.g., traffic jams). Contraction hierarchies [Geisberger et al. 2008, 2012] can alternatively be viewed as an extreme variant of highway-node routing with one level for each node or as an extreme variant of highway hierarchies where we use only (a sophisticated form of) contraction. An even faster routing technique, transit-node routing [Bast et al. 2007] can be viewed as a further development of the distance table optimization described in Section 5.3. Furthermore, its most efficient implementations rely on a hierarchical routing technique such as highway hierarchies for fast preprocessing. The currently fastest method [Abraham et al. 2010] can be viewed as an application of the many-to-many technique [Knopp et al. 2007] to *all* nodes of the network: Preprocessing stores forward and backward search spaces of all nodes. A query then intersects the search spaces of source and destination and selects the best of the candidates defined by the nodes in the intersection.

An alternative, heuristic approach to dealing with dynamic scenarios, which is based on highway hierarchies as well, has been developed by Nannicini et al. [2010]. Several recent techniques for combining hierarchical techniques with goal directed techniques [Bauer et al. 2008; Bauer and Delling 2008] use contraction as a crucial ingredient. The same holds for recent techniques that also work for time-dependent travel times [Delling 2008; Batz et al. 2009].

Several of the experimental techniques first used for highway hierarchies are now routinely used in many works on route planning: the road networks used, the technique for evaluating local queries used in Section 6.4, and the worst-case upper bounds described in Section 6.6.

1.4. Outline

After beginning with some preliminaries in Section 2, we formally define the *highway hierarchy* of a given graph in Section 3. Then, Section 4 deals with both procedures of the preprocessing phase, the edge reduction (i.e., the *construction* of a highway

network) and the node reduction (i.e., the *contraction* of a highway network). The basic query algorithm is introduced in Section 5. Furthermore, several optimizations are presented and some advanced topics, such as outputting complete path descriptions and dealing with turning restrictions, are discussed. In Section 6, we present a wide range of experimental results, dealing with various real-world road networks, parameter settings, and scenarios of application. We do not only give average query times, but also a detailed analysis of queries with different degrees of difficulty, per-instance worst-case upper bounds, and comparisons to other speed-up techniques.

2. PRELIMINARIES

Graphs and Paths. We expect a *directed* graph $G = (V, E)$ with n nodes and m edges (u, v) with *nonnegative* weights $w(u, v)$ as input. The *length* $w(P)$ of a path P is the sum of the weights of the edges that belong to P . $P^* = \langle s, \dots, t \rangle$ is a *shortest path* if there is no path P' from s to t such that $w(P') < w(P^*)$. The *distance* $d(s, t)$ between s and t is the length of a shortest path from s to t or ∞ if there is no path from s to t . If $P = \langle s, \dots, s', u_1, u_2, \dots, u_k, t', \dots, t \rangle$ is a path from s to t , then $P|_{s' \rightarrow t'} = \langle s', u_1, u_2, \dots, u_k, t' \rangle$ denotes the *subpath* of P from s' to t' . We use $u \prec_P v$ to denote that a node u precedes² a node v on a path $P = \langle \dots, u, \dots, v, \dots \rangle$; we just write $u \prec v$ if the path P that is referred to is clear from the context.

Dijkstra's Algorithm. Dijkstra's algorithm [Dijkstra 1959] can be used to solve the *single-source shortest-path (SSSP) problem*, that is, to compute the shortest paths from a single source node s to all other nodes in a given graph. It is covered by virtually any textbook on algorithms (e.g. Cormen et al. [2001] and Skiena [1998]), so that we confine ourselves to introducing our terminology: Starting with the source node s as root, Dijkstra's algorithm grows a *shortest-path tree* that contains shortest paths from s to all other nodes. During this process, each node of the graph is *unreached*, *reached*, or *settled*. A node that already belongs to the tree is *settled*. If a node u is settled, a shortest path P^* from s to u has been found and the distance $d(s, u) = w(P^*)$ is known. A node that is adjacent to a settled node is *reached*. Note that a settled node is also reached. If a node u is reached, a path P from s to u , which might not be the shortest one, has been found and a *tentative distance* $\delta(u) = w(P)$ is known. A node u that is not reached is *unreached*; for such a node, we have $\delta(u) = \infty$.

In case the shortest paths in a graph are not unique, Dijkstra's algorithm can be easily modified to determine *all* shortest paths between s and any node $u \in V$. This means that not a shortest-path tree is grown, but a shortest-path *directed acyclic graph* (DAG).

A bidirectional version of Dijkstra's algorithm can be used to find a shortest path from a given node s to a given node t . Two Dijkstra searches are executed in parallel: One searches from the source node s in the original graph $G = (V, E)$, also called *forward graph* and denoted as $\vec{G} = (V, \vec{E})$; another searches from the target node t backward (i.e., it searches in the *reverse graph* $\overleftarrow{G} = (V, \overleftarrow{E})$, $\overleftarrow{E} := \{(v, u) \mid (u, v) \in E\}$). The reverse graph \overleftarrow{G} is also called *backward graph*. When for the first time a node becomes settled for both searches, a shortest path from s to t has been found.

3. HIGHWAY HIERARCHY

A *highway hierarchy* of a graph G consists of several levels $G_0, G_1, G_2, \dots, G_L$, where the number of levels $L+1$ is given. We will now provide an inductive definition of the levels.

—Base case (G'_0, G_0). Level 0 ($G_0 = (V_0, E_0)$) corresponds to the original graph G ; furthermore, we define $G'_0 := G_0$.

²This does not necessarily mean that u is the direct predecessor of v .

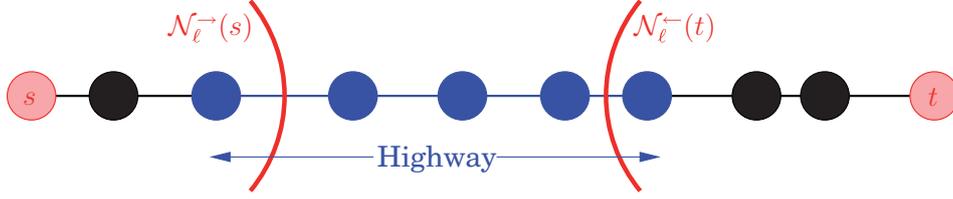


Fig. 1. A shortest path from a node s to a node t . Edges that leave the neighborhood of s or t and edges that are completely outside the neighborhoods of s and t are highway edges.

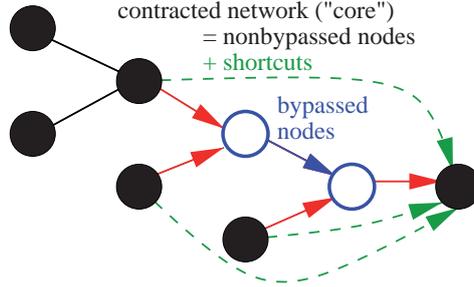


Fig. 2. The core of a highway network consists of the subgraph induced by the set of nonbypassed nodes and additional shortcut edges.

- First step ($G'_\ell \rightarrow G'_{\ell+1}$, $0 \leq \ell < L$). For given *neighborhood radii*, we will define the *highway network* $G'_{\ell+1}$ of a graph G'_ℓ .
- Second step ($G_\ell \rightarrow G'_\ell$, $1 \leq \ell \leq L$). For a given set $B_\ell \subseteq V_\ell$ of *bypassable* nodes, we will define the *core* G'_ℓ of level ℓ .

First Step (Highway Network). For each node u , we choose nonnegative *neighborhood radii* $r_\ell^{\rightarrow}(u)$ and $r_\ell^{\leftarrow}(u)$ for the forward and backward graph, respectively. To avoid some case distinctions, we set $r_\ell^{\rightarrow}(u)$ and $r_\ell^{\leftarrow}(u)$ to infinity for $u \notin V'_\ell$ (**R**adius **P**roperty **R**1) and for $\ell = L$ (**R**2). In all other cases, neighborhood radii have to be $\neq \infty$ (**R**3).

The level- ℓ *neighbourhood* of a node $u \in V'_\ell$ is $\mathcal{N}_\ell^{\rightarrow}(u) := \{v \in V'_\ell \mid d_\ell(u, v) \leq r_\ell^{\rightarrow}(u)\}$ with respect to the forward graph and, analogously, $\mathcal{N}_\ell^{\leftarrow}(u) := \{v \in V'_\ell \mid d_\ell^{\leftarrow}(u, v) \leq r_\ell^{\leftarrow}(u)\}$ with respect to the backward graph, where $d_\ell(u, v)$ denotes the distance from u to v in the forward graph G_ℓ and $d_\ell^{\leftarrow}(u, v) := d_\ell(v, u)$ in the backward graph \overleftarrow{G}_ℓ .

The *highway network* $G_{\ell+1} = (V_{\ell+1}, E_{\ell+1})$ of a graph G'_ℓ is defined by the set $E_{\ell+1}$ of *highway edges*: An edge $(u, v) \in E'_\ell$ belongs to $E_{\ell+1}$ if and only if there are nodes $s, t \in V'_\ell$ such that the edge (u, v) appears in some shortest path $\langle s, \dots, u, v, \dots, t \rangle$ from s to t in G'_ℓ with the property that $v \notin \mathcal{N}_\ell^{\rightarrow}(s)$ and $u \notin \mathcal{N}_\ell^{\leftarrow}(t)$. Figure 1 gives an example. The set $V_{\ell+1}$ is the set of nodes in V'_ℓ , which are adjacent to some edge in $E_{\ell+1}$.

Second Step (Core). For a given set $B_\ell \subseteq V_\ell$ of *bypassable* nodes, we define the set S_ℓ of *shortcut edges* that bypass the nodes in B_ℓ : For each path $P = \langle u, b_1, b_2, \dots, b_k, v \rangle$ with $u, v \in V_\ell \setminus B_\ell$ and $b_i \in B_\ell$, $1 \leq i \leq k$, the set S_ℓ contains an edge (u, v) with $w(u, v) = w(P)$. The *core* $G'_\ell = (V'_\ell, E'_\ell)$ of level ℓ is defined in the following way: $V'_\ell := V_\ell \setminus B_\ell$ and $E'_\ell := (E_\ell \cap (V'_\ell \times V'_\ell)) \cup S_\ell$. This definition is illustrated in Figure 2. Removing all core nodes from G_ℓ yields *connected components of bypassed nodes*.³

³Note that we do not check whether shortcuts are actually shortest paths. This is a simple solution, it speeds up individual contraction steps, and suboptimal shortcuts will not be promoted to the subsequent

The *level* $\ell(e)$ of an edge e is $\max\{\ell \mid e \in E_\ell \cup S_\ell\}$. For an edge (u, v) , we usually write just $\ell(u, v)$ instead of $\ell((u, v))$. The highway hierarchy can be interpreted as a single graph $\mathcal{G} := (V, E \cup \bigcup_{i=1}^L S_i)$, where each node and each edge has additional information on its membership in the various sets $V_\ell, V'_\ell, B_\ell, E_\ell, E'_\ell$, and S_ℓ .

4. CONSTRUCTION

4.1. Computing the Highway Network

Neighborhood Radii. Let us fix any deterministic rule that decides which element Dijkstra’s algorithm removes from the priority queue in the case that there is more than one queued element with the smallest key. Then, during a Dijkstra search from a given node u , all nodes are settled in a fixed order. The *Dijkstra rank* $\text{rk}_u(v)$ of a node v is the rank of v with regard to this order. u has Dijkstra rank $\text{rk}_u(u) = 0$, the closest neighbor v_1 of u has Dijkstra rank $\text{rk}_u(v_1) = 1$, and so on.

We suggest the following strategy to set the neighbourhood radii. For this paragraph, we interpret the graph G'_ℓ as an undirected graph, that is, a directed edge (u, v) is interpreted as an undirected edge $\{u, v\}$ even if the edge (v, u) does not exist in the directed graph. Let $d_\ell^{\leftrightarrow}(u, v)$ denote the distance between two nodes u and v in the undirected graph. For a given parameter H_ℓ , for any node $u \in V'_\ell$, we set $r_\ell^{\rightarrow}(u) := r_\ell^{\leftarrow}(u) := d_\ell^{\leftrightarrow}(u, v)$, where v is the node whose Dijkstra rank $\text{rk}_u(v)$ (with regard to the undirected graph) is H_ℓ . For any node $u \notin V'_\ell$, we set $r_\ell^{\rightarrow}(u) := r_\ell^{\leftarrow}(u) := \infty$ (to fulfill R1).

Originally, we wanted to apply the aforementioned strategy to the forward and backward graph separately in order to define the forward and backward radii, respectively. However, it turned out that using the same value for both forward and backward radii yields a similarly good performance, but needs only half the memory.

Fast Construction: Outline. Given a graph G'_ℓ , we want to construct a highway network $G_{\ell+1}$. We start with an empty set of highway edges $E_{\ell+1}$. For each node $s_0 \in V'_\ell$, two phases are performed: the forward construction of a partial shortest-path DAG B (containing all shortest paths from s_0 to any node $u \in B$) and the backward evaluation of B . The construction is done by an SSSP search from s_0 ; during the evaluation phase, paths from the leaves of B to the root s_0 are traversed and for each edge on these paths, it is decided whether to add it to $E_{\ell+1}$. The crucial part is the specification of an abort criterion for the SSSP search in order to restrict it to a “local search.”

Phase 1: Construction of a Partial Shortest-Path DAG. A Dijkstra search from s_0 is executed. In order to keep track of all shortest paths, for each node in the partial shortest-path DAG B , we manage a list of (tentative) parents: When an edge (u, v) is relaxed such that $d_\ell(s_0, u) + w(u, v) = \delta(v)$, then u is added to the list of tentative parents of v . During the search, a reached node is either in the state *active* or *passive*. The source node s_0 is active; each node that is reached for the first time (*insert*) and each reached node that is updated (*decreaseKey*) is set to active if and only if any of its tentative parents is active. When a node p is settled, we consider all shortest paths P' from s_0 to p , as depicted in Figure 3. The state of p is set to passive if

$$\forall \text{ shortest paths } P' = \langle s_0, \dots, p \rangle :$$

$$s_1 < p \wedge p \notin \mathcal{N}_\ell^{\rightarrow}(s_1) \wedge s_0 \notin \mathcal{N}_\ell^{\leftarrow}(p) \wedge |P' \cap \mathcal{N}_\ell^{\rightarrow}(s_1) \cap \mathcal{N}_\ell^{\leftarrow}(p)| \leq 1. \quad (1)$$

When no active unsettled node is left, the search is *aborted* and the growth of B stops.

An example for Phase 1 of the construction is given in Figure 4. The intuitive reason for s_1 (which is the first successor of s_0 on the path P') to appear in the abort criterion

levels anyway. However, contraction hierarchies [Geisberger et al. 2008] demonstrated that removing non-shortest path edges is surprisingly powerful, so in retrospect, this decision was a mistake.

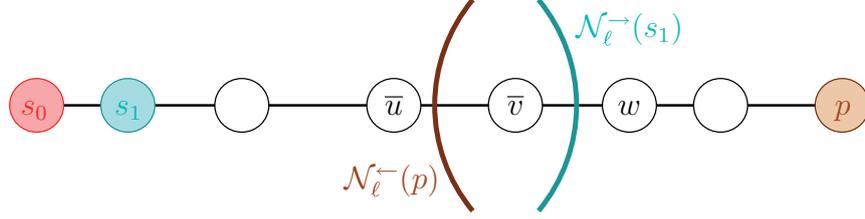


Fig. 3. Abort criterion.

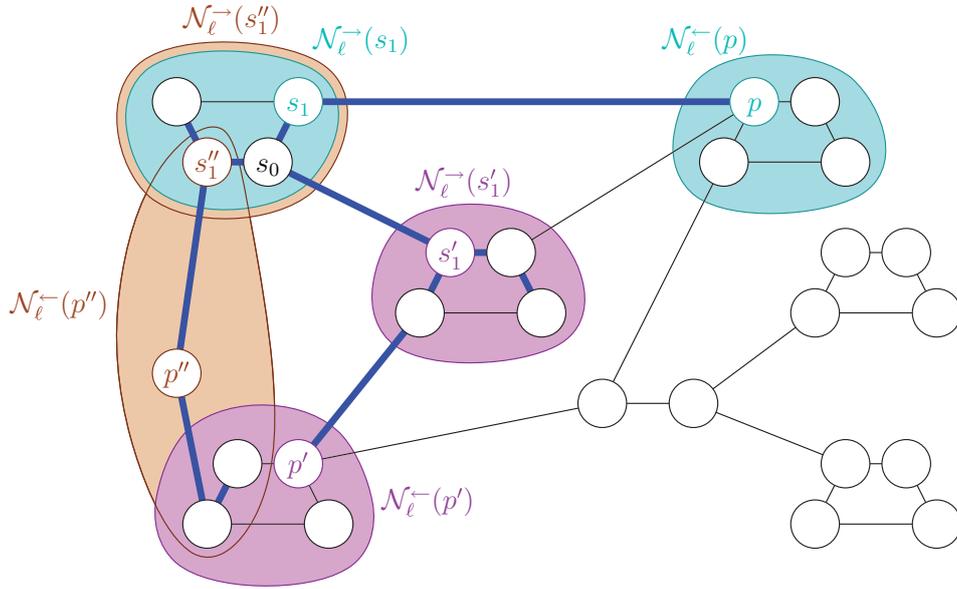


Fig. 4. An example of Phase 1 of the construction. The weight of an edge is the length of the line segment that represents the edge in this figure. The neighborhood size H_ℓ is 3. An SSSP search is performed from s_0 . The abort criterion applies three times, at nodes p , p' , and p'' . All edges that belong to s_0 's partial shortest-path tree are drawn as thick lines.

is the following: When we deactivate a node p during the search from s_0 , we decide to ignore everything that lies behind p . We are free to do this because the abort criterion ensures that s_1 can take “responsibility” for the things that lie behind p , that is, further important edges will be added during the search from s_1 . (Of course, s_1 will refer a part of its “responsibility” to its successor, and so on.)

Phase 2: Selection of the Highway Edges. During Phase 2, exactly all edges (u, v) are added to $E_{\ell+1}$ that lie on paths $\langle s_0, \dots, u, v, \dots, p \rangle$ in the partial shortest-path DAG B with the property that $v \notin N_\ell^rightarrow(s_0)$ and $u \notin N_\ell^leftarrow(p)$. The example from Figure 4 is continued in Figure 5.

THEOREM 4.1. *An edge $(u, v) \in E'_\ell$ is added to $E_{\ell+1}$ by the construction algorithm iff it belongs to some shortest path $P = \langle s, \dots, u, v, \dots, t \rangle$ and $v \notin N_\ell^rightarrow(s)$ and $u \notin N_\ell^leftarrow(t)$.*

PROOF. In this proof, we will refer to the following Neighborhood Property N1 that follows directly from the neighborhood definition: Consider a shortest path $\langle s, \dots, u, \dots, t \rangle$ in G'_ℓ . Then, $t \in N_\ell^rightarrow(s)$ implies $u \in N_\ell^rightarrow(s)$ and $s \in N_\ell^leftarrow(t)$ implies $u \in N_\ell^leftarrow(t)$.

