Scenic Routes Now: Efficiently Solving the Time-Dependent Arc Orienteering Problem

Gregor Jossé*
LMU Munich
Oettingenstraße 67, 80583
Munich, Germany
josse@dbs.ifi.lmu.de

Ying Lu*
University of Southern
California
3737 Watt Way, Los Angeles,
CA 90089-0272, USA
ylu720@usc.edu

Tobias Emrich
LMU Munich
Oettingenstraße 67, 80583
Munich, Germany
emrich@dbs.ifi.lmu.de

Matthias Renz
George Mason University
4400 University Dr, Fairfax, VA
22030, USA
mrenz@gmu.edu

Cyrus Shahabi
University of Southern
California
3737 Watt Way, Los Angeles,
CA 90089-0272, USA
shahabi@usc.edu

Ugur Demiryurek
University of Southern
California
3737 Watt Way, Los Angeles,
CA 90089-0272, USA
demiryur@usc.edu

This paper extends the Arc Orienteering Problem (AOP) to large road networks with time-dependent travel times and time-dependent value gain, termed Twofold Time-Dependent AOP or 2TD-AOP for short. In its original definition, the NP-hard Orienteering Problem (OP) asks to find a path from a source to a destination maximizing the accumulated value while not exceeding a cost budget. Variations of the OP and AOP have many practical applications such as mobile crowdsourcing tasks (e.g., repairing and maintenance or dispatching field workers), diverse logistics problems (e.g., crowd control or controlling wildfires) as well as several tourist guidance problems (e.g., generating trip recommendations or navigating through theme parks). In the proposed 2TD-AOP, travel times and value functions are assumed to be time-dependent. The dynamic values model, for instance, varying rewards in crowdsourcing tasks or varying urgency levels in damage control tasks. We discuss this novel problem, prove the benefit of time-dependence empirically and present an efficient approximative solution, optimized for fast response systems. Our approach is the first time-dependent variant of the AOP to be evaluated on a large scale, fine-grained, real-world road network. We show that optimal solutions are infeasible and solutions to the static problem are often invalid. We propose an approximate dynamic programming solution which produces valid paths and is orders of magnitude faster than any optimal solution.

1. INTRODUCTION

Not long ago, finding the shortest or fastest path in road networks was the core challenge of route planning. Nowadays, challenges are manifold – due to the abundance of sensor data, due to the existence of increasingly complex traffic models and due to the demand for efficient yet convenient solutions. Modern navigation systems often take multiple, sometimes time-dependent, cost criteria into account and solve complex queries in order to increase driver convenience and traffic efficiency. In view of this progress, we propose a novel query with multiple applications. As an extension to the family of Orienteering Problems, we present the Twofold Time-Dependent Arc Orienteering Problem (2TD-AOP).

In its original formulation [22], the Orienteering Problem (OP) requires to find a path from a given source to a given destination abiding by a given cost budget but maximizing the value collected along the way. In most cases, the cost function corresponds to the travel time within the given network. The problem may for instance occur in the field of spatial crowdsourcing [32, 7, 44] where individuals complete advertised tasks in order to collect rewards. A recent popular example of this is the augmented reality game Pokémon Go. A worker who is willing to complete crowdsourced tasks will want to maximize the collected rewards within his chosen time frame. The OP combines the NP-hard Knapsack and Traveling Salesman Problems, hence is NP-hard itself. Numerous variations of the problem have been proposed, for instance, extending it to multiple workers (Team Orienteering Problem) [24] or restricting the collectable values to certain time windows ((Team) Orienteering Problem with Time Windows) [18]. Another variation is the Arc Orienteering Problem (AOP) [2]. In the OP, the value is associated with the vertices of the graph, whereas in the AOP, it is associated with the arcs. Thus, the OP intuitively corresponds to logistic or crowdsourcing tasks, where the value is collected at particular points. The AOP, on the other hand, corresponds to tasks where the value is collected "along the way". Examples of such tasks are the routes of firefighting planes or the planning of scenic bike trips [42].

Until recently, the AOP has only been answered for networks with static travel times. However, assuming static costs is not reasonable for most real-world scenarios. For road networks, the significance of time-dependence was empirically proven in [13] and substantiated by the incorporation into many navigation services such as Waze[1]. In a recent study, Verbeeck et. al [47] first in-

---

*These authors contributed equally to this work.
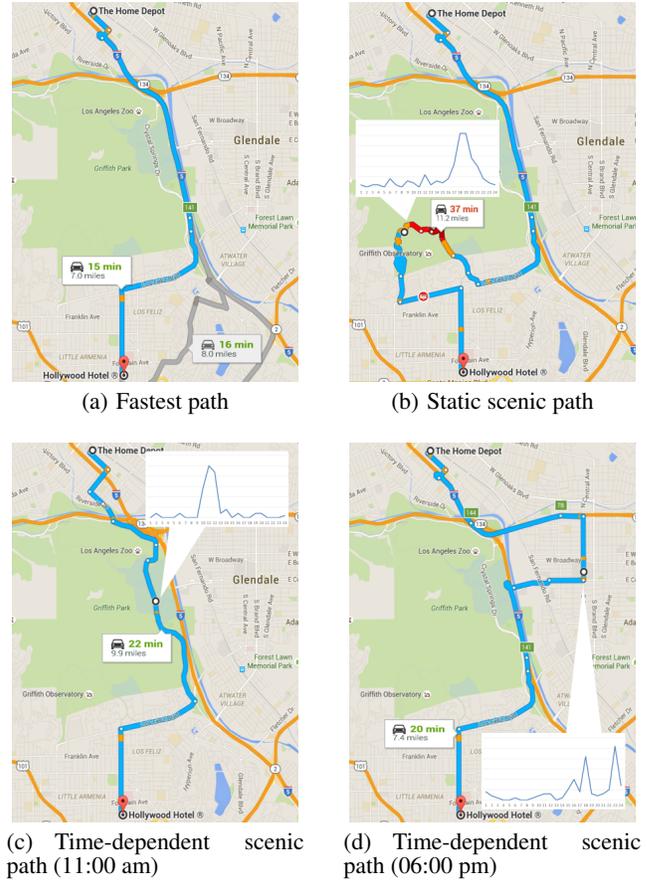
[1] waze.com

troduced time-dependent travel times but the values have been assumed to be static. In this paper, we propose to incorporate time-dependent values in addition to time-dependent travel times, introducing the 2TD-AOP. As will be shown, results benefit greatly from time-dependent values which model the varying benefit for different times. Without loss of generality, we restrict ourselves to the application of scenic route planning. In this use case, a view of the sea might only be worthwhile during day, whereas a view of the skyline might be particularly worthwhile at night. Thus, the value of a coastline road is high during the day but low at night, while a road passing through downtown may have the opposite values. Other use cases include the aforementioned firefighting planes, collecting street view data or gathering Pokémon Go items.

Consider the example given in Figure 1 which shows different result paths for a scenic path query with the same source and destination. Figure 1(a) shows the time-dependent fastest path according to our algorithm for two departure times. While the path is the same at 11 am and 6 pm, the travel times are not the same (15 and 20 minutes, respectively). In this example, the fastest path does not pass through particularly scenic areas, according to the crowd-sourced data set. The other paths pass through such areas, and for particular scenic parts of the respective paths, the hourly value distributions are highlighted. At first glance it is evident that the value functions show great variance over the course of a day. If not allowing for time-dependent values, these functions are commonly averaged, evening out the significant peaks of the distribution.

Suppose a driver is willing to spend 25 minutes for his trip from Hollywood Hotel to the Home Depot north of Griffith Park. Employing an algorithm for the AOP in static networks[38], we obtain the *static scenic path* displayed in Figure 1(b). The path was computed in a static network derived from our time-dependent network. The static travel time along an arc is set as the average of all travel times along this arc in the time-dependent network. The static sceniceness values are derived analogously. The generated result path has a static travel time of 22 minutes which abides by the budget. However, the actual travel time might of course deviate in either direction. For a desired departure at 11 am, the time-dependent travel time coincides with the static time of 22 minutes, but for a desired departure at 6 pm, the time-dependent path takes 37 minutes. Thus, at 6 pm the result path clearly exceeds the budget and is therefore invalid. In a similar way, the averaged scenicness values can deviate from the time-dependent ones. While the static scenic path at 11 am does not exceed the budget, it is not optimal either. Griffith Observatory, which it passes along the way, is a popular sight but particularly scenic during sunset or in the evening. Instead, the best solution at 11 am is displayed in Figure 1(c). This *time-dependent scenic path* takes 20 minutes and passes a golf course, offering particularly nice views at daytime. The best solution at 6 pm is shown in Figure 1(d). This path takes 22 minutes and passes through central Glendale, a bustling area in the evening and at night.

Taking time-dependence into account when computing AOP paths is crucial. Not only result quality but also result validity depends on it. Our experiments show that every second static solution in a moderately complex settings is invalid. Users benefit considerably from time-dependent travel times as every 2TD-AOP path is valid. Furthermore, results are improved significantly when the values reflect time-dependence as well. Our experiments show that 2TD-AOP paths generate three to four times more value than static solutions. We therefore propose to employ value distribution models whenever possible.

