

Bidirectional A^* Search for Time-Dependent Fast Paths^{*}

Giacomo Nannicini^{1,2}, Daniel Delling³, Leo Liberti¹, Dominik Schultes³

¹ *LIX, École Polytechnique, F-91128 Palaiseau, France*

`{giacomon,liberti}@lix.polytechnique.fr`

² *Mediamobile, 10 rue d'Oradour sur Glane, Paris, France*

`giacomo.nannicini@v-traffic.com`

³ *Universität Karlsruhe (TH), 76128 Karlsruhe, Germany*

`{delling,schultes}@ira.uka.de`

Abstract. The computation of point-to-point shortest paths on time-dependent road networks has many practical applications, but there have been very few works that propose efficient algorithms for large graphs. One of the difficulties of route planning on time-dependent graphs is that we do not know the exact arrival time at the destination, hence applying bidirectional search is not straightforward; we propose a novel approach based on A^* with landmarks (ALT) that starts a search from both the source and the destination node, where the backward search is used to bound the set of nodes that have to be explored by the forward search. Extensive computational results show that this approach is very effective in practice if we are willing to accept a small approximation factor, resulting in a speed-up of several times with respect to Dijkstra's algorithm while finding only slightly suboptimal solutions.

1 Introduction

We consider the TIME-DEPENDENT SHORTEST PATH PROBLEM (TDSPP): given a directed graph $G = (V, A)$, a source node $s \in V$, a destination node $t \in V$, an interval of time instants T , a departure time $\tau_0 \in T$ and a time-dependent arc weight function $c : A \times T \rightarrow \mathbb{R}_+$, find a path $p = (s = v_1, \dots, v_k = t)$ in G such that its *time-dependent cost* $\gamma_{\tau_0}(p)$, defined recursively as follows:

$$\gamma_{\tau_0}(v_1, v_2) = c(v_1, v_2, \tau_0) \quad (1)$$

$$\gamma_{\tau_0}(v_1, \dots, v_i) = \gamma_{\tau_0}(v_1, \dots, v_{i-1}) + c(v_{i-1}, v_i, \tau_0 + \gamma_{\tau_0}(v_1, \dots, v_{i-1})) \quad (2)$$

for all $2 \leq i \leq k$, is minimum. We also consider a function $\lambda : A \rightarrow \mathbb{R}_+$ which has the following property:

$$\forall (u, v) \in A, \tau \in T \quad (\lambda(u, v) \leq c(u, v, \tau)).$$

^{*} Partially supported by the Future and Emerging Technologies Unit of EC (IST priority – 6th FP), under contract no. FP6-02123502 (project ARRIVAL), and by DFG grant SA 933/1-3.

In other words, $\lambda(u, v)$ is a lower bound on the travelling time of arc (u, v) for all time instants in T . In practice, this can easily be computed, given an arc length and the maximum allowed speed on that arc. We naturally extend λ to be defined on paths, i.e. $\lambda(p) = \sum_{(v_i, v_j) \in p} \lambda(v_i, v_j)$.

In this paper, we propose a novel algorithm for the TDSPP based on a bidirectional A^* algorithm. Since the arrival time is not known in advance (so c cannot be evaluated on the arcs adjacent to the destination node), our backward search occurs on the graph weighted by the lower bounding function λ . This is used for bounding the set of nodes that will be explored by the forward search.

Many ideas have been proposed for the computation of point-to-point shortest paths on static graphs (see [1, 2] for a review), and there are algorithms capable of finding the solution in a matter of a few microseconds [3]; adaptations of those ideas for dynamic scenarios, i.e. where arc costs are updated at regular intervals, have been tested as well (see [4–6]).

Much less work has been undertaken on the time-dependent variant of the shortest paths problem; this problem has been first addressed by [7] (a good review of this paper can be found in [8], p. 407): Dijkstra’s algorithm [9] is extended to the dynamic case through a recursion formula based on the assumption that the network $G = (V, A)$ has the FIFO property. The FIFO property is also called the *non-overtaking property*, because it basically states that if A leaves u at time τ_0 and B at time $\tau_1 > \tau_0$, B cannot arrive at v before A using the arc (u, v) . The FIFO assumption is usually necessary in order to maintain an acceptable level of complexity: the TDSPP in FIFO networks is polynomially solvable [10], while it is NP-hard in non-FIFO networks [11]. Given source and destination nodes s and t , the problem of maximizing the departure time from node s with a given arrival time at node t is equivalent to the TDSPP (see [12]).