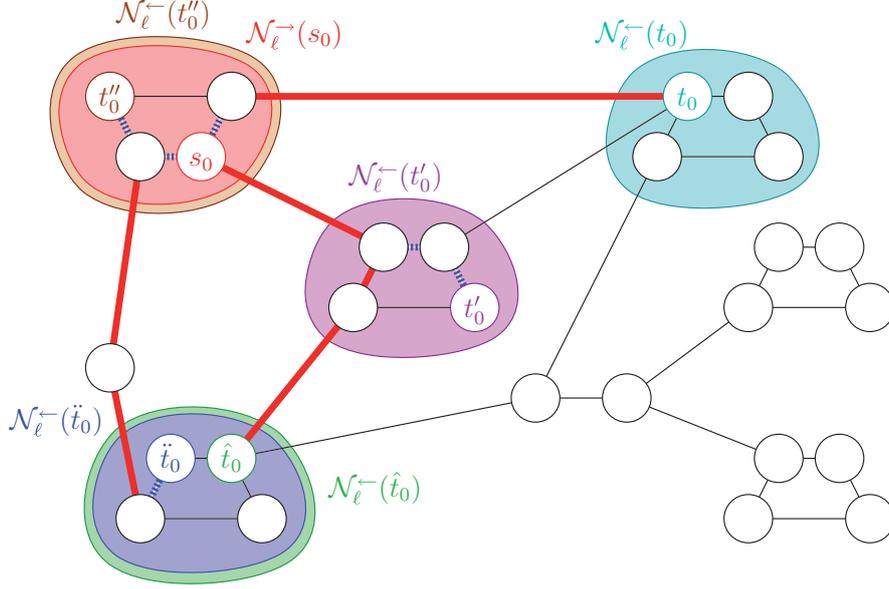


Fig. 5. An example of Phase 2 of the construction. s_0 's partial shortest-path tree (thick lines) has five leaves t_0 , t'_0 , t''_0 , \hat{t}_0 , and \check{t}_0 . The edges that are added to $E_{\ell+1}$ are represented as solid thick lines.

\Leftarrow) Consider the node s_0 on $P|_{s \rightarrow u}$ such that $v \notin \mathcal{N}_\ell^\rightarrow(s_0)$ and $d_\ell(s_0, v)$ is minimal. Such a node s_0 exists because the condition $v \notin \mathcal{N}_\ell^\rightarrow(s_0)$ is always fulfilled for $s_0 = s$. The direct successor of s_0 on P is denoted by s_1 . Note that $v \in \mathcal{N}_\ell^\rightarrow(s_1)$ [*]. We show that the edge (u, v) is added to $E_{\ell+1}$ when Phases 1 and 2 are executed from s_0 . Due to the specification of Phase 2, it is sufficient to prove that after Phase 1 has been completed, the partial shortest-path DAG⁴ B contains a node $p \in P|_{s_0 \rightarrow t}$ such that $v \leq p$ and $u \notin \mathcal{N}_\ell^\leftarrow(p)$.

If $t \in B$, this statement is obviously fulfilled for $p := t$, since $v \leq t$ and $u \notin \mathcal{N}_\ell^\leftarrow(t)$. Otherwise ($t \notin B$), the search is not continued from some node $t_0 < t$ on $P|_{s_0 \rightarrow t}$. We can conclude that t_0 is passive because, otherwise, its successor on $P|_{s_0 \rightarrow t}$ would adopt its active state and the search would not be aborted at that time. Since s_0 is active and t_0 is passive, either t_0 or one of its ancestors must have been switched from active to passive. Let p denote the first passive node on $P|_{s_0 \rightarrow t} = \langle s_0, s_1, \dots, p, \dots, t_0, \dots, t \rangle$. Due to the definition of the abort condition, we have $s_1 < p \wedge p \notin \mathcal{N}_\ell^\rightarrow(s_1) \wedge s_0 \notin \mathcal{N}_\ell^\leftarrow(p) \wedge |P' \cap \mathcal{N}_\ell^\rightarrow(s_1) \cap \mathcal{N}_\ell^\leftarrow(p)| \leq 1$ [***], where $P' = P|_{s_0 \rightarrow p}$. The fact that $v \in \mathcal{N}_\ell^\rightarrow(s_1)$ [see *] and $p \notin \mathcal{N}_\ell^\rightarrow(s_1)$ [see ***] imply $v < p$ due to N1. In order to obtain a contradiction, we assume $u \in \mathcal{N}_\ell^\leftarrow(p)$. Since $s_0 \notin \mathcal{N}_\ell^\leftarrow(p)$ [see ***], this implies $s_0 < u$ by N1. Hence, $s_1 \leq u$. Because $v \in \mathcal{N}_\ell^\rightarrow(s_1)$ [see *], we obtain $u \in \mathcal{N}_\ell^\rightarrow(s_1)$ due to N1. Similarly, we get $v \in \mathcal{N}_\ell^\leftarrow(p)$, since $v < p$ and $u \in \mathcal{N}_\ell^\leftarrow(p)$. Thus, $\{u, v\} \subseteq P' \cap \mathcal{N}_\ell^\rightarrow(s_1) \cap \mathcal{N}_\ell^\leftarrow(p)$. Therefore, $|P' \cap \mathcal{N}_\ell^\rightarrow(s_1) \cap \mathcal{N}_\ell^\leftarrow(p)| \geq 2$, which is a contradiction to [***]. We can conclude that $u \notin \mathcal{N}_\ell^\leftarrow(p)$.

\Rightarrow) Since each path $\langle s_0, \dots, u, v, \dots, p \rangle$ in B is a shortest path, the claim follows directly from the specification of Phase 2. \square

⁴For the shortest-path DAG we keep the terminology that is usually applied to shortest-path trees: A *parent* of a node in a shortest-path DAG is an adjacent node closer to the source node. Terms like *children* and *descendants* are used consistently.

Algorithmic Details: Phase 1. For an efficient implementation, we keep track of a *border distance* $b(x)$ and a *reference distance* $a(x)$ for each node x in B . Along a path P' , as depicted in Figure 3, we assign $b(x)$ the distance from the root to the border of the neighborhood of s_1 as soon as s_1 is settled. This value is passed to all successors on the path, which allows to determine the first node w outside $\mathcal{N}_{\ell}^{\rightarrow}(s_1)$, that is, its direct predecessor \bar{v} is the last node inside $\mathcal{N}_{\ell}^{\rightarrow}(s_1)$. In order to fulfill the abort condition, we have to make sure that \bar{v} is the only node on P' within $\mathcal{N}_{\ell}^{\rightarrow}(s_1) \cap \mathcal{N}_{\ell}^{\leftarrow}(p)$. Therefore, we want to check whether \bar{v} 's direct predecessor \bar{u} belongs to $\mathcal{N}_{\ell}^{\leftarrow}(p)$. To allow an easy check, we determine, store, and propagate the reference distance from s_0 to \bar{u} as soon as w is settled. Knowing the reference distance $d_{\ell}(s_0, \bar{u})$, the current distance $d_{\ell}(s_0, p)$ and p 's neighborhood radius $r_{\ell}^{\leftarrow}(p)$, checking $\bar{u} \notin \mathcal{N}_{\ell}^{\leftarrow}(p)$ is then straightforward. If there are several shortest paths from s_0 to some node x , we determine appropriate maxima of the involved border and reference distances.

More formally, for any node x in B , $\pi(x)$ denotes the set of parent nodes in B . To avoid some case distinctions, we set $\pi(s_0) := \{s_0\}$, that is, the root is its own parent. For the root s_0 , we set $b(s_0) := 0$ and $a(s_0) := \infty$. For any other node $x \neq s_0$, we define $b'(x) := d_{\ell}(s_0, x) + r_{\ell}^{\rightarrow}(x)$ if $s_0 \in \pi(x)$, and 0, otherwise; $b(x) := \max\{b'(x)\} \cup \{b(y) \mid y \in \pi(x)\}$; $a'(x) := \max\{a(y) \mid y \in \pi(x)\}$; and $a(x) := \max\{d_{\ell}(s_0, u) \mid y \in \pi(x) \wedge u \in \pi(y)\}$ if $a'(x) = \infty \wedge d_{\ell}(s_0, x) > b(x)$, and $a'(x)$, otherwise.

Then, we can easily check the following abort criterion at a settled node p :

$$a(p) + r_{\ell}^{\leftarrow}(p) < d_{\ell}(s_0, p). \quad (2)$$

LEMMA 4.2. (2) implies (1).

PROOF. We prove the contraposition “ \neg (1) implies \neg (2)”, that is, we assume that there is some shortest path P' from s_0 to p such that $p \leq s_1 \vee p \in \mathcal{N}_{\ell}^{\rightarrow}(s_1) \vee s_0 \in \mathcal{N}_{\ell}^{\leftarrow}(p) \vee |P' \cap \mathcal{N}_{\ell}^{\rightarrow}(s_1) \cap \mathcal{N}_{\ell}^{\leftarrow}(p)| \geq 2$ and show that $a(p) + r_{\ell}^{\leftarrow}(p) \geq d_{\ell}(s_0, p)$.

Case 1. $p \leq s_1$. If $p = s_0$, then $a(p) = \infty$, which yields \neg (2). Otherwise ($p = s_1$), $b(p) \geq d_{\ell}(s_0, p) + r_{\ell}^{\rightarrow}(p)$, $a'(p) = \infty$, and $a(p) = a'(p)$, since $d_{\ell}(s_0, p) \leq b(p)$, which implies \neg (2).

Case 2. $s_1 < p \wedge p \in \mathcal{N}_{\ell}^{\rightarrow}(s_1)$. Due to N1 (see proof of Theorem 4.1), we have $\forall x, s_1 \leq x \leq p : x \in \mathcal{N}_{\ell}^{\rightarrow}(s_1)$. Hence, $\forall x : d_{\ell}(s_0, x) \leq d_{\ell}(s_0, s_1) + r_{\ell}^{\rightarrow}(s_1) \leq b(x)$. By an inductive proof, we can show that $a(p) = \infty$, which yields \neg (2).

Case 3. $s_1 < p \wedge p \notin \mathcal{N}_{\ell}^{\rightarrow}(s_1) \wedge s_0 \in \mathcal{N}_{\ell}^{\leftarrow}(p)$. We have $d_{\ell}(s_0, p) \leq r_{\ell}^{\leftarrow}(p)$, which directly implies \neg (2).

Case 4. $s_1 < p \wedge p \notin \mathcal{N}_{\ell}^{\rightarrow}(s_1) \wedge s_0 \notin \mathcal{N}_{\ell}^{\leftarrow}(p) \wedge |P' \cap \mathcal{N}_{\ell}^{\rightarrow}(s_1) \cap \mathcal{N}_{\ell}^{\leftarrow}(p)| \geq 2$. The assumption of Case 4 implies that there are two nodes \bar{u} and \bar{v} , $s_1 \leq \bar{u} < \bar{v} \leq p$, that belong to $P' \cap \mathcal{N}_{\ell}^{\rightarrow}(s_1) \cap \mathcal{N}_{\ell}^{\leftarrow}(p)$. If $a(p) = \infty$, we directly have \neg (2). Otherwise, there has to be some node w on P' such that $a'(w) = \infty \wedge d_{\ell}(s_0, w) > b(w)$. Obviously, $w \neq s_0$. Consider such a node w that maximizes $d_{\ell}(s_0, w)$, that is, for all nodes $x > w$ the aforementioned condition does not hold, which implies $a(x) = a'(x) \geq a(w)$. In particular, $a(p) \geq a(w)$. We have $b(w) \geq d_{\ell}(s_0, s_1) + r_{\ell}^{\rightarrow}(s_1)$. We can conclude that $d_{\ell}(s_0, w) > d_{\ell}(s_0, s_1) + r_{\ell}^{\rightarrow}(s_1)$ and, thus, $w \notin \mathcal{N}_{\ell}^{\rightarrow}(s_1)$. We obtain, by N1, $\bar{u} < \bar{v} < w$. Hence, $a(w) \geq d_{\ell}(s_0, \bar{u})$, which implies $a(p) \geq d_{\ell}(s_0, \bar{u})$. Furthermore, since $\bar{u} \in \mathcal{N}_{\ell}^{\leftarrow}(p)$, we have $r_{\ell}^{\leftarrow}(p) \geq d_{\ell}(\bar{u}, p)$. Adding up the last two inequalities yields $a(p) + r_{\ell}^{\leftarrow}(p) \geq d_{\ell}(s_0, p)$, which corresponds to \neg (2). \square

Algorithmic Details: Phase 2. For a node $u \in B$, we define $\mathcal{B}(u) := \{u\} \cup \{v \mid v \text{ is a descendant of } u \text{ in } B\}$ and the *slack* $\Delta(u) := \min_{w \in \mathcal{B}(u)} (r_{\ell}^{\leftarrow}(w) - d_{\ell}(u, w))$. For a leaf b , we have $\mathcal{B}(b) = \{b\}$ and $\Delta(b) = r_{\ell}^{\leftarrow}(b)$. The slack of an inner node u can be computed using only the slacks of its children $C(u)$: $\Delta(u) = \min(r_{\ell}^{\leftarrow}(u), \min_{c \in C(u)} \Delta_c(u))$, where $\Delta_c(u) := \Delta(c) - d_{\ell}(u, c)$. This leads to an equivalent, recursive definition.

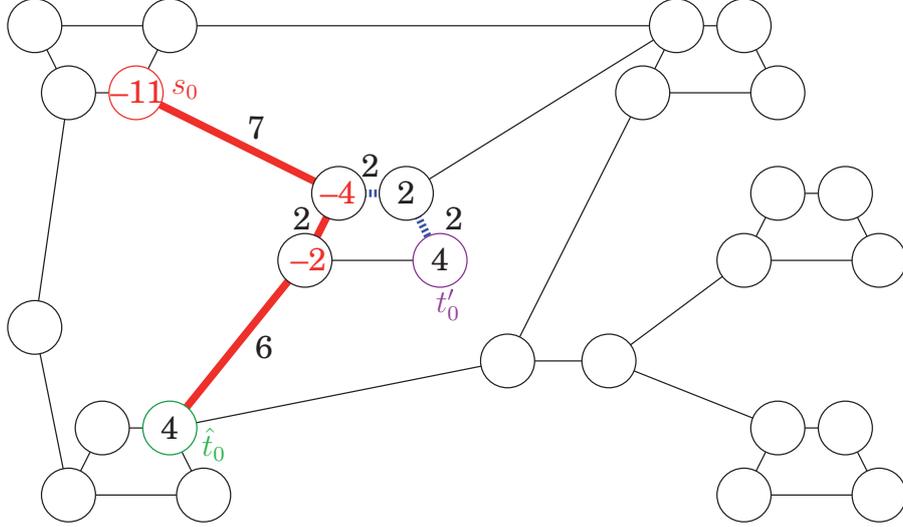


Fig. 6. An example of the *slack-based method* that realizes Phase 2 of the construction. The process is shown only for a part of the tree. As before, the weight of an edge is the length of the line that represents the edge in this figure. For the sake of transparency, the (rounded) weights are given explicitly for the relevant edges. Furthermore, the slacks of the involved nodes are given. Edges that are added to $E_{\ell+1}$ are solid edges that are not added dotted.

The tentative slacks $\hat{\Delta}(u)$ of all nodes u in B are set to $r_\ell^{\leftarrow}(u)$. We process all nodes in the reverse order as they were settled. This guarantees that all descendants of some node u have been processed before u is processed. We can stop as soon as a node $u \in \mathcal{N}_\ell^{\rightarrow}(s_0)$ is encountered. We maintain the invariant that the tentative slack $\hat{\Delta}(u)$ of an element u that is processed is equal to the actual slack $\Delta(u)$. When a node u is processed, for each parent p of u in B , we perform the following steps: compute $\Delta_u(p) = \Delta(u) - d_\ell(p, u)$; if $\Delta_u(p) < 0$, the edge (p, u) is added to $E_{\ell+1}$; if $\Delta_u(p) < \hat{\Delta}(p)$, the tentative slack $\hat{\Delta}(p)$ is set to $\Delta_u(p)$. Figure 6 gives an example.

THEOREM 4.3. *An edge (u, v) is added to $E_{\ell+1}$ by the previously described slack-based method if and only if it lies on a path $\langle s_0, \dots, u, v, \dots, p \rangle$ in the partial shortest-path DAG B with the property that $v \notin \mathcal{N}_\ell^{\rightarrow}(s_0)$ and $u \notin \mathcal{N}_\ell^{\leftarrow}(p)$.*

PROOF. \Leftarrow) From the definition of the slack of a node, it follows that

$$\Delta_v(u) = \Delta(v) - d_\ell(u, v) \leq r_\ell^{\leftarrow}(p) - d_\ell(v, p) - d_\ell(u, v) = r_\ell^{\leftarrow}(p) - d_\ell(u, p) < 0$$

because $u \notin \mathcal{N}_\ell^{\leftarrow}(p)$. Since $v \notin \mathcal{N}_\ell^{\rightarrow}(s_0)$, v is processed at some point. Then, $\Delta_v(u)$ is computed and, since it is negative, the edge (u, v) is added to $E_{\ell+1}$.

\Rightarrow) Only edges that belong to a path in B from s_0 to a node p are considered. The condition $v \notin \mathcal{N}_\ell^{\rightarrow}(s_0)$ is never violated because the traversal from the leaves to the root, and consequently, the addition of edges to $E_{\ell+1}$, is not continued when a node $v \in \mathcal{N}_\ell^{\rightarrow}(s_0)$ is encountered. If an edge (u, v) is added, the condition $\Delta_v(u) < 0$ is fulfilled. Hence, $\Delta(u) = \min_{w \in \mathcal{B}(u)} (r_\ell^{\leftarrow}(w) - d_\ell(u, w)) \leq \Delta_v(u) < 0$. Therefore, there is a node p such that $d_\ell(u, p) > r_\ell^{\leftarrow}(p)$, i.e., $u \notin \mathcal{N}_\ell^{\leftarrow}(p)$. \square

THEOREM 4.4. *Let V_B denote the set of nodes of s_0 's partial shortest-path DAG B . Let $G_B = (V_B, E_B)$ denote the subgraph of G_ℓ that is vertex induced by V_B . The complexity of Phases 1 and 2 started from s_0 is $T_{Dijkstra}(|G_B|)$.*

PROOF. The number of nodes of G_B is denoted by n' , the number of edges by m' . The complexity of Phase 1 corresponds to the complexity of an SSSP search in G_B started from s_0 , that is, $O(n' + m')$ outside the priority queue plus n' *insert* and n' *deleteMin* operations plus at most m' *decreaseKey* operations. During Phase 2, each node and each edge is processed at most once, that is, Phase 2 runs in $O(n' + m')$. \square

Speeding Up the Highway Network Construction. Even a single active endpoint of a long edge can cause a large search space during construction, although most nodes of the search space might already be passive.⁵ To face this undesirable effect, we declare an active node v to be a *maverick* if $d_\ell(s_0, v) > f \cdot r_\ell^{\rightarrow}(s_0)$, where f is a parameter. When all active nodes are mavericks, the search from passive nodes is no longer continued. This way, the construction process is accelerated and $E_{\ell+1}$ becomes a superset of the highway network. Hence, queries will be slower, but they will still compute exact shortest paths. The *maverick factor* f enables us to adjust the trade-off between construction and query time.