Due to their NP-hardness, static variations of the OP and AOP are usually solved employing metaheuristics. Allowing for time-

(a) Fastest path

(b) Static scenic path

(c) Time-dependent scenic path (11:00 am)

(d) Time-dependent scenic path (06:00 pm)

Figure 1: Exemplary paths with hourly value distribution for particular areas along the path.

dependent travel times adds further complexity to the task. Slight modifications of valid solutions may invalidate them. For instance, when taking a minor detour at the beginning of the path, travel times for the rest of the path may change and exceed the budget. Additionally incorporating time-dependent values multiplies this effect. Taking a minor detour at the beginning of a path may result in a change of value and travel time. We propose to solve the 2TD-AOP with a novel dynamic programming approach which makes use of pruning techniques to limit the candidate set and reduce intermediate path computations.

To the best of our knowledge, no previous work has allowed for both travel times and values to be time-dependent. In addition to handling the increased complexity, we require fast response times. We focus on the practicality of the solution where fast response times are crucial and large road network are standard, for instance, in order to meet the requirements of a mobile application. In conclusion, the contributions of this paper are as follows:

1. We motivate and introduce the 2TD-AOP, as a new graph-based routing problem and a Mixed Integer Programming formulation.

2. We present an efficient approximate solution to the 2TD-AOP.

3. We present experiments on a large-scale real-world road network. Our evaluation is the first of this dimension for any

time-dependent AOP. We show the value of twofold time-dependence empirically and compare 2TD-AOP solutions to static solutions as well as to optimal results.

The remainder of this paper is organized as follows. In Section 2 we give an overview of research relevant to our work. In Section 3 the 2TD-AOP is formally defined and formulated as an Mixed Integer Programming. In Section 4 we lay the theoretical foundation for our algorithm which is described in Sections 5 and 6. Subsequently, our algorithm is evaluated in Section 7. A conclusion is provided in Section 8.

2. RELATED WORK

We divide research relevant to this work into three groups. First, we review research on how to attain profit values from data in general and crowdsourced content in particular. Next, we briefly introduce related research from the database community such as the Trip Planning Query (TPQ). These problems have a similar field of application but are not congruent to the problem discussed in this paper. Finally, we present research on the family of OPs which mostly stems from the field of Operations Research. We give an overview of solutions to the static and time-dependent OP and AOP.

2.1 Attaining Value Functions

An important aspect of the OP/AOP is how to attain the value functions, in particular when aiming to solve a time-dependent instance. Most works from the field of Operations Research rely on specific bench mark data sets, some synthetic, others real. The authors of [42], for instance, evaluate their approach to cycle trip planning on real-world tourist data from Flanders, Belgium. The authors of [25] use data from Chinese theme parks. Relying on specifically recorded data like the above has two major drawbacks. First, expert knowledge or an official mandate may be required to attain the data. Second, the data may not be available globally. As this approach usually lacks scalability, there exist advances to mine this knowledge from freely available data.

Crowdsourced data is a rich source of information. It is for instance used for general POI recommendation [27], for categorization [5], for popularity estimation [3], for determining opening hours or visiting times [26], for determining duration of stay [9] or for determining spatial connectivity [41]. The most common data sources for these purposes are Foursquare, Yelp, Twitter, Flickr[2] or the inoperative but oft-cited Gowalla. When interested in how to attain value functions from the raw data, we refer the reader to the respective works. In the context of this paper, we assume the values to be given. For our experimental evaluation, we rely on Flickr data, details are given in Section 7.

2.2 TPQ and Related Queries

The most related query in the database community is the TPQ [8] (alternatively, Route Planning Query [6] or Route Search Query [35]). In this problem setting, the user specifies a different POI categories, e.g., "ATM", "restaurant", "florist", "cinema". The result of a TPQ is the fastest path from a given source to a given destination visiting exactly one instance of each resource. The TPQ is NP-hard because in the case that each category occurs exactly once, the TPQ degenerates to the Traveling Salesman Problem. There exist several variations, for instance introducing order constraints [33, 30] or modeling the fulfillment at a location probabilistically or time-dependent [31, 29, 11]. Solutions to this problem include full enumeration [11], heuristics [6, 33, 11], or backtracking and

---

branch-and-bound approaches [29]. Partially time-dependent variations of these queries exist, however, most of the research pertains to static networks. Most importantly, albeit related, the problem settings of the OP/AOP and the TPQ are not congruent.

2.3 OP and AOP

Blueprint for the OP was the family of orienteering sports where efficient navigation from checkpoint to checkpoint is the goal [22]. Alternative names for the OP include the Selective Traveling Salesman Problem [34] or the Traveling Salesman Problem with Profits [14]. The OP is executed on a graph whose arc weights represent traversal costs and whose vertices may yield certain value. Given a source, a destination and a cost budget, the output of an OP is the path with maximum value among all paths not exceeding the budget. In the variation referred to as AOP or Privatized Rural Postman Problem [1], the profit is not assigned to vertices but to arcs, while the rest is as before. Extensive surveys on these problems and their numerous variants are given in [46, 24, 2] (or from the perspective of the related Tourist Trip Design Problem in [18]).

Although NP-hard [22], exact solutions to the OP exist [34, 40, 14]. Most solutions, however, are approximative. While in older research specific heuristic approaches were proposed [45, 23, 4], the trend shifted towards algorithms following particular meta-heuristics such as Tabu Search [21], Genetic Algorithms [43] or Ant Colony Algorithms [37]. This holds similarly for the AOP where exact algorithms [10, 48] are rare but available. Generating approximative solutions is the standard approach [20], most effectively following the meta-heuristics Iterative Local Search (ILS) [48] and Greedy Randomized Adaptive Search Procedure (GRASP) [42, 38]. Recently, ILS and GRASP have been improved to solve the AOP on large scale real-world road networks in near-interactive time [38].

Introduced in [15], the OP with time-dependent (TD-OP) costs has gained attention recently [25, 17, 19, 47, 36]. Typical applications for the TD-OP include electronic tourist guides [17, 16] or routing in theme parks [25]. The time-dependent AOP (TD-AOP) as presented in [47] is not directly linked to an application but a time-dependent extension of the authors' cycle trip planning application [48] is imaginable.

The above works are all evaluated on small graphs (with the exception of [38]) with no more than a couple of hundred vertices. Most arcs in these graphs connect two non-zero value vertices. For instance, in the theme park data set of [25], vertices correspond to attractions and the arcs are fastest paths between them. The travel times and the values are derived from user data of the theme parks and their attractions. Similarly, [19] and the demonstration paper [16] are evaluated on data sets consisting of Points of Interest (POIs) and fastest paths connecting these POIs (by walking or by public transportation). When considering real-world road networks, this is not feasible. A moderate-sized network usually contains over one million vertices and two to three million arcs. In a time-dependent network, there usually exist numerous fastest paths from a given vertex to another, depending on the departure time. Hence, introducing shortcuts for precomputed fastest paths as in the approaches above is not possible. This is particularly the case, if the travel times are not statistically inferred but real-time, i.e., depending on live information, as is the case in the network of our evaluation.

The most closely related problem to the 2TD-AOP is the TD-AOP with time-dependent costs and static value. Verbeeck et. al recently proposed to solve this problem with an Ant Colony System algorithm [47]. Evaluated on specific benchmark instances with at most 100 vertices, the proposed approach finds near-optimal to optimal solutions within an average runtime of one second. The com-

plex solver has excessive time requirements ("often more than 100 hours" [47]) to generate optimal solutions on these benchmark instances of the TD-AOP. In comparison, we increase complexity of the problem in two ways. First, by including time-dependent value functions 2TD-AOP adds another dimension of time-dependence to the problem. Second, we aim to solve the problem on large-scale real-world road networks with hundred thousands of vertices. Established solutions to related problems like the TD-AOP or the TD-OP rely on metaheuristics such as the Ant Colony System [47], Iterative Local Search [48] or Greedy Randomized Adaptive Search Procedure [42, 38]). Instead, we propose a heuristic dynamic programming approach.

3. PROBLEM DEFINITION

We model the road network as a graph $G = (V, A, val, tt)$, where V denotes the set of vertices (or nodes), $A \subseteq V \times V$ denotes the set of directed arcs (or edges), $val : V \times V \times \mathbb{R}_0^+ \to \mathbb{R}_0^+$ and tt (with the same domains) denote the non-negative functions mapping each arc onto its respective value and the travel time, respectively

$$val_{i,j}(t) := val(v_i, v_j, t) \geq 0, \quad tt_{i,j}(t) := tt(v_i, v_j, t) \geq 0$$

For the representation of the time-dependent travel time along an arc, we use a piecewise constant function with equal step width τ. In the following, we refer to the intervals where travel time is constant as time windows and denote the k-th time window by τ_k, where $0 \leq k \leq K$ for some maximum time window K.