Goal-directed search, also called A^* [13], has been adapted to work on all the previously described scenarios; an efficient version for the static case has been presented in [14], and then developed and improved in [15]. Those ideas have been used in [4] on dynamic graphs as well, while the time-dependent case on graphs with the FIFO property has been addressed in [16] and [4].

Moreover, the recently developed SHARC-algorithm [17] allows fast *unidirectional* shortest-path calculations in large scale networks. Due to its unidirectional nature, it can easily be used in a time-dependent scenario. However, in that case SHARC cannot guarantee to find the optimal solution.

The rest of this paper is organised as follows. In Section 2 we describe A^* search and the ALT algorithm, which are needed for our method. In Section 3 we describe the foundations of our idea, and present an adaptation of the ALT algorithm based on it. In Section 4 we formally prove our method’s correctness. In Section 5 we propose some improvements. In Section 6 we discuss computational experiments and provide computational results.

2 A^* with Landmarks

A^* is an algorithm for goal-directed search, similar to Dijkstra’s algorithm, but which adds a potential function to the priority key of each node in the queue. The A^* algorithm on static graphs can be described as follows. The potential function on a node v is an estimate of the distance to reach the target from v ; A^* then follows the same procedure as Dijkstra’s algorithm, but the use of this potential function has the effect of giving priority to nodes that are (supposedly) closer to target node t . If the potential function π is such that $\pi(v) \leq d(v, t) \forall v \in V$, where $d(v, t)$ is the distance from v to t , then A^* always finds shortest paths. A^* is guaranteed to explore no more nodes than Dijkstra’s algorithm: if $\pi(v)$ is a good approximation from below of the distance to target, A^* efficiently drives the search towards the destination node, and it explores considerably fewer nodes than Dijkstra’s algorithm; if $\pi(v) = 0 \forall v \in V$, A^* behaves exactly like Dijkstra’s algorithm. In [18] it is shown that A^* is equivalent to Dijkstra’s algorithm on a graph with reduced costs, i.e. $w_\pi(u, v) = w(u, v) - \pi(u) + \pi(v)$.

One way to compute the potential function, instead of using Euclidean distances, is to use the concept of *landmarks*. Landmarks have first been proposed in [14]; they are a preprocessing technique which is based on the triangular inequality. The basic principle is as follows: suppose we have selected a set $L \subset V$ of landmarks, and we have precomputed distances $d(v, \ell), d(\ell, v) \forall v \in V, \ell \in L$; the following triangle inequalities hold: $d(u, t) + d(t, \ell) \geq d(u, \ell)$ and $d(\ell, u) + d(u, t) \geq d(\ell, t)$. Therefore $\pi_t(u) = \max_{\ell \in L} \{d(u, \ell) - d(t, \ell), d(\ell, t) - d(\ell, u)\}$ is a lower bound for the distance $d(u, t)$, and it can be used as a potential function which preserves optimal paths. On a static graph (i.e. non time-dependent), bidirectional search can be implemented, using some care in modifying the potential function so that it is consistent for the forward and backward search (see [15]); the consistency condition states that $w_{\pi_f}(u, v)$ in G is equal to $w_{\pi_b}(v, u)$ in the reverse graph \overline{G} , where π_f and π_b are the potential functions for the forward and the backward search, respectively. Bidirectional A^* with the potential function described above is called ALT. It is straightforward to note that, if arc costs can only increase with respect to their original value, the potential function associated with landmarks is still a valid lower bound, even on a time-dependent graph; in [4] this idea is applied to a real road network in order to analyse the algorithm’s performances, but with a unidirectional search.

The choice of landmarks has a great impact on the size of the search space, as it severely affects the quality of the potential function. Several selection strategies exist, although none of them is optimal with respect to random queries, i.e., is guaranteed to yield the smaller search space for random source-destination pairs. The best known heuristics are *avoid* and *maxCover* [19].

3 Bidirectional Search on Time-Dependent Graphs

Our algorithm is based on restricting the scope of a time-dependent A^* search from the source using a set of nodes defined by a time-*independent* A^* search

from the destination, i.e. the backward search is a reverse search in G_λ , which corresponds to the graph G weighted by the lower bounding function λ .

Given a graph $G = (V, A)$ and source and destination vertices $s, t \in V$, the algorithm for computing the shortest time-dependent cost path p^* works in three phases.