4.2. Computing the Core

In order to obtain the core of a highway network, we contract it, which yields several advantages. The search space during the queries gets smaller because bypassed nodes are not touched, and the construction process gets faster because the next iteration only deals with the nodes that have not been bypassed. Furthermore, a more effective contraction allows us to use smaller neighborhood sizes without compromising the shrinking of the highway networks. This improves both construction and query times. However, bypassing nodes involves the creation of shortcuts, that is, edges that represent the bypasses. Due to these shortcuts, the average degree of the graph is increased and the memory consumption grows. In particular, more edges have to be relaxed during the queries. Therefore, we have to carefully select nodes so that the benefits of bypassing them outweigh the drawbacks.

We give an iterative algorithm that combines the selection of the bypassable nodes B_ℓ with the creation of the corresponding shortcuts. We manage a stack that contains all nodes that have to be considered, initially all nodes from V_ℓ . As long as the stack is not empty, we deal with the topmost node u . We check the *bypassability criterion* $\#\text{shortcuts} \leq c \cdot (\deg_{\text{in}}(u) + \deg_{\text{out}}(u))$, which compares the number of shortcuts that would be created when u was bypassed with the sum of the in- and outdegree of u . The magnitude of the contraction is determined by the parameter c . If the criterion is fulfilled, the node is bypassed, that is, it is added to B_ℓ and the appropriate shortcuts are created. Note that the creation of the shortcuts alters the degree of the corresponding endpoints so that bypassing one node can influence the bypassability criterion of another node. Therefore, all adjacent nodes that have been removed from the stack earlier, that have been bypassed yet, and that are bypassable now are pushed on the stack once again.

THEOREM 4.5. *If $c < 2$, $|E'_\ell| = O(|V_\ell| + |E_\ell|)$.*

PROOF. If a node u is bypassed, the number of edges in the (tentative) core is increased by $\mathcal{D}_u := \#\text{shortcuts} - \deg_{\text{in}}(u) - \deg_{\text{out}}(u)$. (We have to subtract $\deg_{\text{in}}(u)$ and $\deg_{\text{out}}(u)$, since the edges incident to u no longer belong to the core.) Note that $\#\text{shortcuts} = \deg_{\text{in}}(u) \cdot \deg_{\text{out}}(u) - \deg_{\leftrightarrow}(u)$, where $\deg_{\leftrightarrow}(u)$ denotes the number of adjacent nodes v that are connected to u by both an edge (u, v) and an edge (v, u) . (We have to subtract

⁵A real-world example is a search that starts in Italy, at Genoa Harbor. We relax the long-distance ferry edge Genoa-Palermo so that, instead of a local area, we would search almost the entire country, since we cannot abort until the arrival point of the ferry has been settled and deactivated.

$\deg_{\leftrightarrow}(u)$ to account for the fact that a “shortcut” that would be a self-loop is not created.) We can conclude that $\mathcal{D}_u \leq \deg_{\text{in}}(u) \cdot \deg_{\text{out}}(u) - \deg_{\text{in}}(u) - \deg_{\text{out}}(u)$. If $\deg_{\text{in}}(u) \leq 1$ or $\deg_{\text{out}}(u) \leq 1$, we obtain $\mathcal{D}_u \leq 0$. Now, we deal with the case that $\deg_{\text{in}}(u) \geq 2$ and $\deg_{\text{out}}(u) \geq 2$. Since $\deg_{\leftrightarrow}(u) \leq \min(\deg_{\text{in}}(u), \deg_{\text{out}}(u))$, a node that fulfills the bypassability criterion also fulfills $\deg_{\text{in}}(u) \cdot \deg_{\text{out}}(u) \leq c \cdot (\deg_{\text{in}}(u) + \deg_{\text{out}}(u)) + \min(\deg_{\text{in}}(u), \deg_{\text{out}}(u))$. The inequality $x \cdot y \leq c(x + y) + \min(x, y)$ has only finitely many solutions (x, y) for $x, y \in \mathbb{N}, x, y \geq 2$ if $c \in \mathbb{R}$ is a constant less than 2. Consider the solution (x, y) that maximizes $k := x \cdot y$. If there is no solution, take $k := 0$. Note that k is a constant that only depends on the constant c . We can conclude that $\mathcal{D}_u \leq k$.

Each node from V_ℓ is bypassed at most once. For each bypassed node, the number of edges in the (tentative) core is increased by at most k . Therefore, $|E'_\ell| \leq k \cdot |V_\ell| + |E_\ell|$. \square

If we used $\#\text{shortcuts} \leq \max(\deg_{\text{in}}(u), \deg_{\text{out}}(u))$ as bypassability criterion, we would get a contraction that would be very similar to our earlier trees-and-lines method [Sanders and Schultes 2005]. The more general version presented earlier allows a more aggressive contraction by setting c appropriately.

Limiting Component Sizes. To reduce the observed maximum query time, we implement a limit on the number of hops a shortcut may represent. By this means, the sizes of the components of bypassed nodes are reduced—in particular, the first contraction step tended to create quite large components of bypassed nodes so that it took a long time to leave such a component when the search was started from within it.

5. QUERY

Our *highway query algorithm* is a modification of the bidirectional version of Dijkstra’s algorithm. We will see that in contrast to the construction, during the query we need not keep track of ambiguous shortest paths. Let us first assume that the search is not aborted when both search scopes meet. This matter is dealt with in Section 5.3. We only describe the modifications of the forward search, since forward and backward search are symmetric. In addition to the *distance* from the source, each node is associated with the search *level* and the *gap* to the “next applicable neighborhood border.” The search starts at the source node s in level 0. First, a local search in the neighborhood of s is performed, that is, the gap to the next border is set to the neighborhood radius of s in level 0. When a node v is settled, it adopts the gap of its parent u minus the length of the edge (u, v) . As long as we stay inside the current neighborhood, everything works as usual. However, if an edge (u, v) crosses the neighborhood border (i.e., the length of the edge is greater than the gap), we switch to a higher search level ℓ . The node u becomes an *entrance point* to the higher level. If the level of the edge (u, v) is less than the new search level ℓ , the edge is not relaxed—this is one of the two restrictions that cause the speed-up in comparison to Dijkstra’s algorithm (Restriction 1). Otherwise, the edge is relaxed: v adopts the new search level ℓ and the gap to the border of the neighborhood of u in level ℓ , since u is the corresponding entrance point to level ℓ .

We have to deal with the special case that an entrance point to level ℓ does not belong to the core of level ℓ . In this case, the search is continued inside a component of bypassed nodes until the level- ℓ core is entered, that is, a node $u \in V'_\ell$ is settled. At this point, u is assigned the gap to the border of the level- ℓ neighborhood of u . Note that before the core is entered (i.e., inside a component of bypassed nodes), the gap has been infinity (according to R1). To increase the speed-up, we introduce another restriction (Restriction 2): When a node $u \in V'_\ell$ is settled, all edges (u, v) that lead to a bypassed node $v \in B_\ell$ in search level ℓ are not relaxed, that is, once entered the core, we will never leave it again.

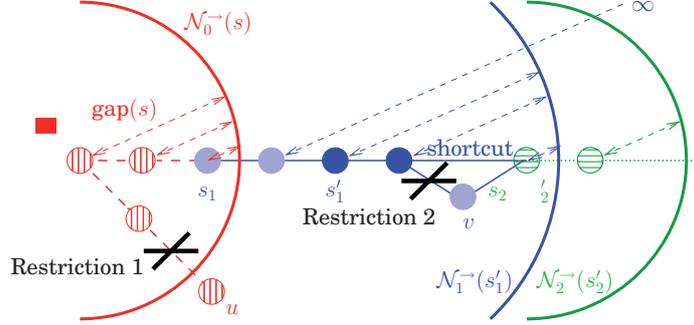


Fig. 7. A detailed example of a highway query. Only the forward search is depicted. Nodes in levels 0, 1, and 2 are vertically striped, solid, and horizontally striped, respectively. In level 1, dark shades represent core nodes, light shades bypassed nodes. Edges in level 0, 1, and 2 are dashed, solid, and dotted, respectively.

Figure 7 gives a detailed example of the forward search of a highway query. The search starts at node s . The gap of s is initialized to the distance from s to the border of the neighborhood of s in level 0. Within the neighborhood of s , the search process corresponds to a standard Dijkstra search. The edge that leads to u leaves the neighborhood. It is not relaxed due to Restriction 1 because the edge belongs only to level 0. In contrast, the edge that leaves s_1 is relaxed because its level allows to switch to level 1 in the search process. s_1 and its direct successor are bypassed nodes in level 1. Their neighborhoods are unbounded, that is, their neighborhood radii are infinity so that the gap is set to infinity as well. At s'_1 , we leave the component of bypassed nodes and enter the core of level 1. Now, the search is continued in the core of level 1 within the neighborhood of s'_1 . The gap is set appropriately. Note that the edge to v is not relaxed due to Restriction 2, since v is a bypassed node. Instead, the direct shortcut to s_2 is used. Here, we switch to level 2. In this case, we do not enter the next level through a component of bypassed nodes, but we get directly into the core. The search is continued in the core of level 2 within the neighborhood of s_2 , and so on.

Despite Restriction 1, we always find the optimal path, since the construction of the highway hierarchy guarantees that the levels of the edges that belong to the optimal path are sufficiently high so that these edges are not skipped. Restriction 2 does not invalidate the correctness of the algorithm, since we have introduced shortcuts that bypass the nodes that do not belong to the core. Hence, we can use these shortcuts instead of the original paths.

5.1. The Basic Algorithm

We use two priority queues \vec{Q} and \overleftarrow{Q} , one for the forward search and one for the backward search. For each search direction, a node u is associated with a triple $(\delta(u), \ell(u), \text{gap}(u))$, which we often call *key*. It consists of the (tentative) distance $\delta(u)$ from s (or t) to u , the search level $\ell(u)$, and the gap $\text{gap}(u)$ to the next applicable neighborhood border. Only the first component $\delta(u)$ is used to decide the priority within the queue.⁶ We use the remaining two components for a tie-breaking rule in the case that the same node is reached with two different keys $k := (\delta, \ell, \text{gap})$ and $k' := (\delta', \ell', \text{gap}')$ such that $\delta = \delta'$. Then, we prefer k to k' if and only if $\ell > \ell'$ or $\ell = \ell' \wedge \text{gap} < \text{gap}'$. Note that any other tie-breaking rule (or even omitting an explicit rule) will yield a correct

⁶If the search direction is not clear from the context, we will explicitly write $\vec{\delta}(u)$ and $\overleftarrow{\delta}(u)$ to distinguish between u 's priority in \vec{Q} and \overleftarrow{Q} .

```

: source node  $s$  and target node  $t$ 
: distance  $d(s, t)$ 

1  $d' := \infty$ ;
2 insert( $\overrightarrow{Q}, s, (0, 0, r_0^{\rightarrow}(s))$ ); insert( $\overleftarrow{Q}, t, (0, 0, r_0^{\leftarrow}(t))$ );
3 while ( $\overrightarrow{Q} \cup \overleftarrow{Q} \neq \emptyset$ ) do {
4     select direction  $\Leftarrow \in \{\rightarrow, \leftarrow\}$  such that  $\overleftarrow{Q} \neq \emptyset$ ;
5      $u := \text{deleteMin}(\overleftarrow{Q})$ ;
6     if  $u$  has been settled from both directions then  $d' := \min(d', \overrightarrow{\delta}(u) + \overleftarrow{\delta}(u))$ ;
7     if  $\text{gap}(u) \neq \infty$  then  $\text{gap}' := \text{gap}(u)$  else  $\text{gap}' := r_{\ell(u)}^{\overleftarrow{Q}}$ ;
8     foreach  $e = (u, v) \in \overleftarrow{E}$  do {
9         for ( $\ell := \ell(u)$ ,  $\text{gap} := \text{gap}'$ ;  $w(e) > \text{gap}$ ;
            $\ell^{++}$ ,  $\text{gap} := r_{\ell}^{\overleftarrow{Q}}(u)$ ); // go "upwards"
10        if  $\ell(e) < \ell$  then continue; // Restriction 1
11        if  $u \in V_{\ell}' \wedge v \in B_{\ell}$  then continue; // Restriction 2
12         $k := (\overrightarrow{\delta}(u) + w(e), \ell, \text{gap} - w(e))$ ;
13        if  $v$  has been reached then decreaseKey( $\overleftarrow{Q}, v, k$ ); else insert( $\overleftarrow{Q}, v, k$ );
14    }
15 }
16 return  $d'$ ;

```

Fig. 8. The highway query algorithm. Differences to the bidirectional version of Dijkstra’s algorithm are marked: Additional and modified lines have a framed line number; in modified lines, the modifications are underlined.

algorithm. However, the chosen rule is most aggressive and has a positive effect on the performance. Figure 8 contains the pseudocode of the highway query algorithm.

Remarks.

- Line 4. The correctness of the algorithm does not depend on the strategy that determines the order in which the forward and the backward searches are processed. However, the choice of the strategy can affect the running time in the case that an abort-on-success criterion is applied (see Section 5.3).
- Line 7. This line deals with the special case that the entrance point did not belong to the core when the current search level ℓ was entered, that is, the gap was set to infinity. In this case, the gap is set to $r_{\ell(u)}^{\overleftarrow{Q}}$. This is correct even if u does not belong to the core, either, because in this case the gap stays at infinity.
- Line 9. It might be necessary to go upward more than one level in a single step.
- Line 13. In the decreaseKey operation, the old key of v is only replaced by k if the above mentioned condition is fulfilled, that is, if (i) the tentative distance is improved or (ii) stays unchanged while the tie-breaking rule succeeds. In the latter case (ii), no priority queue operation is invoked, since the priority (the tentative distance) has not changed.⁷

Algorithmic Details. If we group the outgoing edges (u, v) of each node u by level, we can avoid looking at edges (u, v) in levels $\ell(u, v) < \ell(u)$ since Restriction 1 would always

⁷That way, we also avoid problems that otherwise could arise when an already settled node is reached once again via a zero weight edge.

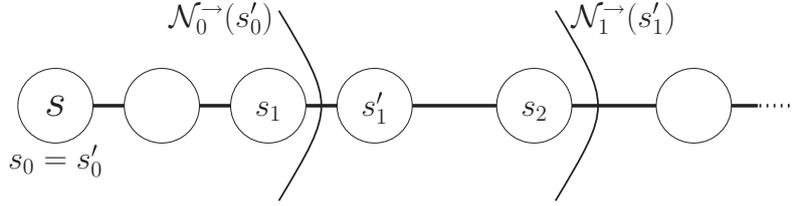


Fig. 9. Example for a forward labeling of a path P . The labels s_0 and s'_0 are set to s (base case). The node s_1 is the last neighbour of s'_0 (denoted by $\vec{\omega}_0^P(s'_0)$), the node s'_1 is the first level-1 core node (denoted by $\vec{\alpha}_1^P(s_1)$), s_2 is the last neighbour of s'_1 , and so on.

apply to them. We can do without explicitly testing Restriction 2 if all edges (u, v) with $k := \ell(u, v)$, $u \in V'_k$, and $v \in B_k$ have been downgraded to level $k - 1$. Then, the test of Restriction 1 also covers Restriction 2.

5.2. Proof of Correctness

Difficulties. Although the basic concepts (e.g., the definition of the highway network) and the algorithm are quite simple, the proof of correctness gets surprisingly complicated. The main reason for that is the fact that we cannot prove that *the* shortest path is found, since there might be several shortest paths of the same length. We could assume that the shortest paths in the input are unique or that the uniqueness can be guaranteed by adding small fractions to the edge weights, as it is done by other authors who face similar problems. However, the former would be too restrictive, since usually, in real-world road networks, there are at least a few ambiguous instances, and a reliable realization of the latter would be rather difficult. Furthermore, the introduction of shortcuts adds a lot of ambiguity even if it was not present in the input.

Therefore, if we pick any shortest path P to show that it is found by the query algorithm, it can happen that a node u on P is settled from another node than its predecessor on P . Of course, in this case, u will still be assigned the optimal distance from the source, but the search level and the distance to the next neighborhood border may be different than expected so that we have to adapt to the new situation.

Outline. We face the above-mentioned difficulties in the following way: First, we show that the algorithm terminates and deal with the special case that no path from the source to the target exists (Section A.1). Then, we introduce some definitions and concepts that will be useful in the main part of the correctness proof. In Section A.2, we define for a given path, a corresponding *contracted* path and an *expanded* path, where subpaths in the original graph are replaced by shortcuts or vice versa, respectively. In Section A.3, we first define the concepts of *last neighbor* and *first core node*, which, iteratively applied, lead to an *unidirectional labeling* of a given path. Figure 9 gives an example. Applying a forward and a backward labeling to the same path then allows the definition of a *meeting level* and a *meeting point* (Figure 10). The latter requires a case distinction, since the forward and backward labeling may either meet in some core or in some component of bypassed nodes. Finally, we introduce the term *highway path*, a path whose properties exactly comply with the two restrictions of the query algorithm. Figure 11 gives an example.

In Section A.4, we deal with the reachability along a highway path. Basically, we show that if the query has settled some node u on a highway path with the appropriate key, then u 's successor on that path can be reached from u with the appropriate key as well (Lemmas A.12 and A.13, which are proved using the auxiliary Lemma A.11). In

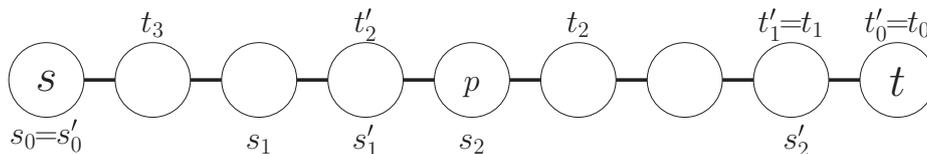


Fig. 10. Example for a forward and backward labeling (depicted below and above the nodes, respectively). The meeting level is 2, the meeting point is p .

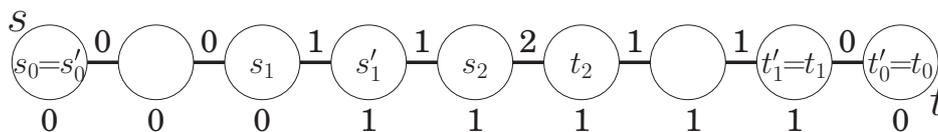


Fig. 11. Example for a highway path. Each edge belongs at least to the given level, each node at least to the given core level.

other words, if there is a highway path, the query can follow the path (at least if there was no ambiguity).