A consecutive set of arcs which visits no arc twice is called a path. A path from source to destination departing at t_0 is denoted by $p = ((v_{p_0}, v_{p_1}), (v_{p_1}, v_{p_2}), \ldots, (v_{p_{N-1}}, v_{p_N}), t_0)$ where v_{p_0} is the source vertex and v_{p_N} the destination vertex. In the remainder of this paper, we simplify notation for the sake of clarity. We omit the double indexes and denote a path p by

$$p = ((v_0, v_1), (v_1, v_2), \ldots, (v_{N-1}, v_N), t_0)$$

For each query, we denote source and destination by v_0 and v_N, respectively. The travel time of a path p is defined as $tt(p) = \sum_{i=0,\ldots,N-1} tt_{i,i+1}(t_i)$, where t_i denotes the time of arrival at (and departure from) vertex v_i. t_i is dependent on the travel times along the preceding arcs and is defined iteratively:

$$t_i := \sum_{j=0}^{i-1} tt_{j,j+1}(t_j)$$

The value of a path p is defined as:

$$val(p) = \sum_{i=0,\ldots,N-1} val_{i,i+1}(t_i)$$

For a given source v_0, destination v_N, and departure time t_0, we denote the set of all paths from v_0 to v_N starting at t_0 by $\mathcal{P}_{0,N}(t_0)$.

The input for the 2TD-AOP is a graph G, a source vertex v_0, a destination vertex v_N, a departure time t_0, and a time budget b. When speaking of a query, we mean the set of these inputs. The optimal path to a 2TD-AOP query is the path maximizing the collected value among all feasible paths. Both notions are defined in the following.

DEFINITION 1 (FEASIBLE PATH). *Given a query, we call any path p departing from v_0 at t_0 and arriving at v_N no later than $t_0 + b$ (i.e., $tt(p) \leq b$) a* feasible *path.*

DEFINITION 2 (OPTIMAL PATH). *Given a query, among all feasible paths, we call one with maximum value an* optimal *path.*

Extending the work in [46, 47], we give a mixed integer programming (MIP) formulation of the problem. We introduce the following two decision variables: $p_{i,j,k}$ is a binary decision variable which equals 1 if the result path enters the arc (v_i, v_j) in time window τ_k and 0 if the arc is not traversed at all or entered in a different time window. $t_{i,j,k}$ is a continuous decision variable containing the time of departure at v_j when entering the arc (v_i, v_j) in time window τ_k. Again, it is 0 if the arc is not traversed at all or entered in a different time window.

$$\max \sum_{i \neq j} \sum_{k=0}^{K} p_{i,j,k} \cdot val_{i,j}(\tau_k) \quad (1a)$$

subject to:
$$\sum_{i \neq j} \sum_{k=0}^{K} p_{i,j,k} \cdot tt_{i,j}(\tau_k) \leq b \quad (1b)$$

$$\forall 1 \leq l \leq N - 1:$$

$$\sum_{i=0}^{N-1} \sum_{k=0}^{K} p_{i,l,k} = \sum_{j=1}^{N} \sum_{k=0}^{K} p_{l,j,k} \leq 1 \quad (1c)$$

$$\sum_{j=1}^{N} p_{0,j,0} = \sum_{i=0}^{N} \sum_{k=0}^{K} p_{i,N,k} = 1 \quad (1d)$$

$$\forall 1 \leq l \leq N - 1:$$

$$\sum_{i=0}^{N-1} \sum_{k=0}^{K} t_{i,l,k} + tt_{i,l}(\tau_k) = \sum_{j=1}^{N} \sum_{k=0}^{K} t_{l,j,k} \quad (1e)$$

$$p_{i,j,k} \in \{0, 1\}, \ \forall i, j, k \quad (1f)$$

$$0 \leq t_{i,j,k} \leq \tau_K, \ \forall i, j, k \quad (1g)$$

$$t_{0,j,0} = t_0, \ \forall j \quad (1h)$$

Equation 1a gives the objective function while Equation 1b gives the budget constraint. Note that due to the time-dependent nature of the problem, all time windows have to be taken into account when summing the collected value and cost. Remember, however, that $p_{i,j,k}$ is a binary decision variable which for each arc can at most once be equal to 1, as each vertex (and thereby also each arc) is at most visited once. This is expressed in Constraint 1c. Constraint 1d ensures that the result path starts at the source and ends at the destination. Next, Constraint 1e ensures that the departure time at each arc is the departure time at its predecessor plus the travel time within the corresponding time window. Note that this condition implicitly forbids "waiting" at vertices. Finally, Constraints 1f and 1g determine the domain of the decision variables and Constraint 1h ensures that the path departs from the source at t_0.
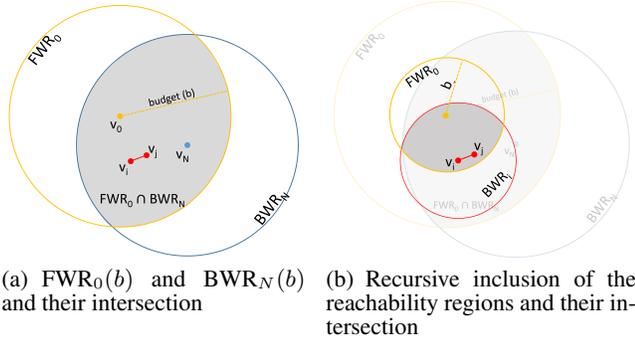
4. SOLVING THE 2TD-AOP

In this section, we present terminology and insights essential to our solution of the 2TD-AOP.

DEFINITION 3 (EARLIEST ARRIVAL). *Given a graph G, a source v_0, a vertex v_i, and a departure time t_0, we define the* earliest arrival (time) *for reaching v_i from v_0 when departing at t_0, denoted by $ea_{0,i}$ (ea_i for short), as follows:*

$$ea_i := t_0 + \min_{p \in \mathcal{P}_{0,i}(t_0)} tt(p)$$

DEFINITION 4 (LATEST DEPARTURE). *Given a graph G, a destination v_N, a vertex v_i, and a latest arrival time $t_0 + b$, we*

(a) $FWR_0(b)$ and $BWR_N(b)$ and their intersection

(b) Recursive inclusion of the reachability regions and their intersection

Figure 2: Reachability regions as computed by FWR and BWR.

define the latest departure (time) for reaching v_N from v_0 no later than $t_0 + b$, denoted by $ld_{i,N}(t_0 + b)$ (ld_i for short), as follows:

$$ld_i := t_0 + b - \min_{p \in \mathcal{P}_{i,N}(-\infty)} tt(p)$$

i.e., the latest departure time among all feasible paths w.r.t. v_i, v_N, and t.

DEFINITION 5 (FORWARD-REACHABLE). *Given a graph G, a source v_0, a departure time t_0, and a time budget b, we refer to the vertices reachable within the budget b from v_0 (when departing at t_0) as* forward-reachable *vertices, denoted by $FWR_0(b)$ (FWR for short if v_0 and b are clear from context).*
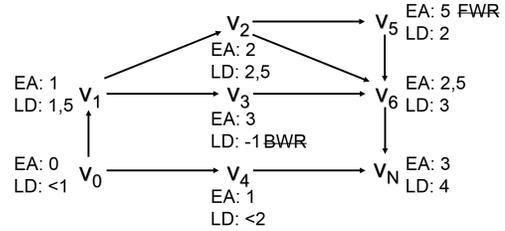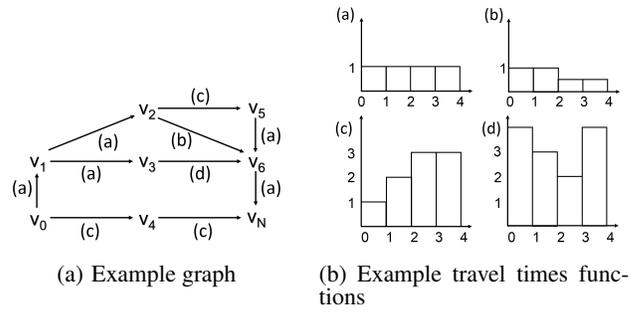
$$FWR_0(b) := \{v_i \in V \mid ea_i \leq t_0 + b\}$$

DEFINITION 6 (BACKWARD-REACHABLE). *Given a graph G, a destination v_N, a departure time t_0, and a time budget b, we refer to the vertices from which the destination can be reached within the budget b when departing no earlier than t_0 as* backward-reachable *vertices, denoted by $BWR_N(b)$ (BWR for short).*

$$BWR_N(b) := \{v_i \in V \mid ld_i \geq t_0\}$$

Intuitively, all forward-reachable vertices lie in a quasi-circular region around the source. This region is similar to the circular expansion of a Dijkstra-search. The difference is that in this case the cost criterion is time-dependent travel time. Analogously, all backward-reachable vertices lie in a quasi-circular region around the destination. This graphical intuition is visualized in Figure 2(a). If the destination is not forward-reachable from the source (and/or vice versa), then the query has no answer.

Our solution to the 2TD-AOP is based on an extension of this observation: For an arc (v_m, v_n) to be part of a feasible path, it is a necessary condition that v_m is forward-reachable from the source v_0 and v_n is backward-reachable from the destination v_N. Hence, if an arc is not contained in the intersection of the forward-reachability region and the backward-reachability region, it cannot be part of a feasible path. This property will be proved later on. In our solution to the 2TD-AOP, arcs which increase the value are recursively inserted until the budget is exhausted. With every insertion of an arc, the overall travel time of the path increases. Hence, the reachability regions iteratively contract until no more arcs can be inserted. Before we detail this approach in the next section, let us explain the above definitions and their computation with an example.