1. A bidirectional A^* search occurs on G , where the forward search is run on the graph weighted by c with the path cost defined by (1)-(2), and the backward search is run on the graph weighted by the lower bounding function λ . All nodes settled by the backward search are included in a set M . Phase 1 terminates as soon as the two search scopes meet.
2. Suppose that $v \in V$ is the first vertex in the intersection of the heaps of the forward and backward search; then the time dependent cost $\mu = \gamma_{\tau_0}(p_v)$ of the path p_v going from s to t passing through v is an upper bound to $\gamma_{\tau_0}(p^*)$. In the second phase, both search scopes are allowed to proceed until the backward search queue only contains nodes whose associated key exceeds μ . In other words: let β be the key of the minimum element of the backward search queue; phase 2 terminates as soon as $\beta > \mu$. Again, all nodes settled by the backward search are included in M .
3. Only the forward search continues, with the additional constraint that only nodes in M can be explored. The forward search terminates when t is settled.

The pseudocode for this algorithm is given in Algorithm 1.

4 Correctness

We denote by $d(u, v, \tau)$ the length of the shortest path from u to v with departure time τ , and by $d_\lambda(u, v)$ the length of the shortest path from u to v on the graph G_λ . We have the following theorems.

Theorem 4.1. *Algorithm 1 computes the shortest time-dependent path from s to t for a given departure time τ_0 .*

Proof. The forward search of Algorithm 1 is exactly the same as the unidirectional version of the A^* algorithm during the first 2 phases, and thus it is correct; we have to prove that the restriction applied during phase 3 does not interfere with the correctness of the A^* algorithm.

Let μ be an upper bound on the cost of the shortest path; in particular, this can be the cost $\gamma_{\tau_0}(p_v)$ of the $s \rightarrow t$ path passing through the first meeting point v of the forward and backward search. Let β be the smallest key of the backward search priority queue at the end of phase 2. Suppose that Algorithm 1 is not correct, i.e. it computes a sub-optimal path. Let p^* be the shortest path from s to t with departure time τ_0 , and let u be the first node on p^* which is not explored by the forward search; by phase 3, this implies that $u \notin M$, i.e. u has not been settled by the backward search during the first 2 phases of Algorithm 1. Hence, we have that $\beta \leq \pi_b(u) + d_\lambda(u, t)$; then we have the chain $\gamma_{\tau_0}(p^*) \leq \mu < \beta \leq \pi_b(u) + d_\lambda(u, t) \leq d_\lambda(s, u) + d_\lambda(u, t) \leq d(s, u, \tau_0) + d(u, t, d(s, u, \tau_0)) = \gamma_{\tau_0}(p^*)$, which is a contradiction. \square

Algorithm 1 Compute the shortest time-dependent path from s to t with departure time τ_0

```

1:  $\vec{Q}$ .insert( $s, 0$ );  $\overleftarrow{Q}$ .insert( $t, 0$ );  $M := \emptyset$ ;  $\mu := +\infty$ ;  $done := \mathbf{false}$ ;  $phase := 1$ .
2: while  $\neg done$  do
3:   if  $(phase = 1) \vee (phase = 2)$  then
4:      $\leftrightarrow \in \{\rightarrow, \leftarrow\}$ 
5:   else
6:      $\leftrightarrow := \rightarrow$ 
7:      $u := \vec{Q}$ .extractMin()
8:     if  $(u = t) \wedge (\leftrightarrow = \rightarrow)$  then
9:        $done := \mathbf{true}$ 
10:    continue
11:   if  $(phase = 1) \wedge (u.dist^{\leftrightarrow} + u.dist^{\leftarrow} < \infty)$  then
12:      $\mu := u.dist^{\leftrightarrow} + u.dist^{\leftarrow}$ 
13:      $phase := 2$ 
14:   if  $(phase = 2) \wedge (\leftrightarrow = \leftarrow) \wedge (\mu < u.key^{\leftarrow})$  then
15:      $phase := 3$ 
16:   continue
17:   for all arcs  $(u, v) \in \vec{A}$  do
18:     if  $\leftrightarrow = \leftarrow$  then
19:        $M.insert(u)$ 
20:     else if  $(phase = 3) \wedge (v \notin M)$  then
21:       continue;
22:     if  $(v \in \overleftarrow{Q})$  then
23:       if  $u.dist^{\leftrightarrow} + c(u, v, u.dist^{\leftrightarrow}) < v.dist^{\leftrightarrow}$  then
24:          $\overleftarrow{Q}$ .decreaseKey( $v, u.dist^{\leftrightarrow} + c(u, v, u.dist^{\leftrightarrow}) + \overleftarrow{\pi}(v)$ )
25:       else
26:          $\overleftarrow{Q}$ .insert( $v, u.dist^{\leftrightarrow} + c(u, v, u.dist^{\leftrightarrow}) + \overleftarrow{\pi}(v)$ )
27: return  $t.dist^{\rightarrow}$ 

```

Theorem 4.2. *Let p^* be the shortest path from s to t . If the condition to switch to phase 3 is $\mu < K\beta$ for a fixed parameter K , then Algorithm 1 computes a path p from s to t such that $\gamma_{\tau_0}(p) \leq K\gamma_{\tau_0}(p^*)$ for a given departure time τ_0 .*

5 Improvements

The basic version of the algorithm can be enhanced by making use of the following results.