In Section A.5, we use all concepts and lemmas introduced in the preceding sections to conduct the actual correctness proof, where we also deal with ambiguous paths. The general idea is to say that at any point the query algorithm has some valid *state* consisting of a shortest s - t -path P and two nodes $u \leq \bar{u}$ that split P into three parts such that the first and the third part are paths in the forward and backward search tree, respectively, and the second part is a contracted path. For such a valid state, we can prove that any node on the first and third part has been settled with the appropriate key (Lemma A.16). Furthermore, we can show that P is a highway path (Lemma A.17).

When the algorithm is started, the nodes s and t are settled and some shortest s - t -path P in the original graph exists. (The special case that no s - t -path exists has already been dealt with.) Consequently, our *initial* state is composed of the contracted version of P and the nodes s and t , which makes it a valid state. A *final* state is a valid state where forward and backward search have met, that is, they have settled a common node $u = \bar{u}$. Originally, we wanted to show that a shortest path is found. Now, we see (in Lemma A.19) that it is sufficient to prove that a final state is reached.

We have already defined the meeting point p on a path. We fall back on this definition and intend to prove that forward and backward search meet at p . When we look at any valid nonfinal state, it is obvious that at least one search direction can proceed to get closer to p , that is, we have $u < p$ or $p < \bar{u}$ (Lemma A.21). We pick such a non-blocked search direction. Let us assume without loss of generality that we picked the forward direction. We know that u has been settled with the appropriate key and that P is an optimal highway path (Lemmas A.16 and A.17). Due to the “reachability along a highway path” (Lemmas A.12 and A.13), we can conclude that u ’s successor v can be reached with the appropriate key as well, in particular with the optimal distance from s . A node that can be reached with the optimal distance will also be settled at some point with the optimal distance. However, we cannot be sure that v is settled with u as its parent, since the shortest path from s to v might be ambiguous. At this point, the state concept gets handy: We just replace the subpath of P from s to v with the path in the search tree that actually has been taken yielding a path P^+ ; we obtain a new state that consists of P^+ and the nodes v and \bar{u} . We prove that the new state is valid (Lemma A.23).

Thus, we can show that from any valid nonfinal state another valid state is reached at some point. We also show in Lemma A.23 that we cannot get into some cycle of states, since in each step, the length of the middle part of the path is decreased. Hence, starting from the initial state, eventually a final state is reached so that a shortest path is found (Theorem A.24).

The actual proof can be found in Appendix A.

5.3. Optimisations

Rearranging Nodes. Similar to Goldberg et al. [2007], after the construction has been completed, we rearrange the nodes by core level, which improves locality for the search in higher levels and, thus, reduces the number of cache misses. In our experiments, this results in 14% to 22% smaller query time.

Speeding Up the Search in the Topmost Level. Let us assume that we have a distance table that contains, for any node pair $s, t \in V'_L$, the optimal distance $d_L(s, t)$. Such a table can be precomputed during the preprocessing phase by $|V'_L|$ SSSP searches in G'_L . Using the distance table, we do not have to search in level L . Instead, when we arrive at a node $u \in V'_L$ that leads to level L , we add u to the initially empty set \vec{T} or \overleftarrow{T} depending on the search direction; we do not relax the edge that leads to level L . After all entrance points have been encountered, we consider all pairs $(u, v) \in \vec{T} \times \overleftarrow{T}$ and compute the minimum path length $D := \vec{\delta}(u) + d_L(u, v) + \overleftarrow{\delta}(v)$. Then, the length of the shortest s - t -path is the minimum of D and the length d' of the tentative shortest path found so far (in case that the search scopes have already met in a level $< L$).

For the sake of a simple incorporation of this optimization into the highway query algorithm, we slightly revise the properties R1 and R2: we use two distinguishable values ∞_1 and ∞_2 that are larger than any real number and set $r_\ell^{\overleftarrow{=}}(u) := \infty_1$ for any ℓ and any node $u \notin V'_L$ (R1) and $r_L^{\overleftarrow{=}}(u) := \infty_2$ for any node $u \in V'_L$ (R2). Then, we just add two lines to Figure 8 and modify Line 16:

between Lines 7 and 8:

7a **if** $\text{gap}' \neq \infty_1 \wedge \ell(u) = L$ **then** $\{\overleftarrow{I} := \overleftarrow{I} \cup \{u\};$ **continue};**

between Lines 11 and 12:

11a **if** $\text{gap} \neq \infty_1 \wedge \ell = L \wedge \ell > \ell(u)$ **then** $\{\overleftarrow{I} := \overleftarrow{I} \cup \{u\};$ **continue};**

16 **return** $\min(\{d'\} \cup \{\vec{\delta}(u) + d_L(u, v) + \overleftarrow{\delta}(v) \mid u \in \vec{T}, v \in \overleftarrow{T}\});$

In Section A.6, we show that our proof of correctness still holds when the distance table optimisation is applied.

Abort on Success. In the bidirectional version of Dijkstra's algorithm, we can abort the search as soon as both search scopes meet. Unfortunately, this would be incorrect for our highway query algorithm. Therefore, we use a more conservative criterion: After a tentative shortest path P' has been encountered (i.e., after both search scopes have met), the forward (backward) search is not continued if the minimum element u of the forward (backward) queue has a key $\delta(u) \geq w(P')$. Obviously, the correctness of the algorithm is not invalidated by this abort criterion. In Sanders and Schultes [2005], we tried using more sophisticated criteria in order to reduce the search space. However, it turned out that this simple criterion, since it can be evaluated so efficiently, yields better query times despite a somewhat larger search space. Note that when the distance table optimization is used and random queries are performed, Our simple abort criterion is very close to an optimal criterion even with respect to the search

space size: Our experiments indicate that less than 1% of the search space is visited after the first meeting of forward and backward search.

5.4. Outputting Complete Path Descriptions

The highway query algorithm in Figure 8 only computes the distance from s to t . In order to determine the actual shortest path, we need to store pointers from each node to its parent in the search tree. Note that the algorithm could be easily modified to compute all shortest paths between s and t by just storing more than one parent pointer in case of ambiguities. However, subsequently, we only deal with a single shortest path.

We face two problems in order to determine a complete description of the shortest path: (i) We have to bridge the gap between the forward and backward topmost core entrance points (in case that the distance table optimisation is used), and (ii) we have to expand the shortcuts on the computed path to obtain the corresponding subpaths in the original graph.

Problem (i) can be solved using a simple algorithm: We start with the forward core entrance point u . As long as the backward entrance point v has not been reached, we consider all outgoing edges (u, w) in the topmost core and check whether $d_L(u, w) + d_L(w, v) = d_L(u, v)$; we pick an edge (u, w) that fulfils the equation, and we set $u := w$. The check can be performed using the distance table. It allows us to greedily determine the next hop that leads to the backward entrance point.

Problem (ii) can be solved without using any extra data (Variant 1). For each shortcut $(u, v) \in S_\ell$ on the shortest path, we perform a search from u to v in order to determine the represented path in G_ℓ . This search can be accelerated by using the knowledge that the first edge of the path enters a component C of bypassed nodes, the last edge leads to v , and all other edges are situated within the component C . The represented path in G_ℓ may contain shortcuts from sets $S_k, k < \ell$, which are expanded recursively. In the end, we obtain the represented path from u to v in the original graph.

However, if a fast output routine is required, it is necessary to spend some additional space to accelerate the unpacking process. We use a rather sophisticated data structure to represent unpacking information for the shortcuts in a space-efficient way (Variant 2). In particular, we do not store a sequence of node IDs that describe a path that corresponds to a shortcut, but we store only *hop indices*: For each edge (u, v) on the path that should be represented, we store its rank within the ordered group of edges that leave u . Since in most cases the degree of a node is very small, these hop indices can be stored using only a few bits. The unpacked shortcuts are stored in a recursive way (e.g., the description of a level-2 shortcut may contain several level-1 shortcuts). Accordingly, the unpacking procedure works recursively.

To obtain a further speed-up, we have a variant of the unpacking data structures (Variant 3) that caches the complete descriptions—without recursions—of all shortcuts that belong to the topmost level, that is, for these important shortcuts that are frequently used, we do not have to use a recursive unpacking procedure, but we can just append the corresponding subpath to the resulting path.

5.5. Turning Restrictions

A turning restriction (in its simplest and most common form) is expressed as an edge pair $((u, v), (v, w))$: The edge (v, w) must not be traversed if the node v has been reached via the edge (u, v) . Dealing with turning restrictions is a well-studied problem [Schmid 2000; Müller 2005]. In principle, there are two basic approaches: modifying the query algorithm or modeling the restrictions into the graph, which introduces additional artificial nodes and edges at affected road junctions. The latter technique can be applied irrespective of the used query algorithm.

In case of highway hierarchies, we expect that modeling turning restrictions into the graph only slightly deteriorates the performance, since the artificial nodes usually have a very small degree so that most of them get bypassed in the very first contraction step. Furthermore, turning restrictions are often encountered at local streets that are not promoted to high levels of the hierarchy so that the negative impact is bounded to the lower levels. With respect to memory consumption, it is important to note that after the preprocessing has been completed, artificial nodes and edges at road junctions that only belong to level 0 can be abandoned provided that the query algorithm (which in level 0 just corresponds to Dijkstra’s algorithm) is modified appropriately to handle turning restrictions.

6. EXPERIMENTS

Apart from Section 6.8, all experimental results refer to the scenario where we only want to compute the shortest-path length between two nodes without outputting the actual route. Turning restrictions are exclusively handled in Section 6.9.

6.1. Implementation

We implemented highway hierarchies in C++, using the C++ Standard Template Library and making extensive use of *generic programming* techniques using C++’s template class mechanism. As graph data structure, we use our own implementation of an *adjacency array* extended by an additional layer that contains level-specific data for each node and level that the node belongs to. We use 32 bits to store edge weights and path lengths. *Binary heaps* are used as priority queues. Note that in case of road networks, only a comparatively small number of *decreaseKey*-operations is performed. Furthermore, the number of nodes that are in the priority queue at the same time is very small in case of highway hierarchies (usually less than 100 nodes). Therefore, using a more sophisticated priority queue implementation is not likely to increase the performance significantly. For more details on the implementation, see Appendix B.

6.2. Environment and Instances

The experiments were done on one core of a single AMD Opteron Processor 270 clocked at 2.0GHz with 8GB main memory and $2 \times 1\text{MB}$ L2 cache, running SuSE Linux 10.0 (kernel 2.6.13). The program was compiled by the GNU C++ compiler 4.0.2 using optimization level 3.

We deal with the road networks of Western Europe⁸ and of the United States (without Hawaii) and Canada. Both networks have been made available for scientific use by the company PTV AG. The original graphs contain for each edge a length and a road category (e.g., motorway, national road, regional road, urban street). We assign average speeds to the road categories⁹ compute for each edge the average travel time, and use it as weight. In addition, we perform experiments on a publicly available version of the U.S. road network (without Alaska and Hawaii) that was obtained from the TIGER/Line Files [U.S. Census Bureau, Washington, DC 2002]. However, in contrast to the PTV data, the TIGER graph is undirected, planarized, and distinguishes only between four road categories (40, 60, 80, 100km/h). In fact 91% of all roads belong to the slowest category so that they cannot be discriminated.

Table I summarizes important properties of the used road networks and the key results of the experiments.

⁸Austria, Belgium, Denmark, France, Germany, Italy, Luxembourg, the Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, and the UK.

⁹For Europe: 10, 20, ..., 130km/h; for US/CAN: 16, 24, 32, 40, 56, 64, 72, 80, 88, 96, 104 and 112km/h.

Table I. Overview of the Used Road Networks and Key Results

		Europe	US/CAN	US (Tiger)
INPUT	#nodes	18,029,721	18,741,705	24,278,285
	#directed edges	42,199,587	47,244,849	58,213,192
	#road categories	13	13	4
PARAM.	average speeds [km/h]	10–130	16–112	40–100
	H	30	40	40
PREPROC.	CPU time [min]	13	17	15
	\emptyset overhead/node [byte]	48	46	34
QUERY	CPU time [ms]	0.61	0.83	0.67
	#settled nodes	709	871	925
	#relaxed edges	2,531	3,376	3,823
	speed-up (CPU time)	9,935	7,259	9,303
	speed-up (#settled nodes)	12,715	10,750	12,889
	worst case (#settled nodes)	2,388	2,428	2,505

Note: “ \emptyset overhead/node” accounts for the *additional* memory that is needed by our highway hierarchy approach (divided by the number of nodes) compared to a space-efficient bidirectional implementation of Dijkstra’s algorithm. Query times are average values based on 10,000 random s - t -queries. “Speed-up” refers to a comparison with Dijkstra’s algorithm (unidirectional). Worst case is an upper bound for any possible query in the respective graph.

6.3. Parameters

Default Settings. Unless otherwise stated, the following default settings apply. We use the *maverick factor* $f = 2(i - 1)$ for the i -th iteration of the construction procedure, the contraction rate $c = 2$, the shortcut hops limit 10, and the neighborhood sizes H as stated in Table I—the same neighborhood size is used for all levels of a hierarchy. First, we contract the original graph.¹⁰ Then, we perform five iterations of our construction procedure, which determines a highway network and its core. Finally, we compute the distance table between all level-5 core nodes.

Self-Similarity. For two levels ℓ and $\ell + 1$ of a highway hierarchy, the *shrinking factor* is the ratio between $|E'_\ell|$ and $|E'_{\ell+1}|$. In our experiments, we observed that the highway hierarchies of Europe and the US were almost *self-similar* in the sense that the shrinking factor remained nearly unchanged from level to level when we used the same neighborhood size H for all levels—provided that H was not too small.

Figure 12 demonstrates the shrinking process for Europe. Note that the first contraction step is not shown. In contrast to our default settings, we do not stop after five iterations. For most levels and $H \geq 70$, we observe an almost constant shrinking factor (which appears as a straight line due to the logarithmic scale of the y -axis). The greater the neighborhood size, the greater the shrinking factor. The last iteration is an exception the highway network collapses, that is, it shrinks very fast because nodes that are close to the border of the network usually do not belong to the next level of the highway hierarchy, and when the network gets small, almost all nodes are close to the border. In case of the smallest neighborhood size ($H = 30$), the shrinking factor gets so small that the network does not collapse even after 14 levels have been constructed.

Varying the Neighborhood Size. Note that in order to simplify the experimental set-up all results in the remainder of Section 6.3 have been obtained without rearranging nodes by level. However, since we want to demonstrate the effects of choosing different parameter settings, the relative performance is already very meaningful.

¹⁰In Section 3, we gave the definition of the highway hierarchies where we first construct a highway network and then contract it. We decided to change this order in the experiments, that is, to start with an initial contraction phase, since we observed a better performance in this case.

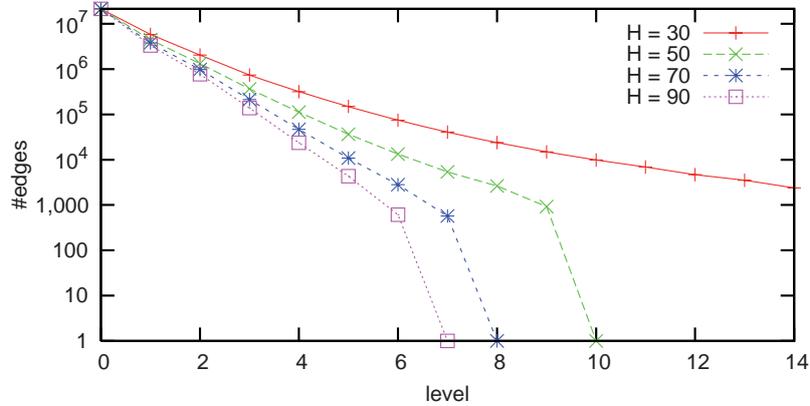


Fig. 12. Shrinking of the highway networks of Europe. For different neighbourhood sizes H and for each level ℓ , we plot $|E_\ell|$, that is, the number of edges that belong to the core of level ℓ .

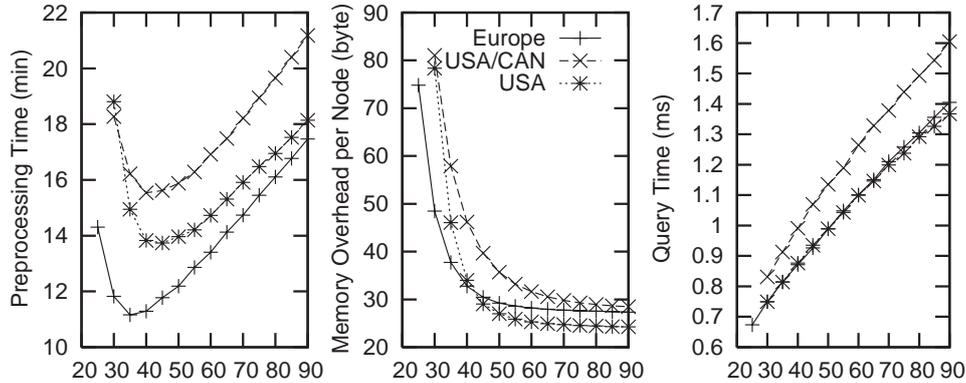


Fig. 13. Preprocessing and query performance depending on the neighborhood size H .

In one test series (Figure 13), we used all the default settings except for the neighborhood size H , which we varied in steps of 5. On the one hand, if H is too small, the shrinking of the highway networks is less effective so that the level-5 core is still quite big. Hence, we need much time and space to precompute and store the distance table. On the other hand, if H gets bigger, the time needed to preprocess the lower levels increases because the area covered by the local searches depends on the neighborhood size. Furthermore, during a query, it takes longer to leave the lower levels in order to get to the topmost level where the distance table can be used. Thus, the query time increases as well. We observe that the preprocessing time is minimized for neighborhood sizes around 40. In particular, the optimal neighborhood size does not vary very much from graph to graph. In other words, if we used the same parameter H , say 40, for all road networks, the resulting performance would be very close to the optimum. Obviously, choosing different neighborhood sizes leads to different space-time trade-offs.

Varying the Contraction Rate. In another test series (Table II), we did not use a distance table, but repeated the construction process until the topmost level was empty or the hierarchy consisted of 15 levels. We varied the contraction rate c from 0.5 to 2.5. In case of $c = 0.5$ (and $H = 30$), the shrinking of the highway networks does not

Table II. Preprocessing and Query Performance for the European Road Network Depending on the Contraction Rate c

contr. rate c	PREPROCESSING			QUERY		
	time (min)	over-head	\varnothing deg.	time (ms)	#settled nodes	#relaxed edges
0.5	83	30	3.2	391.73	472,326	1,023,944
1.0	15	28	3.7	5.48	6,396	23,612
1.5	11	28	3.8	1.93	1,830	9,281
2.0	11	29	4.0	1.85	1,542	8,913
2.5	11	30	4.1	1.96	1,489	9,175

Note: “Overhead” denotes the average memory overhead per node in bytes.