The definitions are visualized in Figure 3. Figure 3(a) shows a graph with source v_0 and destination v_N. Departure time is $t_0 = 0$,

(a) Example graph

(b) Example travel times functions

(c) Earliest arrivals and latest departures

Figure 3: Example graph for a query with source v_0, destination v_N, departure time 0 and time budget $b = 4$.

and let the time budget $b = 4$. The time-dependent travel times along the arcs are visualized in Figure 3(b). The arcs in Figure 3(a) are labeled with letters which correspond to the travel time functions in Figure 3(b). Figure 3(c) shows the earliest arrivals and latest departures according to this query. If not indicated otherwise, vertices are forward-reachable as well as backward-reachable. The earliest arrivals are computed as in the definition. For example, $ea_1 = 0 + tt((v_0, v_1), 0) = 1$, implying that the fastest path from v_0 at $t_0 = 0$ arrives at v_1 at time 1.

Consider the following path:

$$((v_0, v_1), (v_1, v_2), (v_2, v_5), 0) = 1 + 1 + 3 = 5 > b$$

Since there is no shorter path from v_0 to v_5, v_5 is not FWR within the given budget. From v_0 there are two feasible paths to v_N:

$$((v_0, v_4), (v_4, v_N), 0) = 1 + 2 = 3$$

$$((v_0, v_1), (v_1, v_2), (v_2, v_6), (v_6, v_N), 0) = 1 + 1 + 0.5 + 1 = 3.5$$

Which of the two constitutes the optimal path depends on the time-dependent value functions and will be discussed later on.

Let us turn to the computation of the latest departures. For example, $ld_6 = 0 + 4 - 1 = 3$, implying that in order to arrive at v_N no later than $t_0 + b$, one must depart at time 3. Now, consider ld_4: the only path to the destination follows the arc (v_4, v_N). Departing at time 3 results in exceeding the budget, since $tt_{4,N}(3) = 3$. The same holds for time 2. However, when departing at $t < 2$, then $tt_{4,N}(t) \leq 2$, resulting in an arrival within the budget b. Hence, in order to compute the latest departure time of a vertex v_i a linear scan among the time windows might be required. Note that the difference between "at time 2" and "before time 2" is infinitesimal. In theory we set the latest departure at v_4 to $2 - \varepsilon$, in practice, we adjust it by a sufficiently small number. Let us note that this special case only occurs if the latest departure at a vertex by chance coincides with the border between two time windows τ_k which is

Algorithm 1 FWR

```
1:  procedure FWR(G, v_0, t_0, t_0 + b)
2:      Q ← empty queue of vertices v_i sorted asc. by ea_i
3:      result = ∅
4:      ea_0 = t_0
5:      Q.add(v_0)
6:      while Q is not empty do
7:          v_i ← Q.removeFirst
8:          for (v_i, v_j) ∈ G.A do
9:              êa_j ← ea_i + tt_{i,j}(ea_i)
10:             if êa_j < t_0 + b then
11:                 if ¬ result.contains(v_j) then
12:                     ea_j ← êa_j
13:                     Q.add(v_j)
14:                     result.add(v_j)
15:                 else
16:                     if êa_j < ea_j then
17:                         ea_j ← êa_j
18:                     end if
19:                 end if
20:             end if
21:         end for
22:     end while
23:     return result
24: end procedure
```

Algorithm 2 BWR

```
1:  procedure BWR(G, v_N, t_0, t_0 + b)
2:      Q ← empty queue of vertices v_i sorted desc. by ld_i
3:      result = ∅
4:      ld_N = t_0 + b
5:      Q.add(v_N)
6:      while Q is not empty do
7:          v_j ← Q.removeFirst
8:          for (v_i, v_j) ∈ G.A do
9:              τ_k ← arg max_{τ_l}{τ_l < ld_j}
10:             while τ_k + tt_{i,j}(τ_k) > ld_j do
11:                 k ← k − 1
12:             end while
13:             l̂d_i ← ld_j − tt_{i,j}(τ_k) [−ε]
14:             if l̂d_i ≥ t_0 then
15:                 if ¬ result.contains(v_i) then
16:                     ld_i ← l̂d_i
17:                     Q.add(v_i)
18:                     result.add(v_i)
19:                 else
20:                     if l̂d_i > ld_i then
21:                         ld_i ← l̂d_i
22:                     end if
23:                 end if
24:             end if
25:         end for
26:     end while
27:     return result
28: end procedure
```

very unlikely in practice. Finally, we prove that in order to compute all feasible paths it suffices to consider the vertices which are forward-reachable and backward-reachable.

LEMMA 1. *All vertices along all feasible paths are contained in the intersection of the sets of forward-reachable and backward-reachable vertices.*

PROOF. *Suppose there exists a feasible path p visiting a vertex $v_i \notin FWR \cap BWR$. Case 1: $v_i \notin FWR$, this implies $ea_i > t_0 + b$, meaning that the subpath from v_0 to v_i already exceeds the budget, which contradicts the assumption that p is feasible. Case 2: Analogously, if $v_i \notin BWR$, then $ld_i < t_0$, hence p cannot be a feasible path.* □

5. COMPUTING REACHABILITY

Having clarified the importance of the sets of reachable vertices, we will now go into detail on their computation. Both sets are generated by expansions around source and destination with "time-radius b". By the above lemma, all vertices along all feasible paths are contained in the intersection of the two reachability regions.

In order to compute the set of forward-reachable vertices (for a given source v_0, departure time t_0, and a time budget b), we perform a modified time-dependent all-target Dijkstra search starting at v_0 at time t_0, called FWR and illustrated in Algorithm 1. All earliest arrival times are initialized as ∞ except for the source v_0 for which it is 0. Additional to a result set containing all visited vertices, we maintain a priority queue containing vertices sorted in ascending order by the earliest arrival time of each vertex. In every iteration, the top element of the queue is removed. In the for-loop spanning lines 8-21, the outgoing arcs of the current vertex v_i are explored. In line 9, the travel time from v_i to its neighbor v_j is added to the earliest arrival time at v_i. If this value does not exceed the budget b, v_j is forward-reachable. If v_j was previously not visited, ea_j is set accordingly and v_j is added to the queue as well as

the result set. If v_j was previously visited and the updated earliest arrival time is better than the previous, it is updated accordingly. Note that once a vertex v_i has been dequeued, the label ea_i is final. This is equivalent to the Dijkstra property which ensures that any processed vertex is reached by the shortest path, in this case the fastest time-dependent path. Finally, the result set contains all vertices reachable from the source within the budget.

The set of backward-reachable vertices is computed in a similar fashion but slightly more complicated. The procedure is presented in Algorithm 2 and called BWR. In addition to the result set containing all backward-reachable vertices, we maintain a queue of vertices sorted in descending order by ld_i. Initially, the latest departure at the destination v_N is set to the latest possible arrival time, $t_0 + b$. The algorithm operates backwards from the destination, following incoming arcs which are explored in the for-loop spanning lines 8-25. As was illustrated in the above example (Figure 3), in order to find the latest departure, it is required to linearly scan the time windows. When exploring an incoming arc from v_i to v_j, $τ_k$ is initially set to the latest time window not exceeding the latest departure from v_j. If at $τ_k$ the travel time from v_i to v_j is too high (i.e., v_j cannot be reached in time for its latest departure), k is decremented implying an earlier departure at v_i but possibly also a different travel time along the arc. Note the $[−ε]$ in line 13, corresponding to the aforementioned special case that $l̂d_i = τ_l$ for some l. In this case, $l̂d_i$ is decreased by a small number in order to avoid computational errors which may occur at time window boundaries. Once a valid latest departure time for v_i is found, ld_i is set accordingly. If v_i was previously visited, ld_i is only set if it is later than the previous latest departure. Upon termination, all vertices which are backward-reachable from the destination within the given budget are in the result set and labeled with their respective

(a) FWEST   (b) EALD
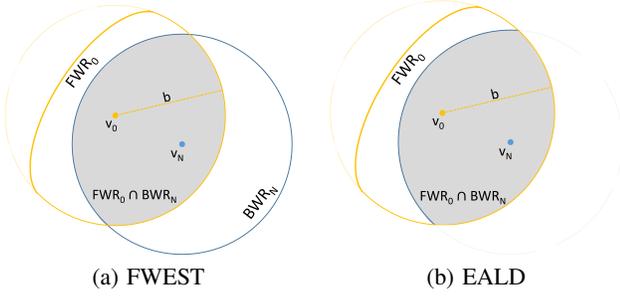
Figure 4: Illustration of the effects of the proposed pruning techniques.

latest departure.

6. COMPUTING SOLUTIONS

Due to the NP-hardness of the problem, we forgo giving an optimal solution as it would not meet our requirements of providing solutions on large graphs in efficient time. Instead, we propose a heuristic solution to the 2TD-AOP. So far no other solution to the 2TD-AOP exists. Our dynamic programming solution recursively fills gaps with arcs in order to improve the value while not exceeding the budget. Upon initialization, the set of arcs `arcs` is empty, the set of gaps `gaps` contains (v_0, v_N), i.e., the gap given by the query. Assume that in the first iteration the arc (v_i, v_j) is inserted. Then `arcs` holds the new arc (v_i, v_j) and `gaps` holds two new gaps, (v_0, v_i) and (v_j, v_N). Into the new gaps, further arcs will be inserted. This procedure is repeated until the budget is exhausted. Hence, the 2TD-AOP is solved in a divide-and-conquer manner by recursive insertions. We therefore refer to the algorithm, pseudo-coded in Algorithm 3, as RECINSERT.