Proposition 5.1. *During phase 2 the backward search does not need to explore nodes that have already been settled by the forward search.*

We can take advantage of the fact that the backward search is used only to bound the set of nodes explored by the forward search, i.e. the backward search does not have to compute shortest paths. This means that we can tighten the bounds used by the backward search, as long as they are still valid lower

bounds, even if doing so would result in an A^* backward search that computes suboptimal distances.

Proposition 5.2. *At a given iteration, let v be the last node settled by the forward search. Then, for each node w which has not been settled by the forward search, $d(s, v, \tau_0) + \pi_f(v) - \pi_f(w) \leq d(s, w, \tau_0)$.*

Let v be as in Prop. 5.2, and w a node which has not been settled by the forward search. Prop. 5.2 suggests that we can use

$$\pi_b^*(w) = \max\{\pi_b(w), d(s, v, \tau_0) + \pi_f(v) - \pi_f(w)\} \quad (3)$$

as a lower bound to $d(s, w, \tau_0)$ during the backward search. However, to prove the algorithm’s correctness when using π_b^* we must assume that the node v used in (3) is fixed at each backward search iteration. Thus, we do the following: we set up 10 checkpoints during the query; when a checkpoint is reached, the node v used to compute (3) is updated, and the backward search queue is flushed and filled again using the updated π_b^* . This is enough to guarantee correctness. The checkpoints are computed comparing the initial lower bound $\pi_f(t)$ and the current distance from the source node, both for the forward search.

6 Experiments

In this section, we present an extensive experimental evaluation of our time-dependent ALT algorithm. Our implementation is written in C++ using solely the STL. As priority queue we use a binary heap. Our tests were executed on one core of an AMD Opteron 2218 running SUSE Linux 10.1. The machine is clocked at 2.6 GHz, has 16 GB of RAM and 2 x 1 MB of L2 cache. The program was compiled with GCC 4.1, using optimization level 3.

Unless otherwise stated, we use 16 maxcover landmarks [14], computed on the input graph using the lower bounding function λ to weight edges, and we use (3) as potential function for the backward search, with 10 checkpoints (see Section 5). When performing random s - t queries, the source s , target t , and the starting time τ_0 are picked uniformly at random and results are based on 10 000 queries.

Inputs. We tested our algorithm on two different road networks: the road network of Western Europe provided by PTV AG for scientific use, which has approximately 18 million vertices and 42.6 million arcs, and the road network of the US, taken from the TIGER/Line Files, with 23.9 million vertices and 58.3 million arcs. A travelling time in uncongested traffic situation was assigned to each arc using that arc’s category (13 categories for Europe, 4 for US) to determine the travel speed.

Modeling Traffic. Unfortunately, we are not aware of a *large* publicly available real-world road network with time-dependent arc costs. Therefore, we have to use artificially generated costs. In order to model the time-dependent costs on each arc, we developed an heuristic algorithm, based on statistics gathered using real-world data on a limited-size road network; we used piecewise linear cost functions, with one breakpoint for each hour over a day. Arc costs are generated assigning, at each node, several random values that represent peak hour (i.e. hour with maximum traffic increase), duration and speed of traffic increase/decrease for a traffic jam; for each node, two traffic jams are generated, one in the morning and one in the afternoon. Then, for each arc in a node’s arc star, a *speed profile* is generated, using the traffic jam’s characteristics of the corresponding node, and assigning a random increase factor between 1.5 and 3 to represent that arc’s slowdown during peak hours with respect to uncongested hours. We do not assign speed profile to arcs that have both endpoints at nodes with level 0 in a pre-constructed Highway Hierarchy [20], and as a result those arcs will have the same travelling time value throughout the day; for all other arcs, we use the traffic jam values associated with the endpoint with smallest ID. The breakpoints of these speed profiles are stored in memory as a multiplication factor with respect to the speed in uncongested hours, which allows us to use only 7 bits for each breakpoint. We assume that all roads are uncongested between 11PM and 4AM, so that we do not need to store the corresponding breakpoints; as a result, we store all breakpoints using 16 additional bytes per edge. The travelling time of an arc at time τ is computed via linear interpolation of the two breakpoints that precede and follow τ .