Table III. Preprocessing and Query Performance for the European Road Network Depending on the Number of Levels

# levels	PREPROC.		QUERY	
	time (min)	over-head	time (ms)	#settled nodes
6	12	48	0.75	709
7	10	34	0.93	852
8	10	30	1.14	991
9	10	30	1.35	1,123
10	10	29	1.54	1,241
11	10	29	1.67	1,326

Note: “Overhead” denotes the average memory overhead per node in bytes.

work properly so that the topmost level is still very big. This yields huge query times. Choosing larger contraction rates reduces the preprocessing and query times, since the cores and search spaces get smaller. However, the memory usage and the average degree are slightly increased since more shortcuts are introduced. Adding too many shortcuts ($c = 2.5$) further reduces the search space, but the number of relaxed edges increases so that the query times get worse.

Varying the Number of Levels. In a third test series (Table III), we used the default settings except for the number of levels, which we varied from 6 to 11. Note that the original graph and its core (i.e., the result of the first contraction step) counts as one level so that, for example, “6 levels” means that only five levels are constructed. In each test case, a distance table was used in the topmost level. The construction of the higher levels of the hierarchy is very fast and has no significant effect on the preprocessing times. In contrast, using only six levels yields a rather large distance table, which somewhat slows down the preprocessing and increases the memory usage. However, in terms of query times, “6 levels” is the optimal choice, since using the distance table is faster than continuing the search in higher levels. We omitted experiments with less levels, since this would yield very large distance tables consuming very much memory.

Results for further combinations of neighborhood size, contraction rate, and number of levels can be found in Tables V and VI in Appendix C.

6.4. Local Queries

For use in applications, it is unrealistic to assume a uniform distribution of queries in large graphs such as Europe or the USA. On the other hand, it would hardly be more realistic to arbitrarily cut the graph into smaller pieces. Therefore, we decided to measure local queries within the big graphs: For each power of two $r = 2^k$, we choose

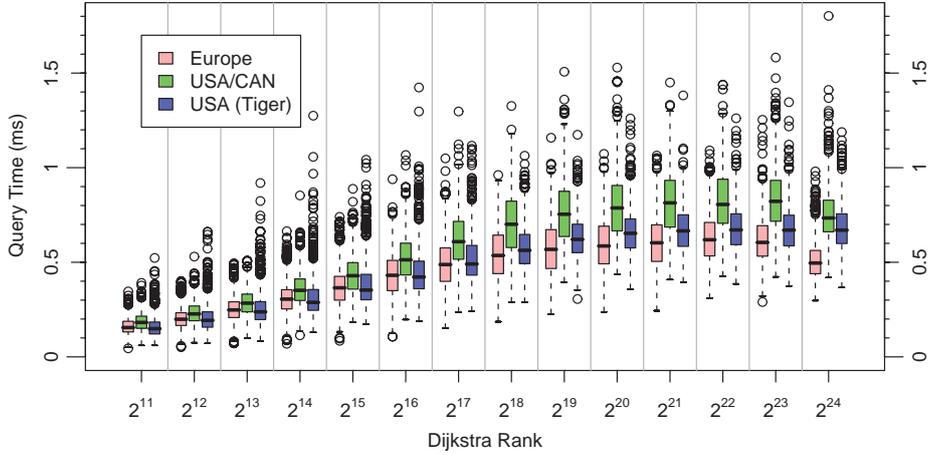


Fig. 14. Local queries. The distributions are represented as box-and-whisker plots [R Development Core Team 2004]: each box spreads from the lower to the upper quartile and contains the median, the whiskers extend to the minimum and maximum value omitting outliers, which are plotted.

random sample points s and then use Dijkstra’s algorithm to find the node t with Dijkstra rank $rk_s(t) = r$. We then use our algorithm to make an s - t -query. By plotting the resulting statistics for each value $r = 2^k$, we can see how the performance scales with a natural measure of difficulty of the query. Figure 14 shows the query times. Note that for ranks up to 2^{18} the median query times are scaling quite smoothly and the growth is much slower than the exponential increase we would expect in a plot with logarithmic x -axis, linear y -axis, and any growth rate of the form r^ρ for Dijkstra rank r and some constant power ρ ; the curve is also not the straight line one would expect from a query time logarithmic in r . For ranks $r \geq 2^{19}$, the query times hardly rise due to the fact that the all-pairs distance table can bridge the gap between the forward and backward search of these queries irrespective of dealing with a small or a large gap. In case of Europe and USA/CAN, the query times drop for $r = 2^{24}$, since r is only slightly smaller than the number of nodes so that the target lies close to the border of the respective road network, which implies some kind of trivial sense of goal direction for the backward search (because, in the beginning, we practically cannot go into the wrong direction).

6.5. Space Saving

If we omit the first contraction step and use a smaller contraction rate (\Rightarrow less shortcuts), use a bigger neighborhood size (\Rightarrow higher levels get smaller), and construct more levels before the distance table is used (\Rightarrow smaller distance table), the memory usage can be reduced considerably. In case of Europe, using seven levels, $H = 100$, and $c = 1$ yields an average overhead per node of 17 bytes. The construction and query times increase to 55min and 1.1ms, respectively.

6.6. Worst-Case Upper Bounds

By executing a query from each node of a given graph to an added isolated dummy node and a query from the dummy node to each actual node in the backward graph, we obtain a distribution of the search space sizes of the forward and backward search, respectively. We can combine both distributions to get an upper bound for the distribution of the search space sizes of bidirectional queries: Denoting by $\mathcal{F}_\rightarrow(x)$ ($\mathcal{F}_\leftarrow(x)$) the number of source (target) nodes whose search space consists of x nodes in a forward (backward)

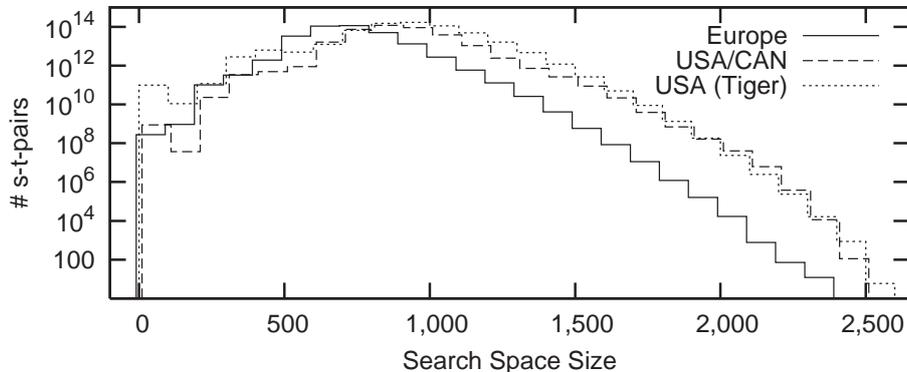


Fig. 15. Histogram of upper bounds for the search space sizes of $s-t$ -queries. To increase readability, only the outline of the histogram is plotted instead of the complete boxes.

search, we define $\mathcal{F}_{\leftrightarrow}(z) := \sum_{x+y=z} \mathcal{F}_{\rightarrow}(x) \cdot \mathcal{F}_{\leftarrow}(y)$, i.e., $\mathcal{F}_{\leftrightarrow}(z)$ is the number of $s-t$ -pairs such that the upper bound of the search space size of a query from s to t is z . In particular, we obtain the upper bound $\max\{z \mid \mathcal{F}_{\leftrightarrow}(z) > 0\}$ for the worst case without performing all n^2 possible queries.

Figure 15 visualizes the distribution $\mathcal{F}_{\leftrightarrow}(z)$ as a histogram.

In a similar way, we obtained a distribution of the number of entries in the distance table that have to be accessed during an $s-t$ -query. While the average values are reasonably small (4,066 in case of Europe), the worst case can get quite large (62,379). For example, accessing 62,379 entries in a table of size $9,351 \times 9,351$ takes about 1.1ms, where 9,351 is the number of nodes of the level-5 core of the European highway hierarchy. Hence, in some cases, the time needed to determine the optimal entry in the distance table might dominate the query time. We could try to improve the worst case by introducing a case distinction that checks whether the number of entries that have to be considered exceeds a certain threshold. If so, we would not use the distance table, but would continue with the normal search process. However, this measures would have only little effect on the average performance.

6.7. Comparisons

In Delling et al. [2009] there is a comparison of 27 variants of route planning algorithms. It turns out that highway hierarchies outperform all methods developed before 2006 with respect to query times on continental sized networks. Even with the advent of faster techniques like transit-node routing [Bast et al. 2007] or techniques easier to dynamize like highway-node routing [Schultes and Sanders 2007], highway hierarchies remained very important as a preprocessing tool. In 2008, contraction hierarchies [Geisberger et al. 2008] and sophisticated combinations of hierarchical and goal directed techniques [Bauer et al. 2008; Bauer and Delling 2008] became available that outperform highway hierarchies for road networks with a travel time objective. They also use contraction in a similar way as highway hierarchies, but it seems that for road networks, the sparsification possible by finding highway edges is not necessary.

In Bauer et al. [2008], a similar comparison is also made for optimizing the travelled distance. Here, highway hierarchies are one of several techniques that are Pareto optimal in the sense that none of the other techniques is better with respect to all three criteria of query time, preprocessing time, and space overhead. This is an indication that highway hierarchies remain an interesting candidate for networks that have a less pronounced hierarchy than road networks with the travel time metric.

Table IV. Outputting Complete Path Descriptions

Additional preprocessing time, additional disk space and query time that is needed to determine a complete description of the shortest path and to traverse it summing up the weights of all edges—assuming that the query to determine its lengths has already been performed. Moreover, the average number of hops, that is, the average path length in terms of number of nodes, is given. The three algorithmic variants are described in Section 5.4.

Var.	Europe				USA (Tiger)			
	preproc. (s)	space (MB)	query (ms)	# hops (avg.)	preproc. (s)	space (MB)	query (ms)	# hops (avg.)
1	0	0	17.22	1,366	0	0	39.69	4,410
2	69	126	0.49	1,366	68	127	1.16	4,410
3	74	225	0.19	1,366	70	190	0.25	4,410

In Bauer et al. [2007], there is a comparison of highway hierarchies with 10 other techniques, including SHARC [Bauer and Delling 2008] and several older techniques. No attempt is made to adapt parameter settings to the input. For road networks with various distance measures, railway networks with and without time information, sparse random unit-disk graphs and two-dimensional grids with random edge weights, highway hierarchies are consistently Pareto optimal because they have a very good compromise between preprocessing time and query time. For graphs with a small world property (the internet at the router level, DPLP citation, and coauthor networks), plain bidirectional Dijkstra wins over all speed-up techniques tried. For higher-dimensional random grid graphs, methods with a better sense of goal direction win over highway hierarchies.

6.8. Outputting Complete Path Descriptions

So far, we have reported only the times needed to compute the shortest-path length between two nodes. Now, we determine a complete description of the shortest path. In Table IV, we give the additional preprocessing time and the additional disk space for the unpacking data structures. Furthermore, we report the additional time that is needed to determine a complete description of the shortest path and to sum up the weights of all edges¹¹—assuming that the query to determine the shortest-path length has already been performed. That means that the total average time to determine a shortest path is the time given in Table IV plus the query time given in previous tables.¹² Note that Variant 1 is no longer supported by the current version of our implementation so that the numbers in the first data row of Table IV have been obtained with an older version and different settings.

We can conclude that even Variant 3 requires little additional preprocessing time and only a moderate amount of space. With Variant 3, the time for outputting the path remains considerably smaller than the time to determine the path length and a factor 3 to 5 smaller than using Variant 2. The U.S. graph profits more than the European graph, since it has paths with considerably larger hop counts, perhaps due to a larger number of degree 2 nodes in the input. Note that due to cache effects, the time for outputting the path using preprocessed shortcuts is likely to be considerably smaller than the time for traversing the shortest path in the original graph.

¹¹Note that we do not traverse the path in the original graph, but we directly scan the assembled description of the path. We sum up edges because otherwise the unpacked data would not really be accessed, possibly leading to unrealistically small running time.

¹²Note that in the current implementation outputting path descriptions and the feature to rearrange nodes by level are mutually exclusive. However, this is not a limitation in principle.

6.9. Turning Restrictions

We did a simple experiment with the German road network (a subgraph of our European network) with 4,380,384 nodes and 10,668,470 directed edges. The HH for this graph has 15,267,563 edges. After adding 206,213 real-world turning restrictions (4.4% more nodes and 1.1% more edges for the input), we get a HH with 0.8% more edges than before (i.e., a smaller increase than for the input). The query time increases by 3%. Overall, we can conclude that turn restrictions have a negligible effect on the outcome.

6.10. Distance Metric

When we apply a distance metric instead of the usual (and for most practical applications more relevant) travel time metric, the hierarchy that is inherent in the road network is less distinct, since the difference between fast and slow roads fades. We no longer observe the self-similarity in the sense that a fixed neighborhood size yields an almost constant shrinking factor. Instead, we have to use an increasing sequence of neighborhood sizes to ensure a proper shrinking. For Europe, we use $H = 100, 200, 300, 400$ and 500 to construct five levels before an all-pairs distance table is built. Constructing the hierarchy takes 34 minutes and entails a memory overhead of 36 bytes per node. On average, a random query then takes 4.8ms, settling 4,810 nodes and relaxing 33,481 edges. Further experiments on different metrics can be found in Delling et al. [2006].

6.11. An Even Larger Road Network

We recently obtained a new version of the European road network that is larger than the old one and covers more countries.¹³ It has been provided for scientific use by the company ORTEC and consists of 33,726,989 nodes and 75,108,089 directed edges. We use the same parameters as for the old version (in particular, $H = 30$) and observe a very good shrinking behavior: We have 1.87 times as many nodes in the beginning, but after the construction of the same number of levels, only 1.04 times as many nodes remain. Thus, the same number of levels is sufficient, only the distance table gets slightly bigger. We arrive at a preprocessing time of 18 minutes, a memory overhead of 37 bytes per node, and query times of 0.6ms for random queries; on average, 685 nodes are settled and 2,457 edges are relaxed.

7. DISCUSSION

Highway hierarchies are a simple, robust, and space-efficient concept that allows very efficient exact fastest-path queries even in huge real-world road Networks. These attributes have been confirmed in an extensive experimental study. The ideas developed for highway hierarchies like shortcuts, contraction, preprocessing by local search, and distance tables are an important ingredient of subsequent techniques like advanced reach-based routing [Goldberg et al. 2006, 2007], highway-node routing [Schultes and Sanders 2007], and combinations of techniques [Bauer and Delling 2008; Bauer et al. 2008]. As of now, it looks like highway hierarchies on static road networks are no longer competitive with the best subsequent techniques like contraction hierarchies [Geisberger et al. 2008]. However, highway hierarchies have an interesting feature that currently is not used by any of its competitors and that might turn out to be useful for networks where techniques like contraction hierarchies produce too many shortcuts: Networks are more sparse at the higher levels because edges that are only needed near source and target need not be present. Therefore, only with highway hierarchies do we observe the strong self-similarity from Figure 12. Probably, to become competitive,

¹³In addition to the old version, the Czech Republic, Finland, Hungary, Ireland, Poland, and Slovakia.

new versions of highway hierarchies would have to incorporate subsequently developed ideas, particularly, the more sophisticated contraction routines from contraction hierarchies [Geisberger et al. 2008] and advanced combinations with goal directed techniques [Bauer et al. 2008].

APPENDIXES

A. QUERY—PROOF OF CORRECTNESS

Additional Notation. “ \circ ” denotes *path concatenation*. $\text{succ}(u, P)$ and $\text{pred}(u, P)$ denote the direct successor and predecessor of u on P , respectively. We just write $\text{succ}(u)$ and $\text{pred}(u)$ if the path is clear from the context. For two nodes u and v on some path, $\min(u, v)$ denotes u if $u \leq v$ and v otherwise. $\max(u, v)$ is defined analogously. $d_P(u, v) := w(P|_{u \rightarrow v})$ denotes the distance from u to v along the path P . Note that for any edge (u, v) on P , we have $w(u, v) = d_P(u, v)$.

A.1. Termination and Special Cases

Since we have set the neighborhood radius in the topmost level to infinity (R2), we are never tempted to go upwards beyond the topmost level. This observation is formalized in the following lemma.

LEMMA A.1. *The for-loop in Line 9 of the highway query algorithm always terminates with $\ell \leq L$ and $(\ell = L \rightarrow \text{gap} = \infty)$.*

PROOF. We only consider iterations where the forward search direction is selected; analogous arguments apply to the backward direction. By an inductive proof, we show that at the beginning of any iteration of the main while-loop, we have $\ell(u) \leq L$ and $(\ell(u) = L \rightarrow \text{gap}(u) = \infty)$ for any node u in \vec{Q} .

Base Case. True for the first iteration, where only s belongs to \vec{Q} : we have $\ell(s) = 0 \leq L$ and $\text{gap}(s) = r_0^\rightarrow(s)$ (Line 2), which is equal to infinity if $L = 0$ (due to R2).

Induction Step. We assume that our claim is true for iteration i and show that it also holds for iteration $i + 1$. Due to the induction hypothesis, we have $\ell(u) \leq L$ and $(\ell(u) = L \rightarrow \text{gap}(u) = \infty)$ in Line 5. If $\ell(u) = L$, we have $\text{gap} = \text{gap}' = r_{\ell(u)}^\rightarrow(u) = \infty$ (Line 7 and 9, R2); thus, the for-loop in Line 9 terminates immediately with $\ell = \ell(u) = L$. Otherwise ($\ell(u) < L$), the for-loop either terminates with $\ell < L$ or reaches $\ell = L$; in the latter case, we have $\text{gap} = r_\ell^\rightarrow(u) = \infty$ (Line 9, R2); hence, the loop terminates.

Thus, in any case, the loop terminates with $\ell \leq L$ and $(\ell = L \rightarrow \text{gap} = \infty)$. Therefore, if the node v adopts the key k in Line 13 (either by a decreaseKey or an insert operation), the new key fulfills the required condition.

This concludes our inductive proof, which also shows that the claim of this lemma holds during any iteration of the main while-loop. \square

It is easy to see that the following property of Dijkstra’s algorithm also holds for the highway query algorithm.