We will now explain in detail how the algorithm proceeds. The two sorted lists `gaps` and `arcs` hold two-tuples of vertices. `gaps` holds pairs between which arcs are to be inserted, and `arcs` holds the converse information, i.e., the arcs already inserted. With every insertion, the budget is decreased by the duration of the detour that has to be taken to insert the new arc. However, this detour is not explicitly computed as a path. Instead, it is implicitly derived from the information generated during FWR and BWR. The earliest arrivals and latest departures of every gap allow to decide whether an arc can be *feasibly inserted*, i.e., without exceeding the budget. This is crucial to the performance of our algorithm. Established meta-heuristics usually compute, improve and perturb solutions. Starting from an initial path, promising arcs are (usually randomly) inserted and under certain conditions (usually randomly) deleted, to avoid local value maxima. Instead, RECINSERT holds a sorted list of arcs which may be feasible inserted into a path. The actual paths between these arcs, however, are not computed until the budget is exhausted.

The arc which is to be inserted into a gap is chosen based on a heuristic. Among all arcs which can be feasibly inserted, the arc is selected which yields the best ratio of value to detour. This step is performed in the `for`-loop spanning lines 9 to 19. As established, all feasibly insertable arcs lie in the intersection of the reachability regions. The intersection is performed in line 8, therefore the set in line 9 comprises all arcs which can be feasibly inserted. In line 10, the detour incurred by inserting an arc is expressed as a decrease in budget. Assume that for the gap between v_i and v_j, the arc (v_m, v_n) is examined for insertion. ea_m is the earliest arrival at v_m coming from v_i, and ld_n is the latest departure from v_n to

---

Algorithm 3 RECINSERT

```
 1: procedure RECINSERT(G, gaps, arcs, b)
 2:     bestArc ← (null, null)
 3:     bestGap ← (null, null)
 4:     bestRatio ← 0
 5:     for (v_i, v_j) ∈ gaps do
 6:         FWR_i ← FWR(G, v_i, b)
 7:         BWR_j ← BWR(G, v_j, b)
 8:         G'_ij ← FWR_i ∩ BWR_j
 9:         for (v_m, v_n) in G'_ij: ea_m + tt_{m,n}(ea_m) < ld_n do
10:             b_- ← ld_n − ea_m − tt_{m,n}(ea_m)
11:             val_+ ← val_{m,n}(ea_m)
12:             ratio ← val_+ · b_-^{-1}
13:             if ratio > bestRatio then
14:                 bestRatio = ratio
15:                 bestArc = (v_m, v_n)
16:                 bestGap = (v_i, v_j)
17:             end if
18:         end for
19:     end for
20:     gaps.remove(bestGap)
21:     if bestArc = null then
22:         compute fastest path to close all gaps
23:         return concatenation of fastest paths
24:     else
25:         gaps.add((v_i, v_m))      ▷ at former index of bestGap
26:         gaps.add((v_n, v_j))      ▷ after (v_i, v_m)
27:         arcs.add((v_i, v_j))
28:     end if
29:     RECINSERT(G, gaps, arcs, b_-)
30: end procedure
```

reach v_j "in time" (depending on the recursive budget b). Thus, $ld_n - ea_m - tt_{m,n}(ea_m)$ corresponds to the budget that would remain when taking the detour of traveling from v_i to v_m, then along (m, n) (at time ea_m) and from v_n to v_j. Hence, if inserting (v_m, v_n), the budget will be set to $ld_n - ea_m - tt_{m,n}(ea_m)$. In line 11, the gain in value is estimated. Note that the actual gain in value may be higher (due to other arcs along the path). However, in order to compute the actual gain in value, the fastest paths to and from each arc would have to be computed. Also, the value of all succeeding arcs would have to be reassessed. For the sake of efficiency, we therefore only estimate the value as proposed. The estimated gain in value divided by the decrease in budget yields the ratio by which we evaluate the arcs. The arc with the best ratio is inserted, i.e., the arc is added to `arcs` as a pair of vertices, and two pairs of vertices are added to `gaps` (the gaps left and right of the new arc). As long as at least one arc qualifies for insertion, RECINSERT is recursively called with the extended sets `gaps` and `arcs` and the reduced budget b_-. If no arc can be inserted (into any of the gaps), the time-dependent fastest paths for all gaps are computed and the concatenation is returned which constitutes our approximate solution to the 2TD-AOP.

6.1 Pruning Strategies

Let us discuss three variations which are omitted in the algorithm description for reasons of brevity. (i), it is recommended to employ a simple forward estimation during FWR. (ii), it is recommended to compute the intersection of FWR$_i$ and BWR$_j$ during computation of BWR$_j$. (iii), it is possible to make the recursive call with a restricted portion of the graph. Variations (i) and (ii) are based on the following observation. The sets FWR$_i$ and BWR$_j$ are not

needed separately, only their intersection is needed. Therefore, we may exclude vertices which are not contained in the intersection during computation of the respective sets.

Regarding (i): We propose to employ a simple forward estimation for pruning. If information about the speed limits is available, the following lower bound can be used with negligible overhead. During FWR, in line 10, the `if`-clause may be tightened:

$$ea_j + \frac{ed_{jN}}{ms} < t_0 + b$$

where ed_{jN} is the Euclidean distance from v_j to the destination v_N and ms is maximum speed in the given network. If traveling from v_j to the destination at maximum speed exceeds the budget, so will any actual path. Any such vertex cannot be forward-reachable. Figure 4(a) illustrates the effect of this pruning strategy yielding a trimmed forward-reachability region. We refer to this pruning strategy as *FWEST* pruning. At marginal computational cost, FWEST yields a perceptible reduction in candidates. In our experiments, over 15% of vertices are pruned during FWR using this approach. Of course, the method holds analogously for BWR and can equally be applied. However, during BWR it is recommended to use the following pruning technique which is superior but exclusive to BWR.

Regarding (ii): One may use that for all vertices in the intersection $\text{FWR}_i \cap \text{BWR}_j$ holds $ea_i \leq ld_i$. This follows from the definition:

$$ea_i > ld_i$$
$$\Leftrightarrow t_0 + tt(p^\star_{0i}) > t_0 + b - tt(p^\star_{iN})$$
$$\Leftrightarrow tt(p^\star_{0i}) + tt(p^\star_{iN}) > b$$

where $p^\star_{ij}$ denotes the fastest path from v_i to v_j at ea_i. Hence, the `if`-clause in line 14 of BWR can be tightened to $ld_i \geq ea_i$. In this case, the simple forward estimation cannot be applied because

$$ea_i \geq t_0 \frac{ed_{0i}}{ms} \geq t_0$$

By construction, ea_i is based on an actual path expansion and therefore serves as a tighter bound than any estimation. We refer to this pruning strategy as EALD pruning. Figure 4(b) illustrates the effect of this pruning strategy, where the vertices which are not forward-reachable are not visited during BWR. Thus, instead of computing both reachability regions separately and subsequently their intersection, the information of FWR may be used in BWR. This limits the search space and implicitly computes the intersection. Every vertex in the intersection will be visited during BWR. Unvisited vertices are either not forward-reachable or backward-reachable. The effect of this pruning strategy is remarkable. Particularly when taking into account that it causes no computational overhead. When visiting a vertex during BWR, it suffices to retrieve its earliest arrival to decide if it qualifies for pruning.

Regarding (iii): It is possible but not necessarily recommended to issue the recursive call of RECINSERT with a restricted portion of the graph. As established in Lemma 1, all vertices along feasible paths are in the intersection of the reachability regions. This property holds recursively. For instance, consider the insertion of a first arc (v_i, v_j) into the initial gap (v_0, v_N). In the next recursion, two gaps will be examined, (v_0, v_i) and (v_j, v_N), as illustrated in Figure 2(b). Consider the first gap (v_0, v_i), the respective reachability regions (forward from v_0 and backward from v_i) are contained in the original reachability regions (forward from v_0 and backward from v_N. Moreover, the intersection is contained in the original intersection. Thus, it suffices to input a restricted graph into the next
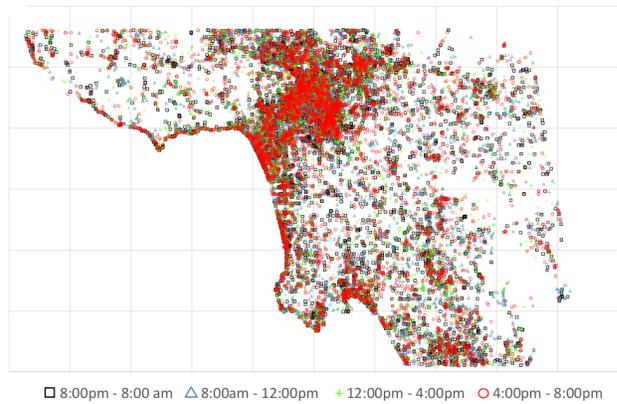
Figure 5: Value arcs in experimental road network of Los Angeles. Each value arc corresponds to a colored shape which indicates the arc's value peak interval.

recursion:

$$G_{\text{next}} = \bigcup \{ G'_{mn} \mid (v_m, v_n) \in \texttt{gaps} \}$$

Albeit correct, actually computing the union of these graphs may cause more computational overhead than it saves. In our experiments, computing the restricted graph has not proved efficient. We input the initial intersection of FWR_0 and BWR_N into all recursions as it restricts the graph significantly while not requiring costly unions of graphs.