This method was developed to ensure spatial coherency between traffic increases, i.e. if a certain arc is congested at a given time, then it is likely that adjacent arcs will be congested too. This is a basic principle of traffic analysis [21].

Random Queries. Table 1 reports the results of our bidirectional ALT variant on time-dependent networks for different approximation values K using the European and the US road network as input. For the European road network, preprocessing takes approximately 75 minutes and produces 128 *additional* bytes per node (for each node we have to store distances to and from all landmarks); for the US road network, the corresponding figures are 92 minutes and 128 bytes per node. For comparison, we also report the results on the same road network for the time-dependent versions of Dijkstra, unidirectional ALT, and the SHARC algorithm [17].

As the performed ALT-queries compute approximated results instead of optimal solutions, we record three different statistics to characterize the solution quality: error rate, average relative error, maximum relative error. By *error rate* we denote the percentage of computed suboptimal paths over the total number of queries. By *relative error* on a particular query we denote the relative percentage increase of the approximated solution over the optimum, computed as $\omega/\omega^* - 1$, where ω is the cost of the approximated solution computed by our algorithm and ω^* is the cost of the optimum computed by Dijkstra’s algorithm. We report

Table 1. Performance of the time-dependent versions of Dijkstra, unidirectional ALT, SHARC, and our bidirectional approach. For SHARC, we use approximation values of 1.001 and 1.002 (cf. [17] for details).

input	method	K	ERROR			QUERY			time [ms]
			rate	avg	max	# settled nodes			
						phase 1	phase 2	phase 3	
EUR	Dijkstra	-	0.0%	0.000%	0.00%	-	-	8 908 300	6 325.8
	uni-ALT	-	0.0%	0.000%	0.00%	-	-	2 192 010	1 775.8
	1.001-SHARC	-	57.1%	0.686%	34.31%	-	-	140 945	60.3
	1.002-SHARC	-	42.8%	0.583%	34.31%	-	-	930 251	491.4
	ALT	1.00	0.0%	0.000%	0.00%	125 068	2 784 540	3 117 160	3 399.3
		1.02	1.0%	0.003%	1.13%	125 068	2 154 900	2 560 370	2 723.3
		1.05	4.0%	0.029%	4.93%	125 068	1 333 220	1 671 630	1 703.6
		1.10	18.7%	0.203%	8.10%	125 068	549 916	719 769	665.1
		1.13	30.5%	0.366%	12.63%	125 068	340 787	447 681	385.5
		1.15	36.4%	0.467%	13.00%	125 068	265 328	348 325	287.3
		1.20	44.7%	0.652%	18.19%	125 068	183 899	241 241	185.3
		1.30	48.2%	0.804%	23.63%	125 068	141 358	186 267	134.6
		1.50	48.8%	0.844%	25.70%	125 068	130 144	172 157	121.9
		2.00	48.9%	0.886%	48.86%	125 068	125 071	165 650	115.7
USA	Dijkstra	-	0.0%	0.000%	0.00%	-	-	12 435 900	8 020.6
	uni-ALT	-	0.0%	0.000%	0.00%	-	-	2 908 170	2 403.9
	ALT	1.00	0.0%	0.000%	0.00%	272 790	4 091 050	4 564 030	4 534.2
		1.10	21.5%	0.135%	7.02%	272 790	633 758	829 176	656.3
		1.15	54.4%	0.402%	9.98%	272 790	312 575	405 699	289.6
		1.20	62.0%	0.482%	9.98%	272 790	278 345	359 190	251.1
		1.50	64.8%	0.506%	13.63%	272 790	272 790	351 865	247.5
	2.00	64.8%	0.506%	16.00%	272 790	272 791	351 854	246.8	

average and *maximum* values of this quantity over the set of all queries. We also report the number of nodes settled at the *end* of each phase of our algorithm, denoting them with the labels phase 1, phase 2 and phase 3.