PROPOSITION A.2. *For each search direction, the sequence of distances $\delta(u)$ of settled nodes u is monotonically increasing.*

Now, we can prove that

LEMMA A.3. *The highway query algorithm terminates.*

PROOF. The for-loop in Line 9 always terminates due to Lemma A.1. The for-loop in Line 8 terminates, since the edge set is finite. The main while-loop in Line 3 terminates, since each node v is inserted into each priority queue at most once, namely if it is unreachable (Line 13); if it is reached, it either already belongs to the priority queue or

it has already been settled; in the latter case, we know that $\delta(v) \leq \delta(u) \leq \delta(u) + w(e)$ (Proposition A.2; edge weights are nonnegative) so that no priority queue operation is performed due to the specification of the decreaseKey operation. \square

The *special case* that there is no path from s to t is trivial. The algorithm terminates due to Lemma A.3 and returns ∞ since no node can be settled from both search directions (otherwise, there would be, some path from s to t). For the remaining proof, we assume that a shortest path from s to t exists in the original graph G .

A.2. Contracted and Expanded Paths

LEMMA A.4. *Shortcuts do not overlap, that is, if there are four nodes $u < u' < v < v'$ on a path P in G , then there cannot exist both a shortcut (u, v) and a shortcut (u', v') at the same time.*

PROOF. Let us assume that there is a shortcut $(u, v) \in S_\ell$ for some level ℓ . All inner nodes, in particular u' , belong to B_ℓ . Since u' does not belong to V'_ℓ , a shortcut that starts from u' can belong only to some level $k < \ell$. If there was a shortcut $(u', v') \in S_k$, the inner node v would have to belong to B_k , which is a contradiction, since $v \in V'_\ell$. \square

Definition A.5. For a given path P in a given highway hierarchy \mathcal{G} , the *contracted path* $\text{ctr}(P)$ is defined in the following way: While there is a subpath $\langle u, b_1, b_2, \dots, b_k, v \rangle$ with $u, v \in V'_\ell$ and $b_i \in B_\ell$, $1 \leq i \leq k$, $k \geq 1$, for some level ℓ , replace it by the shortcut edge $(u, v) \in S_\ell$.

Note that this definition terminates, since the number of nodes in the path is reduced by at least one in each step and the definition is unambiguous due to Lemma A.4.

Definition A.6. For a given path P in a given highway hierarchy \mathcal{G} , the *level- ℓ expanded path* $\text{exp}(P, \ell)$ is defined in the following way: While the path contains a shortcut edge $(u, v) \in S_k$ for some $k > \ell$, replace it by the represented path in G_k .

Note that this definition terminates, since an expanded subpath can only contain shortcuts of a smaller level.

A.3. Highway Path

Consider a given highway hierarchy \mathcal{G} and an arbitrary path $P = \langle s, \dots, t \rangle$. In the following, we will bring out the structure of P with regard to \mathcal{G} .

Last Neighbor and First Core Node. For any level ℓ and any node u on P , we define the *last succeeding level- ℓ neighbor* $\vec{w}_\ell^P(u)$ and the *first succeeding level- ℓ core node* $\vec{\alpha}_\ell^P(u)$: $\vec{w}_\ell^P(u)$ is the node $v \in \{x \in P \mid u \preceq x \wedge d_P(u, x) \leq r_\ell^{\vec{}}(u)\}$ that maximizes $d_P(u, v)$, and $\vec{\alpha}_\ell^P(u)$ is the node $v \in \{t\} \cup \{x \in P \cap V'_\ell \mid u \preceq x\}$ that minimizes $d_P(u, v)$. The *last preceding neighbor* $\overleftarrow{w}_\ell^P(u)$ and the *first preceding core node* $\overleftarrow{\alpha}_\ell^P(u)$ are defined analogously.

Unidirectional Labeling. Now, we inductively define a forward *labelling* of the path P . The labels s_0 and s'_0 are set to s and for ℓ , $0 \leq \ell < L$, we set $s_{\ell+1} := \vec{w}_\ell^P(s'_\ell)$ and $s'_{\ell+1} := \vec{\alpha}_{\ell+1}^P(s_{\ell+1})$. Furthermore, in order to avoid some case distinctions, $s_{L+1} := t$. For an example, we refer to Figure 9.

PROPOSITION A.7. *The following properties apply to the (Unidirectional) forward labelling of P :*

- U1: $s = s_0 = s'_0 \leq s_1 \leq s'_1 \leq \dots \leq s_L \leq s'_L \leq s_{L+1} = t$
- U2a: $\forall \ell, 0 \leq \ell \leq L : \forall u, s'_\ell \leq u \leq s_{\ell+1} : d_P(s'_\ell, u) \leq r_\ell^{\vec{}}(s'_\ell)$
- U2b: $\forall \ell, 0 \leq \ell \leq L : \forall u \succ s_{\ell+1} : d_P(s'_\ell, u) > r_\ell^{\vec{}}(s'_\ell)$

- U3: $\forall \ell, 0 \leq \ell \leq L : \forall u, s_\ell \leq u < s'_\ell : u \notin V'_\ell$
- U4: $\forall \ell, 0 \leq \ell \leq L : s'_\ell = t \vee s'_\ell \in \bar{V}'_\ell$

A backward labelling (specifying nodes t_ℓ and t'_ℓ) is defined analogously.

Meeting Level and Point. The *meeting level* λ of P is 0 if $s = t$ and $\max\{\ell \mid s_\ell \leq t_\ell\}$ if $s \neq t$. Note that $\lambda \leq L$ (in any case) and $t_{\lambda+1} < s_{\lambda+1}$ (in case that $s \neq t$). The *meeting point* p of P is either t_λ (if $t_\lambda \leq s'_\lambda$) or $\min(s_{\lambda+1}, t'_\lambda)$ (otherwise). Figure 10 gives an example.

PROPOSITION A.8. *The following properties apply to the Meeting point p of P :*

- M1: $s_\lambda \leq p \leq t_\lambda$
- M2: $t_{\lambda+1} \leq p \leq s_{\lambda+1}$
- M3: $\forall \ell, 0 \leq \ell \leq L : (s'_\ell < p \rightarrow p \leq t'_\ell) \wedge (p < t'_\ell \rightarrow s'_\ell \leq p)$

PROOF. The case $s = t$ is trivial. Subsequently, we assume $s \neq t$. In order to prove M1, M2, and (M3 for $\ell = \lambda$), we distinguish between two cases.

Case 1. $t_\lambda \leq s'_\lambda$. Then, $p = t_\lambda$. M1 is fulfilled due to the definition of the meeting level, which implies $s_\lambda \leq t_\lambda$. Furthermore, due to U1, we have $t_{\lambda+1} \leq t'_\lambda \leq t_\lambda = p \leq s'_\lambda \leq s_{\lambda+1}$ so that M2 and (M3 for $\ell = \lambda$) are fulfilled.

Case 2. $s'_\lambda < t_\lambda$. Then, $p = \min(s_{\lambda+1}, t'_\lambda)$.

Subcase 2.1. $s_{\lambda+1} \leq t'_\lambda$. Then, $p = s_{\lambda+1}$. We have $s_\lambda \leq s'_\lambda \leq s_{\lambda+1} = p \leq t'_\lambda \leq t_\lambda$ so that M1 and (M3 for $\ell = \lambda$) are fulfilled. Furthermore, M2 holds due to $t_{\lambda+1} < s_{\lambda+1}$.

Subcase 2.2. $t'_\lambda < s_{\lambda+1}$. Then, $p = t'_\lambda$. Since $s'_\lambda < t_\lambda \leq t$, we know that $s'_\lambda \in V'_\lambda$ (due to U4). Thus, we have $s'_\lambda \leq t'_\lambda \leq t_\lambda$ (otherwise $(t'_\lambda < s'_\lambda \leq t_\lambda)$, we would have a contradiction with U3). Hence, $s_\lambda \leq s'_\lambda \leq t'_\lambda = p \leq t_\lambda$ so that M1 and (M3 for $\ell = \lambda$) are fulfilled. M2 holds as well, since $t_{\lambda+1} \leq t'_\lambda = p < s_{\lambda+1}$.

It remains to show M3 for $\ell < \lambda$ and for $\ell > \lambda$. In the former case, M3 holds due to M1, which implies $s'_\ell \leq s_\lambda \leq p \leq t_\lambda \leq t'_\ell$ (U1). In the latter case, M3 holds due to M2, which implies $t'_\ell \leq t_{\lambda+1} \leq p \leq s_{\lambda+1} \leq s'_\ell$ (U1). \square

Highway Path. $P = \langle s, \dots, t \rangle$ is a *highway path* (Figure 11) if and only if the following two Highway properties are fulfilled:

- H1: $\forall \ell, 0 \leq \ell \leq L : \text{H1}(\ell)$
- H2: $\forall \ell, 0 \leq \ell \leq L : \text{H2}(\ell)$

where

- H1(ℓ): $\forall(u, v), s'_\ell \leq u < v \leq t'_\ell : u, v \in V'_\ell$
- H2(ℓ): $\forall(u, v), s_\ell \leq u < v \leq t_\ell : \ell(u, v) \geq \ell$

A.4. Reachability Along a Highway Path

We consider a path $P = \langle s, \dots, t \rangle$. For a node u on P , we define the *reference level* $\bar{\ell}(u) := \max(\{0\} \cup \{i \mid s_i < u\})$.

PROPOSITION A.9. *For any two nodes u and v with $u \leq v$, the following reference Level properties apply:*

- L1: $0 \leq \bar{\ell}(u) \leq L$
- L2: $s_{\bar{\ell}(u)} \leq u$
- L3: $u \leq s_{\bar{\ell}(u)+1}$
- L4: $\bar{\ell}(u) \leq \bar{\ell}(v)$

Definition A.10. A node u is said to be Appropriately reached/settled with the key $k = (\delta(u), \ell(u), \text{gap}(u))$ on the path P if and only if all of the following conditions are fulfilled:

- A₁(k, u): $\delta(u) = d_0(s, u)$

- $A_2(k, u)$: $\ell(u) = \bar{\ell}(u)$
- $A_3(k, u)$: $\text{gap}(u) = \begin{cases} \infty & \text{if } u \preceq s'_{\ell(u)} \\ r_{\ell(u)}^{\rightarrow}(s'_{\ell(u)}) - d_P(s'_{\ell(u)}, u) & \text{otherwise} \end{cases}$
- $A_4(u)$: $\forall i : t \neq s'_i \preceq u \rightarrow u \in V'_i$

The following (somewhat technical) lemma will be used to prove Lemmas A.12 and A.13. Basically, it states that in the highway query algorithm the search level and the gap to the next applicable neighborhood border are set correctly.

LEMMA A.11. *Consider a path $P = (s, \dots, t)$ and an edge (u, v) on P . Assume that u is settled by the highway query algorithm appropriately with some key k . We consider the attempt to relax the edge (u, v) . After Line 9 has been executed, the following Invariants apply with regard to the variables ℓ and gap :*

- I1: (a) $s_\ell \preceq u \wedge$ (b) $v \preceq s_{\ell+1}$
- I2: $\ell = \bar{\ell}(v)$
- I3: $\text{gap} = \begin{cases} \infty & \text{if } v \preceq s'_\ell, \\ r_{\ell}^{\rightarrow}(s'_\ell) - d_P(s'_\ell, u) & \text{otherwise.} \end{cases}$

PROOF. We distinguish between two cases in order to prove I1 and I3.

Case 1. Zero iterations of the for-loop in Line 9 take place ($\ell = \ell(u)$).

In this case, we have $\ell = \ell(u)$ and $w(u, v) \leq \text{gap}'$. Hence, $s_\ell \preceq u$ due to $A_2(k, u)$ and L2 (\Rightarrow I1a). In order to show I1b and I3, we distinguish between three subcases.

- Subcase 1.1.* $u < s'_\ell \Rightarrow v \preceq s'_\ell \preceq s_{\ell+1}$ (U1) (\Rightarrow I1b). Furthermore, because of $\text{gap}(u) = \infty$ ($A_3(u, k)$), we have $\text{gap} = \text{gap}' = r_{\ell(u)}^{\rightarrow}(u) = \infty$ due to U3 and R1 (\Rightarrow I3, since $v \preceq s'_\ell$).
- Subcase 1.2.* $u = s'_\ell \Rightarrow \text{gap}(u) = \infty$ ($A_3(u, k)$) $\Rightarrow w(u, v) \leq \text{gap}' = r_{\ell}^{\rightarrow}(u)$ (Line 7) $\Rightarrow d_P(s'_\ell, v) \leq r_{\ell}^{\rightarrow}(s'_\ell)$ (since $u = s'_\ell$) $\Rightarrow v \preceq s_{\ell+1}$ (U2b) (\Rightarrow I1b). Furthermore, $\text{gap} = \text{gap}' = r_{\ell}^{\rightarrow}(u) = r_{\ell}^{\rightarrow}(s'_\ell) - d_P(s'_\ell, u)$ (since $u = s'_\ell$) implies I3, since $s'_\ell < v$.
- Subcase 1.3.* $u > s'_\ell \Rightarrow \text{gap}(u) = r_{\ell}^{\rightarrow}(s'_\ell) - d_P(s'_\ell, u)$ ($A_3(u, k)$). By Lemma A.1, $\ell \leq L$ and ($\ell = L \rightarrow \text{gap} = \infty$). If $\ell = L$, we have $v \preceq t = s_{L+1} = s_{\ell+1}$ (\Rightarrow I1b) and $\text{gap} = \infty = r_{\ell}^{\rightarrow}(s'_\ell) - d_P(s'_\ell, u)$ (R2) (\Rightarrow I3, since $s'_\ell < v$). Subsequently, we deal with the remaining case $\ell < L$. The facts that $u \preceq t$ and $s'_\ell < u$ imply $s'_\ell \neq t$, which yields $s'_\ell \in V'_\ell$ due to U4. Hence, due to R3, $\text{gap}(u) \neq \infty \Rightarrow w(u, v) \leq \text{gap}' = \text{gap}(u)$ (Line 7) $\Rightarrow d_P(u, v) \leq r_{\ell}^{\rightarrow}(s'_\ell) - d_P(s'_\ell, u) \Rightarrow d_P(s'_\ell, v) \leq r_{\ell}^{\rightarrow}(s'_\ell) \Rightarrow v \preceq s_{\ell+1}$ (U2b) (\Rightarrow I1b). Furthermore, $\text{gap} = \text{gap}' = \text{gap}(u) = r_{\ell}^{\rightarrow}(s'_\ell) - d_P(s'_\ell, u)$ implies I3, since $s'_\ell < v$.

Case 2. At least one iteration of the for-loop takes place ($\ell > \ell(u)$).

We claim that after any iteration of the for-loop, we have $u = s_\ell$. Proof by induction:

Base Case. We consider the first iteration of the for-loop. Line 9 and the fact that an iteration takes place imply $w(u, v) > \text{gap}'$, which means that $\text{gap}' \neq \infty$. We distinguish between two subcases to show that $d_P(s'_{\ell(u)}, v) > r_{\ell(u)}^{\rightarrow}(s'_{\ell(u)})$.

- Subcase 2.1.* $u \preceq s'_{\ell(u)} \Rightarrow \text{gap}(u) = \infty$ ($A_3(u, k)$) $\Rightarrow w(u, v) > \text{gap}' = r_{\ell(u)}^{\rightarrow}(u)$ (Line 7) $\Rightarrow r_{\ell(u)}^{\rightarrow}(u) \neq \infty$. We have $s_{\ell(u)} \preceq u \preceq s'_{\ell(u)}$ due to L2, $A_2(u, k)$, and the assumption of Subcase 2.1. However, we can exclude that $s_{\ell(u)} \preceq u < s'_{\ell(u)}$: This would imply $u \notin V'_{\ell(u)}$ (U3) and, thus, $r_{\ell(u)}^{\rightarrow}(u) = \infty$ (R1). Therefore, $u = s'_{\ell(u)} \Rightarrow d_P(s'_{\ell(u)}, v) > r_{\ell(u)}^{\rightarrow}(s'_{\ell(u)})$
- Subcase 2.2.* $u > s'_{\ell(u)} \Rightarrow s'_{\ell(u)} \neq t \Rightarrow s'_{\ell(u)} \in V'_{\ell(u)}$ (U4). Furthermore, $\text{gap}(u) = r_{\ell(u)}^{\rightarrow}(s'_{\ell(u)}) - d_P(s'_{\ell(u)}, u)$ ($A_3(u, k)$) $\Rightarrow \text{gap}(u) \neq \infty$ (due to R3, since $\ell(u) < L$ (Lemma A.1) and $s'_{\ell(u)} \in V'_{\ell(u)} \Rightarrow d_P(u, v) = w(u, v) > \text{gap}' = \text{gap}(u) = r_{\ell(u)}^{\rightarrow}(s'_{\ell(u)}) - d_P(s'_{\ell(u)}, u)$ (Line 7) $\Rightarrow d_P(s'_{\ell(u)}, v) > r_{\ell(u)}^{\rightarrow}(s'_{\ell(u)})$

From $d_P(s'_{\ell(u)}, v) > r_{\ell(u)}^{\rightarrow}(s'_{\ell(u)})$, it follows that $s_{\ell(u)+1} < v$ (U2a), which implies $s_{\ell(u)+1} \preceq u$. Hence, $u = s_{\ell(u)+1}$ (since $u \preceq s_{\ell(u)+1}$ due to L3 and $A_2(k, u)$).

Induction Step. Let us now deal with the iteration from level i to level $i + 1$ for $i \geq \ell(u) + 1$. We have $w(u, v) > \text{gap} = r_i^{\rightarrow}(u)$, which implies $r_i^{\rightarrow}(u) \neq \infty$. Starting with $u = s_i \leq s'_i \leq s_{i+1}$ (induction hypothesis, U1), we can conclude that $u = s'_i$ (U3, R1) $\Rightarrow d_P(s'_i, v) > r_i^{\rightarrow}(s'_i) \Rightarrow s_{i+1} < v$ (U2a) $\Rightarrow s_{i+1} \leq u \Rightarrow u = s_{i+1}$ (since $u \leq s_{i+1}$). This completes our inductive proof.

After the last iteration, we have $u = s_\ell \leq s'_\ell$ (\Rightarrow I1a). Furthermore, $w(u, v) \leq r_\ell^{\rightarrow}(u)$. If $u < s'_\ell$, we obtain $v \leq s'_\ell \leq s_{\ell+1}$ (\Rightarrow I1b) and $\text{gap} = r_\ell^{\rightarrow}(u) = \infty$ due to U3 and R1 (\Rightarrow I3, since $v \leq s'_\ell$). Otherwise ($u = s'_\ell$), we get $d_P(s'_\ell, v) \leq r_\ell^{\rightarrow}(s'_\ell)$, which implies $v \leq s_{\ell+1}$ as well (U2b) (\Rightarrow I1b); furthermore, $\text{gap} = r_\ell^{\rightarrow}(u) = r_\ell^{\rightarrow}(s'_\ell) - d_P(s'_\ell, u)$ (since $u = s'_\ell$) implies I3 since $s'_\ell < v$. This completes the proof of I1 and I3.