Let us note that Gunawan et. al use very similar terminology in their approach to solving the TD-OP [25]. However, there are significant differences to the approach presented here. First, the authors solve TD-OP where travel times are time-dependent but values are assumed to be static and associated with vertices. Second, the authors propose a different algorithmic approach. While also using the terms earliest arrival and latest departure, they are computed differently and for different purposes. Given a specific path, a routine called FORWARDPROPAGATION computes the earliest arrival times at all vertices along the path. Analogously, BACKWARDPROPAGATION computes the latest departures. In comparison, our FWR and BWR procedures compute earliest arrivals and latest departures for all vertices which are feasibly insertable into a given gap. From this information, the travel time for a detour to any feasibly insertable arc is derived. This is done without explicit path computation. While we use dynamic programming in a divide-and-conquer manner, Gunawan et. al follow an Iterative Local Search (optionally Variable Neighbor Descent) approach. where actual paths are computed, improved and perturbed. In their experiments, they produce near-optimal to optimal solutions with a maximum runtime of one second. Their approach is evaluated on small data sets with at most 40 vertices and 40 time windows. The arcs correspond to precomputed fastest paths. This restrictive scenario is not comparable to the use case we are modeling in this paper where fast response times are required for a more complex problem on graphs which are orders of magnitude larger.

7. EXPERIMENTAL EVALUATION

We evaluate our solution to the 2TD-AOP on a time-dependent road network of Los Angeles, CA, USA [13, 28] which contains about 500K vertices and 1M arcs. The travel time functions in this

peak interval	# of arcs	% of all arcs
8:00 am - 12:00 pm	6,238	0.57%
12:00 pm - 4:00 pm	7,366	0.67%
4:00 pm - 8:00 pm	5,851	0.53%
8:00 pm - 8:00 am	6,417	0.59%
all intervals	25,872	2.38%

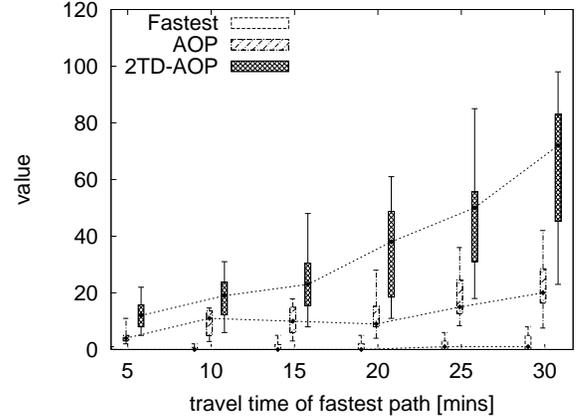**Table 1: Statistics about arcs with non-zero value.**
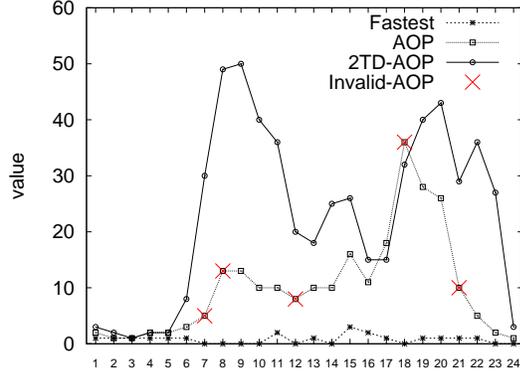
Figure 6: Value results for the introductory example in Figure 1.

Figure 7: Value results for varying fastest path lengths.

network are derived from large-scale and high-resolution (both spatial and temporal) sensor data (both live and historic) from different transportation authorities in California. The step length of the piecewise constant travel time functions is 5 minutes, details can be found in [12]. Based on the data, streets are uncongested between 9 pm and 6 am. Therefore, travel times are assumed to be static in this interval.

The time-dependent value functions are derived from a geo-tagged set of Flickr photos [39] consisting of 217,391 photos in the area of Los Angeles. For each arc (v_i, v_j) the value $val_{i,j}(t)$ is the number of photos which have (v_i, v_j) as their nearest arc (but not further than one kilometer) and which were taken in the hour-interval that contains t. Hence, the step length of the piecewise constant value functions is 60 minutes. Figure 5 illustrates the size of the network and the distribution of non-zero value arcs. Each arc which has positive value in some time window is visualized as a colored shape. Shape and color depend on the time interval in which the value of the arc is the highest. For example, an arc with highest value between 6:00 and 7:00 pm will be displayed as red circle. Table 1 presents the absolute and relative numbers of arcs with non-zero value. The number of photos taken during the day is about three times higher than during night.

For the following experiments, we arrange paths into six time buckets. We randomly draw source and destination vertices from the network and compute the respective time-dependent fastest paths using [12]. If the travel time of a path is k minutes $\pm\varepsilon$, its source and destination pair is assigned to the k minutes time bucket for $k \in \{5, 10, 15, 20, 25, 30\}$. Each time bucket consists of twenty source and destination pairs. Since travel times are empirically static at night and value peaks less frequent at night, we randomly sample departure times from the interval 8 am to 8 pm. We make a case for twofold time-dependence, therefore, a network with significant time-dependent impact on travel time and value is essential. Our standard experimental setting has the following inputs. Stan-

dard time bucket is 20 minutes, standard query budget is 200% (i.e., 40 minutes), and departure times are randomly drawn from 8 am to 8 pm. Deviating values are explicitly mentioned. For each query, the departure time is the same across time-dependent algorithms.

We omit providing results of the MIP formulation given in Section 3 generated by a complex solver. This is due to the exorbitant computatial requirements. Even for small instances with at most 100 vertices, solving a less complex MIP often requires days [47, 25]. However, for comparison, we compute optimal solutions with a custom dynamic programming method.

7.1 Experimental Results

In a first set of experiments, we compare the value of result paths generated by different algorithms taking varying degrees of information into account. As in the introductory example, we compare 2TD-AOP paths computed by RECINSERT with static scenic paths and to time-dependent fastest paths. For the static scenic paths, we use the AOP solution presented in [38] which employs an Iterative Local Search approach combined with spatial pruning techniques. In accordance with the metaheuristic, it generates an initial solution which is subsequently improved by inserted arcs and perturbed by deleting arcs. The insertion follows a heuristic, the deletion is pseudo-random. We refer to solutions created by this approach as AOP paths and call this algorithm AOP-ILS. It operates on a network with static travel time and static value which are both derived by averaging the respective time-dependent functions. Note that AOP queries are given the same time budget but require no departure time. For the time-dependent fastest paths, we use the solution proposed in [12] which relies on hierarchical routing and forward estimations to conduct efficient bidirectional path computation. We refer to the results as Fastest paths. The results generated by RECINSERT are referred to as 2TD-AOP paths.

The values and travel times of all paths are computed w.r.t. the time-dependent functions. For a given path, the value and travel time are the sum of value and travel time of each arc along the path upon arrival. Recall the introductory example in Figure 1. For varying departure times in this example, the result values are displayed in Figure 6. As mentioned in the example, it is possible that paths computed in the static network are invalid in the time-dependent network. This is the case if the time-dependent travel time exceeds the budget. Static paths where this is the case are marked with a red cross. For instance, when departing at 6 pm, the time-dependent travel time is 37 minutes, exceeding the budget of 25 minutes by almost 50%.
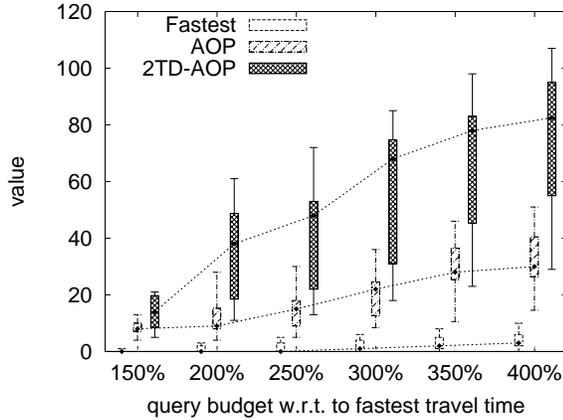
Figure 8: Value results for varying time budget.

Figure 9: Percentage of visited vertices when employing the proposed pruning techniques.

Figure 10: Processing time for varying path lengths.

Figure 7 shows the value for the different approaches. Fastest paths serve a baseline which shows how much value can be attained "by chance" as it does not optimize for value. Indeed, Fastest paths collect negligible value. Since the non-zero value arcs are rather scarce, this is not surprising. 2TD-AOP and AOP, in contrast, optimize for value. It can be observed that on average 2TD-AOP generates three times the value of AOP. This gap widens for increasing path length.

Figure 8 illuminates the effect of the query budget. We observe that with increasing budget, 2TD-AOP paths generate significantly more value than AOP solutions. This is particularly noteworthy as the problem becomes more complex with increasing budget. A higher budget allows for greater detours which in turn increases the search space. RECINSERT leverages this effect to create better solutions. In contrast, AOP results barely improve. Therefore, albeit increasing complexity, twofold time-dependence is able to enhance results considerably.