As expected, we observe a clear trade-off between the quality of the computed solution and query performance. If we are willing to accept an approximation factor of $K = 2.0$, on the European road network queries are on average 55 times faster than Dijkstra’s algorithm, but almost 50% of the computed paths will be suboptimal and, although the average relative error is still small, in the worst case the approximated solution has a cost which is 50% larger than the optimal value. The reason for this poor solution quality is that, for such high approximation values, phase 2 is very short. As a consequence, nodes in the middle of the shortest path are not explored by our approach, and the meeting point of the two search scopes is far from being the optimal one. However, by decreasing the value of the approximation constant K we are able to obtain solutions that are very close to the optimum, and performance is significantly better than for unidirectional ALT or Dijkstra. In our experiments, it seems as if the best trade-off between quality and performance is achieved with an

approximation value of $K = 1.15$, which yields average query times smaller than 300 ms on both road networks with a maximum recorded relative error of 13% (on the European road network, while the corresponding figure is 9.98% for the US instance). By decreasing K to values < 1.05 it does not pay off to use the bidirectional variant any more, as the unidirectional variant of ALT is faster and is always correct.

Comparing results for $K > 1.15$ for the US with those for Europe, we observe that number of queries that return suboptimal paths increases, but the average and maximum error rates are smaller than the corresponding values on the European road network with the same values of K . Moreover, the speed-ups of our algorithm with respect to plain Dijkstra are lower on the US instance: the maximum recorded speed-up (for $K = 2.0$) is only of a factor 33. This behaviour has also been observed in the static scenario [4]. However, with $K = 1.15$, which is a good trade-off between quality and speed, query performance is very similar on both networks.

An interesting observation is that for $K = 2.0$ switching from a static to a time-dependent scenario increases query times only of a factor of ≈ 2 : on the European road network, in a static scenario, ALT-16 has query times of 53.6 ms (see [4]), while our time-dependent variant yields query times of 115 ms. We also note that for our bidirectional search there is an additional overhead which increases the time spent per node with respect to unidirectional ALT: on the European road network, using an approximation factor of $K = 1.05$ yields similar query times to unidirectional ALT, but the number of nodes settled by the bidirectional approach is almost 30% smaller. We suppose that this is due to the following facts: in the bidirectional approach, one has to check at each iteration if the current node has been settled in the opposite direction, and during phase 2 the upper bound has to be updated from time to time. The cost of these operations, added to the phase-switch checks, is probably not negligible.

Comparing the time-dependent variant of SHARC with our approach, we observe that SHARC with an approximation value of 1.001 settles as many nodes as ALT with $K = 2.0$. However, query performance is better for SHARC due to its small computational overhead. By increasing the approximation value, computational times are slowed by almost one order of magnitude, but the solution quality merely improves. The reason for this poor performance is that SHARC uses a contraction routine which cannot bypass nodes incident to time-dependent edges. As in our scenario about half of the edges are time-dependent, the preprocessing of SHARC takes quite long (≈ 12 hours) and query performance is poor. Summarizing, ALT seems to work much better in a time-dependent scenario.

Local Queries. For random queries, our bidirectional ALT algorithm (with $K = 1.15$) is roughly 6.7 times faster than unidirectional ALT on average. In order to gain insight whether this speed-up derives from small or large distance queries, Fig. 1 reports the query times with respect to the Dijkstra rank⁴. These values

⁴ For an s - t query, the Dijkstra rank of node t is the number of nodes settled before t is settled. Thus, it is some kind of distance measure.

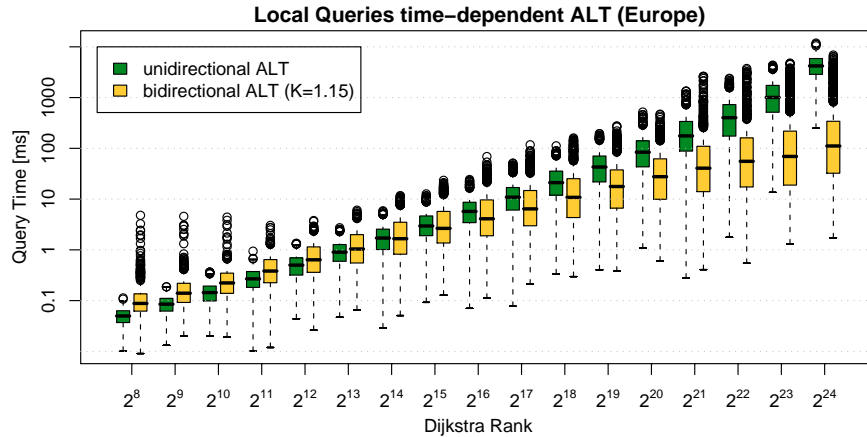


Fig. 1. Comparison of uni- and bidirectional ALT using the Dijkstra rank methodology [20]. The results are represented as box-and-whisker plot: each box spreads from the lower to the upper quartile and contains the median, the whiskers extend to the minimum and maximum value omitting outliers, which are plotted individually.