I2 ($\ell(v) = \ell$) directly follows from $s_\ell < v \leq s_{\ell+1}$ (I1). \square

LEMMA A.12. *Consider a highway path $P = \langle s, \dots, t \rangle$ and an edge (u, v) on P such that u precedes the meeting point p . Assume that u has been appropriately settled. Then, the edge (u, v) is not skipped, but relaxed.*

PROOF. We consider the relaxation of the edge (u, v) . Due to Lemma A.11, the Invariants I1–I3 apply after Line 9 has been executed. Now, we consider Line 10 of the highway query algorithm.

I1 and M2 imply $s_\ell \leq u < p \leq s_{\lambda+1}$. Hence, $\ell \leq \lambda$. Thus, $u < p \leq t_\lambda \leq t_\ell$ (M1). By H2, we obtain $\ell(u, v) \geq \ell$. Therefore, the edge (u, v) is not skipped at this point.

Moreover, we prove that the condition in Line 11 is not fulfilled, since (u, v) belongs to a highway path. This means that the edge (u, v) is not skipped at this point either. We have to show that $u \notin V'_\ell \vee v \notin B_\ell$. We have $s_\ell \leq u$ (I1). If $u < s'_\ell$, we get $u \notin V'_\ell$ (U3). Otherwise, we have $s'_\ell \leq u < v \leq p \leq t'_\ell$ (M3), which yields $v \notin B_\ell$ (H1).

Therefore, (u, v) is not skipped, but relaxed. \square

LEMMA A.13. *Consider a shortest path $P = \langle s, \dots, t \rangle$ and an edge (u, v) on P . Assume that u has been appropriately settled with some key k . Furthermore, assume that the edge (u, v) is not skipped, but relaxed. Then, v can be appropriately reached from u with key k' .*

PROOF. We consider the relaxation of the edge (u, v) . Due to Lemma A.11, the Invariants I1–I3 apply after Line 9 has been executed. Therefore, since (u, v) is not skipped but relaxed, the node v can be reached with the key

$$k' = (\delta'(v), \ell'(v), \text{gap}'(v)) := (\delta(u) + w(u, v), \ell, \text{gap} - w(u, v)).$$

Thus, $A_1(k', v)$, $A_2(k', v)$, and $A_3(k', v)$ hold, since P is a shortest path and due to $A_1(k, u)$, I2, and I3.

Consider an arbitrary i such that $t \neq s'_i \leq v$. To prove $A_4(v)$, we have to show that $v \in V'_i$. Due to U4, this is true for $s'_i = v$. Now, we deal with the remaining case $s'_i \leq u < v$. Since $v \leq s_{\ell+1} \leq s'_{\ell+1}$ (I1, U1), we have $i \leq \ell$. The case $\ell = 0$ is trivial; hence, we assume $\ell > 0$. Since the edge (u, v) is not skipped, we know that Restriction 1 does not apply. Therefore, we have $\ell(u, v) \geq \ell$, which implies $v \in V_\ell \subseteq V'_{\ell-1}$. For $i < \ell$, we have $V'_{\ell-1} \subseteq V'_i$ and are done. For $i = \ell$, we have $u \in V'_\ell$ due to $A_4(u)$. This implies $v \notin B_\ell$, since Restriction 2 does not apply as well. $v \in V_\ell$ and $v \notin B_\ell$ yield $v \in V'_\ell$. \square

Analogous considerations hold for the backward search.

A.5. Finding an Optimal Path

Source and target nodes s and t are given such that a shortest path from s to t exists.¹⁴

¹⁴The special case that there is no path from s to t is treated in Section A.1.

Definition A.14. A state z is a triple (P, u, \bar{u}) , where P is a s - t -path, $u, \bar{u} \in V \cap P$, and $u \leq \bar{u}$.

Definition A.15. A state $z = (P, u, \bar{u})$ is *valid* if and only if all of the following valid State properties are fulfilled:

- S1: $w(P) = d_0(s, t)$
- S2: $P|_{u \rightarrow \bar{u}}$ is contracted, i.e., $P|_{u \rightarrow \bar{u}} = \text{ctr}(P|_{u \rightarrow \bar{u}})$
- S3: $P|_{s \rightarrow u}$ and $P|_{\bar{u} \rightarrow t}$ are paths in the forward and backward search tree, respectively.

LEMMA A.16. Consider a valid state $z = (P, u, \bar{u})$ and an arbitrary node x , $s \leq x \leq u$, on P . Then, x has been appropriately settled. Analogously for backward search.

PROOF.

Base Case: True for s .

Induction Step: We assume that $y, s \leq y < u$, has been appropriately settled and show that $x = \text{succ}(y)$ is appropriately settled as well. Since (y, x) belongs to the forward search tree (S3), we know that (y, x) is not skipped, but relaxed. The other prerequisites of Lemma A.13 are fulfilled as well (due to the induction hypothesis and S1). Thus, we can conclude that x can be appropriately *reached* from y . Since (y, x) belongs to the forward search tree, we know that x is also *settled* from y . \square

LEMMA A.17. If $z = (P, u, \bar{u})$ is a valid state, then P is a highway path.

PROOF. All labels (e.g., s'_ℓ) in this proof refer to P . We show that the highway properties H1 and H2 are fulfilled by induction over the level ℓ .

Base Case: H2(0) trivially holds, since $\ell(u, v) \geq 0$ for any edge (u, v) .

Induction Step (a): H2(ℓ) \rightarrow H1(ℓ). We assume $s'_\ell < t'_\ell$. (Otherwise, H1(ℓ) is trivially fulfilled.) This implies $s'_\ell \neq t$. Consider an arbitrary node x on $P|_{s'_\ell \rightarrow t'_\ell}$. We distinguish between three cases.

Case 1. $x \leq u$. According to Lemma A.16, $A_4(x)$ holds. Hence, $x \in V'_\ell$ since $s'_\ell \leq x$.

Case 2. $u \leq x \leq \bar{u}$. We have $y := \max(u, s'_\ell) \in V'_\ell$ (either by Lemma A.16: $A_4(u)$ or by U4). Analogously, $\bar{y} := \min(\bar{u}, t'_\ell) \in V'_\ell$. Since $u \leq y \leq x \leq \bar{y} \leq \bar{u}$ and $P|_{u \rightarrow \bar{u}} = \text{ctr}(P|_{u \rightarrow \bar{u}})$ (S2), we can conclude that $x \notin B_\ell$. Furthermore, we have $x \in V_\ell$ (due to H2(ℓ)). Thus, $x \in V'_\ell$.

Case 3. $\bar{u} \leq x$. Analogous to Case 1.

Induction Step (b): H1(ℓ) \wedge H2(ℓ) \rightarrow H2($\ell + 1$). Let \bar{P} denote $\text{exp}(P|_{s'_\ell \rightarrow t'_\ell}, \ell)$ and consider an arbitrary edge (x, y) on \bar{P} . If (x, y) is part of an expanded shortcut, we have $\ell(x, y) \geq \ell + 1$ and $x, y \in V_{\ell+1} \subseteq V'_\ell$. Otherwise, (x, y) belongs to $P|_{s'_\ell \rightarrow t'_\ell}$, which is a subpath of $P|_{s_\ell \rightarrow t_\ell}$, which implies $x, y \in V'_\ell$ and $\ell(x, y) \geq \ell$ by H1(ℓ) and H2(ℓ). Thus, in any case, $\ell(x, y) \geq \ell$, $x, y \in V'_\ell$, and (x, y) is not a shortcut of some level $> \ell$. Hence, \bar{P} is a path in G'_ℓ . Now, consider an arbitrary edge (u, v) , $s_{\ell+1} \leq u < v \leq t_{\ell+1}$, on P . If (u, v) is a shortcut of some level $> \ell$, we directly have $\ell(u, v) \geq \ell + 1$. Otherwise, (u, v) is on \bar{P} as well. Since $s_{\ell+1} < v$, we have $d_P(s'_\ell, v) > r_\ell^{\rightarrow}(s'_\ell)$ (U2b). Moreover, S1 implies that \bar{P} is a shortest path in G'_ℓ and, in particular, $d_{\bar{P}}(s'_\ell, v) = w(\bar{P}|_{s'_\ell \rightarrow v}) = d_\ell(s'_\ell, v)$. Using the fact that $d_{\bar{P}}(s'_\ell, v) = d_P(s'_\ell, v)$, we obtain $d_\ell(s'_\ell, v) > r_\ell^{\rightarrow}(s'_\ell)$ and, thus, $v \notin \mathcal{N}_\ell^{\rightarrow}(s'_\ell)$.

Analogously, we have $u \notin \mathcal{N}_\ell^{\leftarrow}(t'_\ell)$. Hence, the definition of the highway network $G_{\ell+1}$ implies $(u, v) \in E_{\ell+1}$. Thus, $\ell(u, v) \geq \ell + 1$. \square

Definition A.18. A valid state is either a *final* state (if $u = \bar{u}$) or a *non-final* state (otherwise).

We pick any shortest s - t -path P . The state $(\text{ctr}(P), s, t)$ is the *initial* state. Since forward and backward search run completely independently of each other, any serialization of both search processes will yield exactly the same result. Therefore, in our proof, we are free to pick—without loss of generality—any order of forward and backward steps.

We assume that at first one forward and one backward iteration is performed, which implies that s and t are settled. At this point, the highway query algorithm is in the initial state. It is easy to see that the initial state is a valid state. Due to the following lemma, it is sufficient to prove that a final state is eventually reached.

LEMMA A.19. *Getting to a final state is equivalent to finding a shortest s - t -path.*

PROOF. $u = \bar{u}$ means that forward and backward search meet. Due to Lemma A.16, we can conclude that both u and \bar{u} are settled with the optimal distance (A_1), that is, $\overrightarrow{\delta}(u) = d_0(s, u)$ and $\overleftarrow{\delta}(\bar{u}) = d_0(\bar{u}, t)$. Since $u = \bar{u}$ lies on a shortest path (due to S1), we have $d(s, t) = d_0(s, u) + d_0(\bar{u}, t)$. Line 6 implies $d' \leq \overrightarrow{\delta}(u) + \overleftarrow{\delta}(\bar{u}) = d(s, t)$. In fact, this means that the algorithm returns $d' = d(s, t)$, since this is already optimal. \square

Definition A.20. For a valid state $z = (P, u, \bar{u})$, the forward direction is said to be *blocked* if $p \leq u$. Analogously, the backward direction is blocked if $\bar{u} \leq p$.

LEMMA A.21. *For a nonfinal state $z = (P, u, \bar{u})$, at most one direction is blocked.*

PROOF. Since z is a nonfinal state, we have $u < \bar{u}$, which implies $u < p$ or $p < \bar{u}$. \square

Definition A.22. The *rank* $\rho(z)$ of a state $z = (P, u, \bar{u})$ is $|\{x \in P \mid u \leq x \leq \bar{u}\}|$.

LEMMA A.23. *From any nonfinal state $z = (P, u, \bar{u})$, another valid state z^+ is reached at some point such that $\rho(z^+) < \rho(z)$.*

PROOF. We pick any nonblocked direction—due to Lemma A.21, we know that there is at least one such direction. Subsequently, we assume that the forward direction was picked; the backward direction can be dealt with analogously.

We have $u < p$ and observe that all prerequisites of Lemma A.12 are fulfilled due to Lemmas A.17 and A.16. Hence, we can conclude that the edge $(u, v := \text{succ}(u))$ is not skipped, but relaxed. Thus, since P is a shortest path (S1), v can be reached with the optimal distance due to Lemma A.13 (A_1). The fact that the algorithm terminates (Lemma A.3) implies that the queue \overrightarrow{Q} gets empty at some point, that is, every element has been deleted from \overrightarrow{Q} . In particular, we can conclude that v is deleted at some point. Since v has been reached with the optimal distance, it will also be settled with the optimal distance (due to the specification of the decreaseKey operation, which guarantees that tentative distances are never increased). Let P' denote the path from s to v in the forward search tree. We set $z^+ := (P^+ := P' \circ P|_{v \rightarrow t}, v, \bar{u})$. We have $w(P^+) = w(P') + w(P|_{v \rightarrow t}) = d_0(s, v) + d_0(v, t) = d_0(s, t)$ (\Rightarrow S1). S2 is fulfilled, since $P^+|_{v \rightarrow \bar{u}}$ is a subpath of $P|_{u \rightarrow \bar{u}}$. S3 holds due to the construction of P^+ . Hence, z^+ is valid. Furthermore, $\rho(z^+) = \rho(z) - 1$. \square

THEOREM A.24. *The highway query algorithm finds a shortest s - t -path.*

PROOF. From Lemma A.23 and the fact that the codomain of the rank function is finite, it follows that eventually a final state is reached, which is equivalent to finding a shortest s - t -path due to Lemma A.19. \square

A.6. Distance Table Optimisation

To prove the correctness of the distance table optimization, we introduce the following new lemma and adapt a few definitions and proofs from Section A.5 to the new situation.

LEMMA A.25. *Consider a valid state $z = (P, u, \bar{u})$ with $u < s'_t$. When u 's edges are relaxed, neither the condition in Line 7a nor the condition in Line 11a is fulfilled.*

PROOF. Due to Lemma A.16, u has been appropriately settled with some key k . We distinguish between two cases.

Case 1. $u < s_L$. From $s_{\ell(u)} = s_{\bar{\ell}(u)} \leq u < s_L$ ($A_2(k, u)$, L2), it follows that $\ell(u) < L$ (U1). Hence, the condition in Line 7a is not fulfilled. Furthermore, we have $s_{\ell} \leq u < s_L$ after Line 9 has been executed (Lemma A.11: I1). Thus, $\ell < L$, which implies that the condition in Line 11a is not fulfilled as well.

Case 2. $s_L \leq u < s'_L$. First, we show that the condition in Line 7a is not fulfilled. We assume $\ell(u) = L$. (Otherwise, the condition cannot be fulfilled.) Due to $A_3(k, u)$, we have $\text{gap}(u) = \infty$. Hence, $\text{gap}' = r_{\ell(u)}^{\rightarrow}(u) = r_L^{\rightarrow}(u) = \infty_1$ by R1 since $u \notin V'_L$ (U3). Now, we prove that the condition in Line 11a is not fulfilled. We assume $\ell = L \wedge \ell > \ell(u)$. (Otherwise, the condition cannot be fulfilled.) Due to Line 9, we get $\text{gap} = r_{\ell}^{\rightarrow}(u) = r_L^{\rightarrow}(u) = \infty_1$ (as above). \square

Definition A.18'. A valid state is either a *final* state (if $u = \bar{u}$ or $s'_L \leq u \wedge \bar{u} \leq t'_L$) or a *nonfinal* state (otherwise).

LEMMA A.19. *Getting to a final state is equivalent to finding a shortest s-t-path.*

PROOF. In the proof of this lemma in Section A.5, we have already dealt with the case $u = \bar{u}$. Now, consider the new case $u < \bar{u} \wedge s'_L \leq u \wedge \bar{u} \leq t'_L$. We show that s'_L is added to the set \vec{T} . Since $s'_L \leq u$, s'_L has been appropriately settled with some key k (due to Lemma A.16). We consider the attempt to relax the edge $(s'_L, v := \text{succ}(s'_L))$ and distinguish between two cases.

Case 1. $s_L = s'_L$. $\ell = \bar{\ell}(v)$ (I2), $s_L = s'_L < v$, and $\bar{\ell}(v) \leq L$ (L1) imply $\ell = \bar{\ell}(v) = L$. Furthermore, $A_2(k, s'_L)$ and the assumption of Case 1 yield $\ell(s'_L) = \bar{\ell}(s'_L) < L = \ell$. In addition, $\text{gap} = \infty_2 \neq \infty_1$ by I3 (since $s'_L < v$), the fact that $s'_L \in V'_L$ (U4), and R2. Hence, the condition in Line 11a is fulfilled so that s'_L is added to \vec{T} .

Case 2. $s_L < s'_L$. By $A_2(k, s'_L)$, $A_3(k, s'_L)$, the assumption of Case 2, and $\bar{\ell}(s'_L) \leq L$ (L1), we get $\ell(s'_L) = \bar{\ell}(s'_L) = L$ and $\text{gap}(s'_L) = \infty$. Thus, $\text{gap}' = r_L^{\rightarrow}(s'_L) = \infty_2 \neq \infty_1$ (R2). Hence, the condition in Line 7a is fulfilled so that s'_L is added to \vec{T} .

Analogously, we can prove that t'_L is added to the set \overleftarrow{T} . Since P is a highway path (due to Lemma A.17), the subpath $P|_{s'_L \rightarrow t'_L}$ is a path in G'_L and, thus, $d_0(s'_L, t'_L) = d_L(s'_L, t'_L)$. Hence, $w(P) = d_0(s, s'_L) + d_L(s'_L, t'_L) + d_0(t'_L, t)$ is the length of a shortest s-t-path and, since the algorithm finds a path with a length $\leq \vec{\delta}(s'_L) + d_L(s'_L, t'_L) + \overleftarrow{\delta}(t'_L)$ and since $\vec{\delta}(s'_L) = d_0(s, s'_L)$ and $\overleftarrow{\delta}(t'_L) = d_0(t'_L, t)$ (due to Lemma A.16: A1), we can conclude that a shortest s-t-path is found. \square

Definition A.20'. For a valid state $z = (P, u, \bar{u})$, the forward direction is said to be *blocked* if $p \leq u$ or $s'_L \leq u$. Analogously, the backward direction is blocked if $\bar{u} \leq p$ or $\bar{u} \leq t'_L$.

LEMMA A.21. *For a nonfinal state $z = (P, u, \bar{u})$, at most one direction is blocked.*

PROOF. Since z is a nonfinal state, we have $u < \bar{u}$ and $(u < s'_L \vee t'_L < \bar{u})$. To obtain a contradiction, let us assume that both directions are blocked, that is, $(p \leq u$ or $s'_L \leq u)$ and $(\bar{u} \leq p$ or $\bar{u} \leq t'_L)$. Consider the case $p \leq u$ and $\bar{u} \leq t'_L$. Hence, $p \leq u < \bar{u} \leq t'_L$. Due to M3, we can conclude that $s'_L \leq p \leq u$. Since $s'_L \leq u$ and $\bar{u} \leq t'_L$, we have a contradiction. The remaining three cases are analogous or straightforward. \square

LEMMA A.23. *From any nonfinal state $z = (P, u, \bar{u})$, another valid state z^+ is reached at some point such that $\rho(z^+) < \rho(z)$.*

PROOF. The proof of this lemma in Section A.5 still works, since the added two lines (7a and 11a) have no effect due to Definition A.20' and Lemma A.25. \square

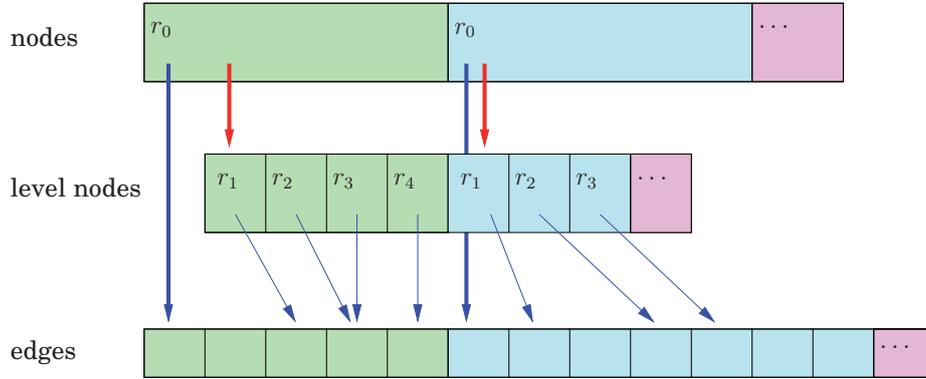


Fig. 16. An adjacency array, extended by a level-node layer.