As mentioned before, static paths may prove invalid when considering time-dependent travel time. Given the departure times of 2TD-AOP and Fastest, we determine the portion of static paths that prove invalid when evaluated in the time-dependent network. The percentage of AOP which are invalid increases from 25% to almost 50% across the increasing time buckets. Note that Figures 7 and 8 show the values of all AOP paths, both valid and invalid, as there is no significant difference. However, taking into account the percentage of invalid results and the low overall value, static results cannot compete with twofold time-dependent results.

The proposed pruning strategies are evaluated in Figure 9. It shows the percentage of visited vertices relative to the number of visited vertices when employing no pruning technique. In this case, all vertices in FWR and BWR have to be visited. In Section 6 we propose two pruning techniques for RECINSERT, FWEST and EALD. FWEST is a basic forward estimation which can be applied during FWR and BWR. EALD uses the information generated during FWR to limit the search space of BWR and can only be applied during BWR. During BWR, EALD is superior to FWEST. Hence, when both are available, FWEST is obsolete during BWR. In the experiments displayed in Figure 9, we show the effect of each pruning technique separately as well as of both pruning techniques combined. When only employing FWEST (during both FWR and BWR), almost 40% of the vertices can be pruned. Although EALD can only be applied during BWR, it prunes about 45% of the vertices. As FWEST causes negligible and EALD causes no compu-
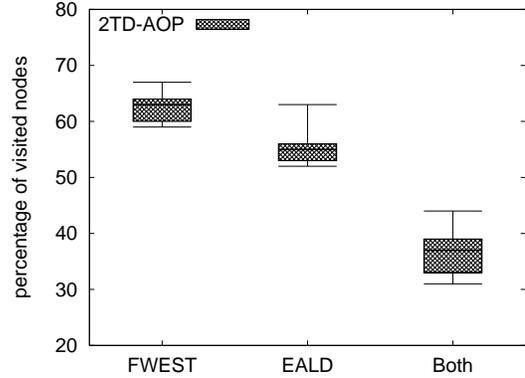
tational overhead, it is recommended to employ both. Combining both approaches, 65% of all vertices in FWR and BWR can be pruned.
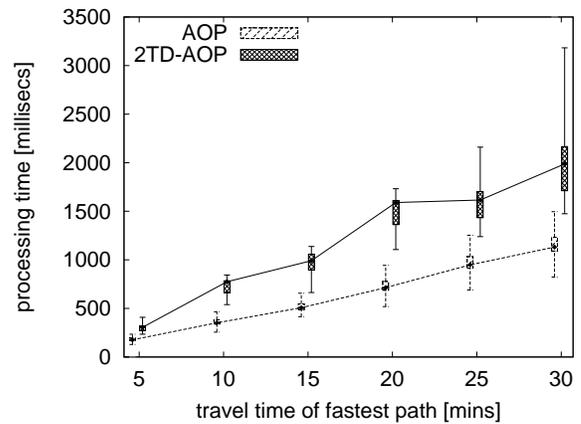
In terms of general complexity, the 2TD-AOP is considerably more intricate than the AOP. While both problems are NP-hard, twofold time-dependence incurs additional complexity. This is mainly due to two aspects. First, fastest path computation is more efficient in static networks. Second, in order to assess the time-dependent value of an arc, it does not suffice to simply retrieve its value. The time frame of arrival has to be considered. In light of these challenges, the efficiency of RECINSERT is all the more surprising. Figure 10 shows the processing times of RECINSERT solving the 2TD-AOP and AOP-ILS solving the AOP. 2TD-AOP solutions can be computed in less than twice the processing time of AOP solution, usually under two seconds. The increase in processing time for both algorithms in linear in the travel time of the fastest path. Taking the result quality into consideration, the additional processing time of RECINSERT compensates for the often invalid AOP paths with less value. It should also be noted that in an application where fixed response times are required, it would be possible to terminate the arc insertion in RECINSERT prematurely and construct an early result path.

Finally, we compare 2TD-AOP and AOP paths to optimal solutions in terms of their value. The optimal solutions are computed with a custom dynamic programming approach using a depth-first
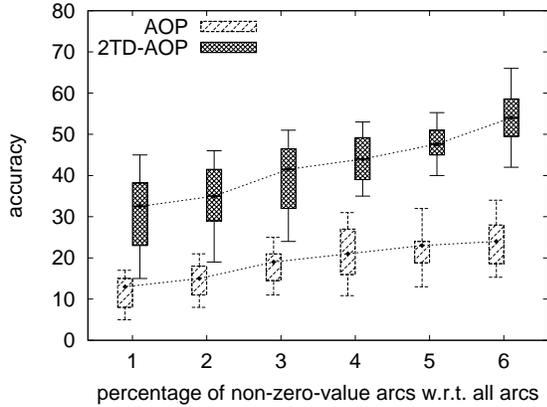
Figure 11: Accuracy of result values relative to optimal results when varying value arc density.

search and several pruning techniques. Any naïve algorithm, much like a complex solver, would not be able to compute solutions in feasible time. Our custom approach takes between 3 to 6 hours per optimal path computation for the standard query setting. Even for minor instances where the fastest path takes 5 minutes and the budget is 10 minutes, computing an optimal solution takes between 1 and 2 hours. Hence, generating optimal solutions to the 2TD-AOP is infeasible. This is irrespective of the application. Users will not tolerate response times in the hours, and precomputation is not possible in time-dependent networks. Figure 11 compares the accuracy of 2TD-AOP paths generated by RECINSERT and of AOP paths generated by AOP-ILS. Values are shown relative to values of optimal paths. Figure 11 also illustrates the effect that the density of non-zero value arcs has on the result. The value is given relative to the optimal value for different networks when varying the percentage of non-zero value arcs. In our original network about 2% of the arcs have non-zero value during some time window. For this experiment, we randomly copied value functions to arcs with zero value to increase density to 4%, 6%, 8% and 10% of all arcs. Also, we set some of the value functions to zero, generating a network with 1% value arcs. As both algorithms follow a heuristic during arc insertion, they are more easily sidetracked in a network with many value arcs. When the density of value arcs decreases, however, the employed heuristics prove effective. In this case, RECINSERT result paths attain between 50% and 60% accuracy. The static solutions, in contrast, hardly exceed 25%. Again, this establishes the superiority of 2TD-AOP over AOP solutions. While computing an optimal solution takes up to 6 hours, RECINSERT takes about 2 seconds. Thus, RECINSERT produces paths with about 50% accuracy in 10^{-5} the processing time. As the 2TD-AOP cannot be expected to be solved optimally in efficient time, RECINSERT yields a promising trade-off.

In conclusion, our experimental evaluation has established the importance of twofold time-dependence. Solutions to the 2TD-AOP are superior to solutions to the static AOP. For increasingly complex query settings, this trend is intensified. Furthermore, RECINSERT computes solutions that achieve 50% accuracy compared to optimal results within seconds. The efficiency can essentially be attributed to the pruning techniques employed in RECINSERT. As of now, no time-dependent AOP has been evaluated on large-scale real-world road networks. We prove that even in such networks, the higly complex 2TD-AOP can be feasibly solved.

8. CONCLUSIONS

In this paper, we present the Twofold Time-Dependent Arc Orienteering Problem (2TD-AOP) a novel extension of the family of Orienteering Problems (OP). In this family of NP-hard graph problems, the goal is to find a path from a given source to a given destination within a given time budget which, additionally, maximizes the value collected along the way. In comparison to static versions of this problem, 2TD-AOP allows for travel times and value functions to be time-dependent. The incorporation of time-dependent values is a novel extension which has not been studied yet. The importance of time-dependent values is showcased and experimentally substantiated. Due to the NP-hardness and the twofold time-dependence, we provide a heuristic approximation for solving the 2TD-AOP. Our approach is evaluated on a large-scale real-world road network. Thus far, no time-dependent variant of the AOP has been evaluated on a network of this dimension. Twofold time-dependent results gain significantly more value than static results, showing the importance of twofold time-dependence empirically. Our algorithm produces result paths with 50% accuracy where no optimal solution is feasible.

9. ADDITIONAL AUTHORS

Additional authors: Matthias Schubert (LMU Munich, Oettingenstraße 67, 80583 Munich, Germany) schubert@dbs.ifi.lmu.de)

10. REFERENCES

[1] J. Aráoz, E. Fernández, and C. Zoltan. Privatized rural postman problems. *Computers & operations research*, 33(12):3432–3449, 2006.

[2] C. Archetti and M. G. Speranza. Arc routing problems with profits. *Arc Routing: Problems, Methods, and Applications, MOS-SIAM Series on Optimization*, pages 257–284, 2013.

[3] I. Brilhante, J. A. Macedo, F. M. Nardini, R. Perego, and C. Renso. Where shall we go today?: Planning touristic tours with tripbuilder. In *Proceedings of the 22Nd ACM International Conference on Information & Knowledge Management*, CIKM '13, pages 757–762, New York, NY, USA, 2013. ACM.

[4] I.-M. Chao, B. L. Golden, and E. A. Wasil. A fast and effective heuristic for the orienteering problem. *European Journal of Operational Research*, 88(3):475–489, 1996.

[5] C. Chen, D. Zhang, B. Guo, X. Ma, G. Pan, and Z. Wu. Tripplanner: Personalized trip planning leveraging heterogeneous crowdsourced digital footprints. *IEEE Transactions on Intelligent Transportation Systems*, 16(3):1259–1273, 2015.