were gathered on the European road network instance. Note that we use a logarithmic scale due to the fluctuating query times of bidirectional ALT. Comparing both ALT version, we observe that switching from uni- to bidirectional queries pays off especially for long-distance queries. This is not surprising, because for small distances the overhead for bidirectional routing is not counterbalanced by a significant decrease in the number of explored nodes: unidirectional ALT is faster for local queries. For ranks of 2^{24} , the median of the bidirectional variant is almost 2 orders of magnitude lower than for the unidirectional variant. Another interesting observation is the fact that some outliers of bidirectional ALT are almost as slow as the unidirectional variant.

Number of Landmarks. In static scenarios, query times of bidirectional ALT can be significantly reduced by increasing the number of landmarks to 32 or even 64 (see [4]). In order to check whether this also holds for our time-dependent variant, we recorded our algorithm’s performance using different numbers of landmarks. Tab. 2 reports those results on the European road network. We evaluate 8 maxcover landmarks (yielding a preprocessing effort of 33 minutes and an overhead of 64 bytes per node), 16 maxcover landmarks (75 minutes, 128 bytes per node) and 32 avoid landmarks (29 minutes, 256 bytes per node). Note that we do not report error rates here, as it turned out that the number of landmarks has almost no impact on the quality of the computed paths. Surprisingly, the number of landmarks has a very small influence on the performance of time-dependent ALT. Even worse, increasing the number of landmarks even yields larger average query times for unidirectional ALT and for bidirectional ALT with low K -values. This is due the fact that the search space decreases only

Table 2. Performance of uni- and bidirectional ALT with different number of landmarks in a time-dependent scenario.

	K	8 landmarks		16 landmarks		32 landmarks	
		# settled	time [ms]	# settled	time [ms]	# settled	time [ms]
uni-ALT	-	2 321 760	1 739.8	2 192 010	1 775.8	2 111 090	1 868.5
ALT	1.00	3 240 210	3 270.6	3 117 160	3 399.3	3 043 490	3 465.1
	1.10	863 526	736.5	719 769	665.1	681 836	669.7
	1.15	495 649	382.1	348 325	287.3	312 695	280.0
	1.20	389 096	286.3	241 241	185.3	204 877	170.1
	1.50	320 026	228.4	172 157	121.9	133 547	98.3
	2.00	313 448	222.2	165 650	115.7	126 847	91.1

slightly, but the additional overhead for accessing landmarks increases when there are more landmarks to take into account. However, when increasing K , a larger number of landmarks yields faster query times: with $K = 2.0$ and 32 landmarks we are able to perform time-dependent queries 70 times faster than plain Dijkstra, but the solution quality in this case is as poor as in the 16 landmarks case. Summarizing, for $K > 1.10$ increasing the number of landmarks has a positive effect on computational times, although switching from 16 to 32 landmarks does not yield the same benefits as from 8 to 16, and thus in our experiments is not worth the extra memory. On the other hand, for $K \leq 1.10$ and for unidirectional ALT increasing the number of landmarks has a negative effect on computational times, and thus is never a good choice in our experiments.

7 Conclusion and Future Work

We have presented an algorithm which applies bidirectional search on a time-dependent road network, where the backward search is used to bound the set of nodes that have to be explored by the forward search; this algorithm is based on the ALT variant of the A^* algorithm. We have discussed related theoretical issues, and we proved the algorithm’s correctness. Extensive computational experiments show that this algorithm is very effective in practice if we are willing to accept a small approximation factor: the exact version of our algorithm is slower than unidirectional ALT, but if we can accept a decrease of the solution quality of a few percentage points with respect to the optimum then our algorithm is several times faster. For practical applications, this is usually a good compromise. We have compared our algorithm to existing methods, showing that this approach for bidirectional search is able to significantly decrease computational times.

Future research will include the possibility of an initial contraction phase for a time-dependent graph, which would be useful for several purposes, and algorithm engineering issues such as the balancing of the forward and backward search, and the update of the available upper bound on the optimal solution cost. The idea of bidirectional routing on time-dependent graphs, using a time-dependent forward

search and a time-independent backward search, may be applied to other static routing algorithms, in order to generalize them in a time-dependent scenario.

Acknowledgements

We would like to acknowledge Mediamobile’s financial support, and thank Daniele Pretolani for providing useful discussions.