Table V. Preprocessing and Query Performance for the European Road Network Depending on the Contraction Rate c and the Neighbourhood Size H

contr. rate c	nbh. size H	PREPROCESSING			QUERY		
		time (min)	\emptyset overhead/ node (byte)	\emptyset deg.	time (ms)	#settled nodes	#relaxed edges
0.5	30	83	30	3.2	391.73	472,326	1,023,944
	40	83	28	3.2	267.57	334,287	711,082
	50	87	27	3.2	188.55	242,787	506,543
	60	86	27	3.2	135.27	177,558	362,748
	70	87	26	3.2	101.36	135,560	271,324
	80	89	26	3.1	73.40	99,857	196,150
	90	87	25	3.1	55.02	75,969	146,247
1.0	30	15	28	3.7	5.48	6,396	23,612
	40	15	28	3.7	2.62	3,033	11,315
	50	17	27	3.6	2.13	2,406	8,902
	60	18	27	3.6	1.93	2,201	8,001
	70	19	26	3.6	1.80	2,151	7,474
	80	20	26	3.6	1.79	2,193	7,392
	90	22	26	3.6	1.78	2,221	7,268
1.5	30	11	28	3.8	1.93	1,830	9,281
	40	12	28	3.8	1.72	1,628	7,672
	50	13	27	3.7	1.56	1,593	6,975
	60	14	27	3.7	1.53	1,645	6,697
	70	15	27	3.7	1.51	1,673	6,590
	80	17	27	3.7	1.51	1,726	6,719
	90	18	27	3.7	1.54	1,782	6,655
2.0	30	11	29	4.0	1.85	1,542	8,913
	40	11	29	3.9	1.64	1,475	7,646
	50	12	28	3.9	1.48	1,470	6,785
	60	14	28	3.8	1.46	1,506	6,650
	70	15	28	3.8	1.45	1,547	6,649
	80	16	27	3.8	1.49	1,611	6,935
	90	17	27	3.8	1.53	1,675	6,988
2.5	30	11	30	4.1	1.96	1,489	9,175
	40	11	29	4.0	1.70	1,453	7,822
	50	12	29	4.0	1.58	1,467	7,119
	60	14	29	3.9	1.57	1,493	7,035
	70	15	28	3.9	1.54	1,536	6,905
	80	16	28	3.9	1.55	1,583	7,094
	90	18	28	3.9	1.58	1,645	7,204

Note: We do not use a distance table, but repeat the construction process until the topmost level is empty or the hierarchy consists of 15 levels.

Table VI. Preprocessing and Query Performance for the European Road Network Depending on the Number of Levels and the Neighborhood Size H

#levels	nbh. size H	PREPROCESSING			QUERY		
		time (min)	\emptyset overhead/ node (byte)	\emptyset deg.	time (ms)	#settled nodes	#relaxed edges
5	40	14	60	3.9	0.67	691	2,398
	50	13	40	3.9	0.77	818	2,892
	60	14	32	3.8	0.87	938	3,361
	70	15	30	3.8	0.96	1,058	3,837
	80	16	28	3.8	1.05	1,165	4,278
	90	17	28	3.8	1.13	1,269	4,697
6	30	12	48	4.0	0.75	709	2,531
	40	11	33	3.9	0.87	867	3,171
	50	12	29	3.9	0.99	1,015	3,759
	60	13	28	3.8	1.10	1,157	4,299
	70	15	28	3.8	1.21	1,292	4,837
	80	16	28	3.8	1.30	1,414	5,311
7	90	17	27	3.8	1.40	1,521	5,817
	30	10	34	4.0	0.93	852	3,195
	40	11	29	3.9	1.07	1,025	3,894
	50	12	28	3.9	1.20	1,187	4,538
	60	13	28	3.8	1.32	1,344	5,166
	70	15	28	3.8	1.39	1,462	5,689
8	80	16	27	3.8	1.47	1,578	6,179
	90	18	27	3.8	1.53	1,668	6,661
	30	10	30	4.0	1.14	991	3,853
	40	11	29	3.9	1.27	1,171	4,624
	50	12	28	3.9	1.36	1,321	5,283
	60	14	28	3.8	1.43	1,455	5,887
9	70	15	28	3.8	1.46	1,546	6,338
	80	16	27	3.8	1.48	1,611	6,935
	90	18	27	3.8	1.53	1,675	6,988
	30	10	30	4.0	1.35	1,123	4,532
	40	11	29	3.9	1.45	1,289	5,338
	50	12	28	3.9	1.48	1,417	5,931
10	60	14	28	3.8	1.47	1,506	6,429
	70	15	28	3.8	1.46	1,547	6,649
	30	10	29	4.0	1.54	1,241	5,214
	40	11	29	3.9	1.57	1,380	6,012
11	50	12	28	3.9	1.51	1,468	6,470
	60	14	28	3.8	1.46	1,506	6,650
	30	10	29	4.0	1.67	1,326	5,847
	40	11	29	3.9	1.65	1,445	6,627
	50	13	28	3.9	1.49	1,470	6,785

Note: In the topmost level, a distance table is used.

B. IMPLEMENTATION

The graph is represented as *adjacency array*, which is a very space-efficient data structure that allows fast traversal of the graph. There are two arrays, one for the nodes and one for the edges. The edges (u, v) are grouped by the source node u and store only the ID of the target node v and the weight $w(u, v)$. Each node u stores the index of its first outgoing edge in the edge array. In order to allow a search in the backward graph, we have to store an edge (u, v) also as backward edge (v, u) in the edge group of node v . In order to distinguish between forward and backward edges, each edge has a forward and a backward flag. By this means, we can also store two-way edges $\{u, v\}$ (which make up the large majority of all edges in a real-world road network) in a space-efficient way: We keep only one copy of (u, v) and one copy of (v, u) , in each case setting both direction flags.

The basic adjacency array has to be extended in order to incorporate the level data that is specific to highway hierarchies. In addition to the index of the first outgoing edge, each node u stores its level-0 neighborhood radius $r_0(u)$. Moreover, for each node u , all outgoing edges (u, v) are grouped by their level $\ell(u, v)$. Between the node and the edge array, we insert another layer: for each node u and each level $\ell > 0$ that u belongs to, there is a *level node* u_ℓ that stores the radius $r_\ell(u)$ and the index of the first outgoing edge (u, v) in level ℓ . All level nodes are stored in a single array. Each node u keeps the index of the level node u_1 . Figure 16 illustrates the graph representation.

To obtain a robust implementation, we include extensive consistency checks in assertions and perform experiments that are checked against reference implementations, that is, queries are checked against Dijkstra’s algorithm and the fast preprocessing algorithm is checked against a naive implementation.

C. EXPERIMENTS

In addition to the experiments presented in Section 6.3, we have considered many more combinations of neighborhood size, contraction rate, and number of levels. The results are given in Tables V and VI.

REFERENCES

- ABRAHAM, I., DELLING, D., GOLDBERG, A. V., AND WERNECK, R. F. 2010. A hub-based labeling algorithm for shortest paths on road networks. Tech. rep. MSR-TR-2010-165, Microsoft Research.
- BAST, H., FUNKE, S., MATLJEVIC, D., SANDERS, P., AND SCHULTES, D. 2007. In transit to constant time shortest-path queries in road networks. In *Proceedings of the Workshop on Algorithm Engineering and Experiments (ALENEX)*. 46–59.
- BATZ, G., DELLING, D., SANDERS, P., AND VETTER, C. 2009. Time-dependent contraction hierarchies. In *Proceedings of the 10th Workshop on Algorithm Engineering and Experiments (ALENEX)*. 97–105. <http://algo2.iti.uni-karlsruhe.de/1222.php>.
- BAUER, R. AND DELLING, D. 2008. SHARC: fast and robust unidirectional routing. In *Proceedings of the Workshop on Algorithm Engineering and Experiments (ALENEX’08)*.
- BAUER, R., DELLING, D., SANDERS, P., SCHIEFERDECKER, D., SCHULTES, D., AND WAGNER, D. 2008. Combining hierarchical and goal-directed speed-up techniques for Dijkstra’s algorithm. In *Proceedings of the 7th Workshop on Experimental Algorithms (WEA’08)*. Lecture Notes in Computer Science, vol. 5038, Springer, 303–318.
- BAUER, R., DELLING, D., AND WAGNER, D. 2007. Experimental study on speed-up techniques for timetable information systems. In *Proceedings of the 7th Workshop on Algorithmic Approaches for Transportation Modeling, Optimization, and Systems (ATMOS’07)*.
- CORMEN, T. H., LEISERSON, C. E., RIVEST, R. L., AND STEIN, C. 2001. *Introduction to Algorithms*. 2nd Ed. MIT Press.
- DELLING, D. 2008. Time-dependent SHARC-routing. In *Proceedings of the 16th Annual European Symposium on Algorithms (ESA’08)*. Lecture Notes in Computer Science, vol. 5193, Springer, 332–343.
- DELLING, D., SANDERS, P., SCHULTES, D., AND WAGNER, D. 2006. Highway hierarchies star. In *Proceedings of the 9th DIMACS Implementation Challenge*.
- DELLING, D., SANDERS, P., SCHULTES, D., AND WAGNER, D. 2009. Engineering route planning algorithms. In *Proceedings of the Algorithmics of Large and Complex Networks*. Lecture Notes in Computer Science, vol. 5515, Springer, 117–139.
- DIJKSTRA, E. W. 1959. A note on two problems in connexion with graphs. *Numerische Mathematik* 1, 269–271.
- DIMACS 9TH IMPLEMENTATION CHALLENGE. 2006. Shortest Paths. <http://www.dis.uniroma1.it/~challenge9/>.
- FAKCHAROENPHOL, J. AND RAO, S. 2001. Planar graphs, negative weight edges, shortest paths, and near linear time. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*. 232–241.
- GEISBERGER, R., SANDERS, P., SCHULTES, D., AND DELLING, D. 2008. Contraction hierarchies: faster and simpler hierarchical routing in road networks. In *Proceedings of the 7th Workshop on Experimental Algorithms*. Lecture Notes in Computer Science, vol. 5038, Springer, 319–333.
- GEISBERGER, R., SANDERS, P., SCHULTES, D., AND VETTER, C. 2012. Exact routing in large road networks using contraction hierarchies. *Transport. Sci.*

- GOLDBERG, A., KAPLAN, H., AND WERNECK, R. F. 2006. Reach for A^* : efficient point-to-point shortest path algorithms. In *Proceedings of the Workshop on Algorithm Engineering and Experiments (ALENEX'06)*. 129–143.
- GOLDBERG, A. V. AND HARRELSON, C. 2005. Computing the shortest path: A^* meets graph theory. In *Proceedings of the 16th ACM-SIAM Symposium on Discrete Algorithms*. 156–165.
- GOLDBERG, A. V., KAPLAN, H., AND WERNECK, R. F. 2007. Better landmarks within reach. In *Proceedings of the 6th Workshop on Experimental Algorithms (WEA'07)*. Lecture Notes in Computer Science, vol. 4525, Springer, 38–51.
- GOLDBERG, A. V. AND WERNECK, R. F. 2005. Computing point-to-point shortest paths from external memory. In *Proceedings of the Workshop on Algorithm Engineering and Experiments (ALENEX'05)*. 26–40.
- GUTMAN, R. 2004. Reach-based routing: A new approach to shortest path algorithms optimized for road networks. In *Proceedings of the Workshop on Algorithm Engineering and Experiments (ALENEX'04)*. 100–111.
- HART, P. E., NILSSON, N. J., AND RAPHAEL, B. 1968. A formal basis for the heuristic determination of minimum cost paths. *IEEE Trans. Syst. Sci. Cybern.* 4, 2, 100–107.
- ISHIKAWA, K., OGAWA, M., AZUME, S., AND ITO, T. 1991. Map navigation software of the electro multivision of the '91 Toyota Soarer. In *Proceedings of the IEEE International Conference on Vehicle Navigation Information System*. 463–473.
- JAGADEESH, G., SRIKANTHAN, T., AND QUEK, K. 2002. Heuristic techniques for accelerating hierarchical routing on road networks. *IEEE Trans. Intell. Transport. Syst.* 3, 4, 301–309.
- KNOPP, S., SANDERS, P., SCHULTES, D., SCHULZ, F., AND WAGNER, D. 2007. Computing many-to-many shortest paths using highway hierarchies. In *Proceedings of the Workshop on Algorithm Engineering and Experiments (ALENEX'07)*.
- KÖHLER, E., MÖHRING, R. H., AND SCHILLING, H. 2005. Acceleration of shortest path and constrained shortest path computation. In *Proceedings of the 4th International Workshop on Efficient and Experimental Algorithms (WEA'05)*.
- KÖHLER, E., MÖHRING, R. H., AND SCHILLING, H. 2006. Fast point-to-point shortest path computations with arc-flags. In *Proceedings of the 9th DIMACS Implementation Challenge*.
- LAUTHER, U. 2004. An extremely fast, exact algorithm for finding shortest paths in static networks with geographical background. In *Proceedings of the Geoinformation und Mobilität – von der Forschung zur praktischen Anwendung*. Vol. 22. IFGI prints, Institut für Geoinformatik, Münster, 219–230.
- LAUTHER, U. 2006. An experimental evaluation of point-to-point shortest path calculation on roadnetworks with precalculated edge-flags. In *Proceedings of the 9th DIMACS Implementation Challenge*.
- MAUE, J., SANDERS, P., AND MATIJEVIC, D. 2006. Goal directed shortest path queries using Precomputed Cluster Distances. In *Proceedings of the 5th Workshop on Experimental Algorithms (WEA'06)*. Lecture Notes in Computer Science, vol. 4007, Springer, 316–328.
- MAUE, J., SANDERS, P., AND MATIJEVIC, D. 2009. Goal directed shortest path queries using Precomputed Cluster Distances. *ACM J. Exp. Algor.* (submitted for special issue on WEA'06).
- MÖHRING, R., SCHILLING, H., SCHÜTZ, B., WAGNER, D., AND WILLHALM, T. 2005. Partitioning graphs to speed up Dijkstra's algorithm. In *Proceedings of the 4th International Workshop on Efficient and Experimental Algorithms (WEA'05)*. 189–202.
- MÖHRING, R. H., SCHILLING, H., SCHÜTZ, B., WAGNER, D., AND WILLHALM, T. 2007. Partitioning graphs to speed up dijkstra's algorithm. *ACM J. Exp. Algor.* 11.
- MÜLLER, K. 2005. Berechnung kürzester Pfade unter Beachtung von Abbiegeverboten. Student Research Project, Universität Karlsruhe (TH).
- MULLER, L. F. AND ZACHARIASEN, M. 2007. Fast and compact oracles for approximate distances in planar graphs. In *Proceedings of the 15th European Symposium on Algorithms (ESA'07)*. Lecture Notes in Computer Science, vol. 4698, Springer, 657–668.
- NANNICINI, G., BAPTISTE, P., BARBIER, G., KROB, D., AND LIBERTI, L. 2010. Fast paths in large-scale dynamic road networks. *Comput. Optim. Appl.* 45, 1, 143–158.
- R DEVELOPMENT CORE TEAM. 2004. R: A language and environment for statistical computing. <http://www.r-project.org>.
- SANDERS, P. AND SCHULTES, D. 2005. Highway hierarchies hasten exact shortest path queries. In *Proceedings of the 13th European Symposium on Algorithms (ESA'05)*. Lecture Notes in Computer Science, vol. 3669, Springer, 568–579.
- SANDERS, P. AND SCHULTES, D. 2006. Engineering highway hierarchies. In *Proceedings of the 14th European Symposium on Algorithms (ESA'06)*. Lecture Notes in Computer Science, vol. 4168, Springer, 804–816.

- SCHMID, W. 2000. Berechnung kürzester wege in straßennetzen mit wegeverboten. Ph.D. thesis, Universität Stuttgart.
- SCHULTES, D. 2008. Route planning in road networks. Ph.D. thesis, Universität Karlsruhe (TH).
- SCHULTES, D. AND SANDERS, P. 2007. Dynamic highway-node routing. In *Proceedings of the 6th Workshop on Experimental Algorithms (WEA'07)*. Lecture Notes in Computer Science, vol. 4525, Springer, 66–79.
- SCHULZ, F., WAGNER, D., AND WEIHE, K. 1999. Dijkstra's algorithm on-line: An empirical case study from public railroad transport. In *Proceedings of the 3rd Workshop on Algorithm Engineering (WAE'99)*. Lecture Notes in Computer Science, vol. 1668, Springer, 110–123.
- SCHULZ, F., WAGNER, D., AND ZAROLIAGIS, C. D. 2002. Using multi-level graphs for timetable information. In *Proceedings of the Workshop on Algorithm Engineering and Experiments (ALENEX'02)*. Lecture Notes in Computer Science, vol. 2409, Springer, 43–59.
- SKIENA, S. S. 1998. *The Algorithm Design Manual*. Springer.
- THORUP, M. 2001. Compact oracles for reachability and approximate distances in planar digraphs. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*. 242–251.
- U.S. CENSUS BUREAU, WASHINGTON, DC. 2002. UA Census 2000 TIGER/Line Files. http://www.census.gov/geol/www/tiger/tigerua/ua_tgr2k.html.
- WAGNER, D. AND WILLHALM, T. 2003. Geometric speed-up techniques for finding shortest paths in large sparse graphs. In *Proceedings of the 11th European Symposium on Algorithms (ESA'03)*. Lecture Notes in Computer Science, vol. 2832, Springer, 776–787.

Repository KITopen

Dies ist ein Postprint/begutachtetes Manuskript.

Empfohlene Zitierung:

Sanders, P.; Schultes, D.
[Engineering highway hierarchies](#).
2012. Journal of experimental algorithmics, 17
[doi:10.5445/IR/1000139529](https://doi.org/10.5445/IR/1000139529)

Zitierung der Originalveröffentlichung:

Sanders, P.; Schultes, D.
[Engineering highway hierarchies](#).
2012. Journal of experimental algorithmics, 17 (1), Article no: 1.6.
[doi:10.1145/2133803.2330080](https://doi.org/10.1145/2133803.2330080)

Lizenzinformationen: [KITopen-Lizenz](#)