References

1. Wagner, D., Willhalm, T.: Speed-up techniques for shortest-path computations. In Thomas, W., Weil, P., eds.: 24th Annual Symposium on Theoretical Aspects of Computer Science (STACS 2007). Volume 4393 of LNCS., New York, Springer (2007) 23–36
2. Sanders, P., Schultes, D.: Engineering fast route planning algorithms. In Demetrescu, C., ed.: Proceedings of the 6th Workshop on Experimental Algorithms (WEA 2007). Volume 4525 of LNCS., New York, Springer (2007) 23–36
3. Bast, H., Funke, S., Sanders, P., Schultes, D.: Fast routing in road networks with transit nodes. *Science* **316**(5824) (2007) 566
4. Delling, D., Wagner, D.: Landmark-based routing in dynamic graphs. In Demetrescu, C., ed.: Proceedings of the 6th Workshop on Experimental Algorithms (WEA 2007). Volume 4525 of LNCS., New York, Springer (2007)
5. Sanders, P., Schultes, D.: Dynamic highway-node routing. In Demetrescu, C., ed.: Proceedings of the 6th Workshop on Experimental Algorithms (WEA 2007). Volume 4525 of LNCS., New York, Springer (2007) 66–79
6. Wagner, D., Willhalm, T., Zaroliagis, C.: Geometric containers for efficient shortest-path computation. *ACM Journal of Experimental Algorithmics* **10** (2005) 1–30
7. Cooke, K., Halsey, E.: The shortest route through a network with time-dependent internodal transit times. *Journal of Mathematical Analysis and Applications* **14** (1966) 493–498
8. Dreyfus, S.: An appraisal of some shortest-path algorithms. *Operations Research* **17**(3) (1969) 395–412
9. Dijkstra, E.: A note on two problems in connexion with graphs. *Numerische Mathematik* **1** (1959) 269–271
10. Kaufman, D.E., Smith, R.L.: Fastest paths in time-dependent networks for intelligent vehicle-highway systems application. *Journal of Intelligent Transportation Systems* **1**(1) (1993) 1–11
11. Orda, A., Rom, R.: Shortest-path and minimum delay algorithms in networks with time-dependent edge-length. *Journal of the ACM* **37**(3) (1990) 607–625
12. Daganzo, C.: Reversibility of time-dependent shortest path problem. Technical report, Institute of Transportation Studies, University of California, Berkeley (1998)
13. Hart, E., Nilsson, N., Raphael, B.: A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems, Science and Cybernetics* **SSC-4**(2) (1968) 100–107
14. Goldberg, A., Harrelson, C.: Computing the shortest path: A^* meets graph theory. In: Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2005), SIAM (2005)

15. Goldberg, A., Kaplan, H., Werneck, R.: Reach for A^* : Efficient point-to-point shortest path algorithms. In Demetrescu, C., Sedgwick, R., Tamassia, R., eds.: Proceedings of the 7th Workshop on Algorithm Engineering and Experimentation (ALENEX 05), SIAM (2005)
16. Chabini, I., Shan, L.: Adaptations of the A^* algorithm for the computation of fastest paths in deterministic discrete-time dynamic networks. IEEE Transactions on Intelligent Transportation Systems **3**(1) (2002) 60–74
17. Bauer, R., Delling, D.: SHARC: Fast and Robust Unidirectional Routing. In: Proceedings of the 10th Workshop on Algorithm Engineering and Experiments (ALENEX 08), SIAM (2008) to appear.
18. Ikeda, T., Tsu, M., Imai, H., Nishimura, S., Shimoura, H., Hashimoto, T., Tenmoku, K., Mitoh, K.: A fast algorithm for finding better routes by ai search techniques. In: Proceedings for the IEEE Vehicle Navigation and Information Systems Conference. (2004) 291–296
19. Goldberg, A., Werneck, R.: An efficient external memory shortest path algorithm. In Demetrescu, C., Sedgwick, R., Tamassia, R., eds.: Proceedings of the 7th Workshop on Algorithm Engineering and Experimentation (ALENEX 05), SIAM (2005) 26–40
20. Sanders, P., Schultes, D.: Highway hierarchies hasten exact shortest path queries. In Stølting Brodal, G., Leonardi, S., eds.: 13th Annual European Symposium on Algorithms (ESA 2005). Volume 3669 of Lecture Notes in Computer Science., Springer (2005) 568–579
21. Kerner, B.S.: The Physics of Traffic. Springer, Berlin (2004)