[6] H. Chen, W.-S. Ku, M.-T. Sun, and R. Zimmermann. The multi-rule partial sequenced route query. In *Proceedings of the 16th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (ACMGIS'08)*, pages 10:1–10:10, 2008.

[7] Z. Chen, R. Fu, Z. Zhao, Z. Liu, L. Xia, L. Chen, P. Cheng, C. C. Cao, Y. Tong, and C. J. Zhang. gmission: A general spatial crowdsourcing platform. *Proceedings of the VLDB Endowment*, 7(13):1629–1632, 2014.

[8] D. Cheng, M. Hadjieleftheriou, G. Kollios, F. Li, and S.-H. Teng. On trip planning queries in spatial databases. In *Proceedings of the 9th International Conference on Advances in Spatial and Temporal Databases (SSTD'05)*, pages 273–290, 2005.

[9] M. De Choudhury, M. Feldman, S. Amer-Yahia, N. Golbandi, R. Lempel, and C. Yu. Automatic construction of travel itineraries using social breadcrumbs. In *Proceedings of the 21st ACM Conference on Hypertext and Hypermedia*, HT '10, pages 35–44, New York, NY, USA, 2010. ACM.

[10] R. Deitch and S. P. Ladany. The one-period bus touring problem: Solved by an effective heuristic for the orienteering

tour problem and improvement algorithm. *European Journal of Operational Research*, 127(1):69–77, 2000.

[11] A. Delis, V. Efstathiou, and V. Verroios. Reaching available public parking spaces in urban environments using ad hoc networking. In *Proceedings of the 12th International Conference on Mobile Data Management (MDM'11)*, pages 141–151, 2011.

[12] U. Demiryurek, F. Banaei-Kashani, C. Shahabi, and A. Ranganathan. Online computation of fastest path in time-dependent spatial networks. In *International Symposium on Spatial and Temporal Databases*, pages 92–111. Springer, 2011.

[13] U. Demiryurek, F. B. Kashani, and C. Shahabi. A case for time-dependent shortest path computation in spatial networks. In *18th ACM SIGSPATIAL International Symposium on Advances in Geographic Information Systems, ACM-GIS 2010, November 3-5, 2010, San Jose, CA, USA, Proceedings*, pages 474–477, 2010.

[14] D. Feillet, P. Dejax, and M. Gendreau. Traveling salesman problems with profits. *Transportation science*, 39(2):188–205, 2005.

[15] F. V. Fomin and A. Lingas. Approximation algorithms for time-dependent orienteering. *Information Processing Letters*, 83(2):57–62, 2002.

[16] A. García, O. Arbelaitz, M. T. Linaza, P. Vansteenwegen, and W. Souffriau. Personalized tourist route generation. In *Current Trends in Web Engineering - 10th International Conference on Web Engineering, ICWE 2010 Workshops, Vienna, Austria, July 2010, Revised Selected Papers*, pages 486–497, 2010.

[17] A. Garcia, M. T. Linaza, O. Arbelaitz, and P. Vansteenwegen. Intelligent routing system for a personalised electronic tourist guide. *Information and Communication Technologies in Tourism 2009*, pages 185–197, 2009.

[18] D. Gavalas, C. Konstantopoulos, K. Mastakas, and G. Pantziou. A survey on algorithmic approaches for solving tourist trip design problems. *Journal of Heuristics*, 20(3):291–328, 2014.

[19] D. Gavalas, C. Konstantopoulos, K. Mastakas, G. Pantziou, and N. Vathis. Efficient heuristics for the time dependent team orienteering problem with time windows. In *Proceedings of the First International Conference on Applied Algorithms - Volume 8321*, ICAA 2014, pages 152–163, New York, NY, USA, 2014. Springer-Verlag New York, Inc.

[20] D. Gavalas, C. Konstantopoulos, K. Mastakas, G. Pantziou, and N. Vathis. Approximation algorithms for the arc orienteering problem. *Information Processing Letters*, 115(2):313–315, 2015.

[21] M. Gendreau, G. Laporte, and F. Semet. A tabu search heuristic for the undirected selective travelling salesman problem. *European Journal of Operational Research*, 106(2):539–545, 1998.

[22] B. L. Golden, L. Levy, and R. Vohra. The orienteering problem. *Naval research logistics*, 34(3):307–318, 1987.

[23] B. L. Golden, Q. Wang, and L. Liu. A multifaceted heuristic for the orienteering problem. *Naval Research Logistics (NRL)*, 35(3):359–366, 1988.

[24] A. Gunawan, H. C. Lau, and P. Vansteenwegen. Orienteering problem: A survey of recent variants, solution approaches and applications. *European Journal of Operational Research*, 2016.

[25] A. Gunawan, Z. Yuan, and H. C. Lau. A mathematical model and metaheuristics for time dependent orienteering problem. pages 202–217, 2014.

[26] H. Hsieh and C. Li. Mining and planning time-aware routes from check-in data. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management, CIKM 2014, Shanghai, China, November 3-7, 2014*, pages 481–490, 2014.

[27] H. P. Hsieh, C. T. Li, and S. D. Lin. Exploiting large-scale check-in data to recommend time-sensitive routes. In *Proceedings of the ACM SIGKDD International Workshop on Urban Computing*. ACM, 2012.

[28] H. Jagadish, J. Gehrke, A. Labrinidis, Y. Papakonstantinou, J. M. Patel, R. Ramakrishnan, and C. Shahabi. Big data and its technical challenges. *Communications of the ACM*, 57(7):86–94, 2014.

[29] G. Jossé, K. A. Schmid, and M. Schubert. Probabilistic resource route queries with reappearance. In *EDBT: 18th International Conference on Extending Database Technology, Brussels, Belgium*, 2015.

[30] Y. Kanza, R. Levin, E. Safra, and Y. Sagiv. Interactive route search in the presence of order constraints. *The International Journal on Very Large Data Bases (VLDB'10)*, pages 117–128, 2010.

[31] Y. Kanza, E. Safra, and Y. Sagiv. Route search over probabilistic geospatial data. In *Proceedings of the 11th International Symposium on Advances in Spatial and Temporal Databases (SSTD'09)*, pages 153–170, 2009.

[32] L. Kazemi and C. Shahabi. Geocrowd: enabling query answering with spatial crowdsourcing. In *Proceedings of the 20th International Conference on Advances in Geographic Information Systems*, pages 189–198. ACM, 2012.

[33] M. Kolahdouzan, C. Shahabi, and M. Sharifzadeh. The optimal sequenced route query. *The International Journal on Very Large Data Bases (VLDB'08)*, 17(4):765–787, 2008.

[34] G. Laporte and S. Martello. The selective travelling salesman problem. *Discrete applied mathematics*, 26(2-3):193–207, 1990.

[35] R. Levin, Y. Kanza, E. Safra, and Y. Sagiv. An interactive approach to route search. In *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (ACMGIS'09)*, pages 408–411, 2009.

[36] J. Li. Research on team orienteering problem with dynamic travel times. *Journal of Software*, 7(2):249–255, 2012.

[37] Y.-C. Liang, S. Kulturel-Konak, and A. E. Smith. Meta heuristics for the orienteering problem. In *Evolutionary Computation, 2002. CEC'02. Proceedings of the 2002 Congress on*, volume 1, pages 384–389. IEEE, 2002.

[38] Y. Lu and C. Shahabi. An arc orienteering algorithm to find the most scenic path on a large-scale road network. In *Proceedings of the 23rd SIGSPATIAL International Conference on Advances in Geographic Information Systems*, page 46. ACM, 2015.

[39] H. Mousselly-Sergieh, D. Watzinger, B. Huber, M. Döller, E. Egyed-Zsigmond, and H. Kosch. World-wide scale geotagged image dataset for automatic image annotation and reverse geotagging. In *ACM Multimedia Systems*, pages 47–52, 2014.

[40] R. Ramesh and K. M. Brown. An efficient four-phase heuristic for the generalized orienteering problem. *Computers & Operations Research*, 18(2):151–165, 1991.

[41] G. Skoumas, K. A. Schmid, G. Jossé, M. Schubert, M. A. Nascimento, A. Züfle, M. Renz, and D. Pfoser. Knowledge-enriched route computation. In *Advances in Spatial and Temporal Databases - 14th International Symposium, SSTD 2015, Hong Kong, China, August 26-28, 2015. Proceedings*, pages 157–176, 2015.

[42] W. Souffriau, P. Vansteenwegen, G. V. Berghe, and D. V. Oudheusden. The planning of cycle trips in the province of east flanders. *Omega*, 39(2):209 – 213, 2011.

[43] M. F. Tasgetiren. A genetic algorithm with an adaptive penalty function for the orienteering problem. *Journal of Economic and Social Research*, 4(2):1–26, 2001.

[44] H. To, C. Shahabi, and L. Kazemi. A server-assigned spatial crowdsourcing framework. *ACM Transactions on Spatial Algorithms and Systems*, 1(1):2, 2015.

[45] T. Tsiligirides. Heuristic methods applied to orienteering. *Journal of the Operational Research Society*, 35(9):797–809, 1984.

[46] P. Vansteenwegen, W. Souffriau, and D. Van Oudheusden. The orienteering problem: A survey. *European Journal of*

Operational Research, 209(1):1–10, 2011.

[47] C. Verbeeck, K. Sörensen, E. Aghezzaf, and P. Vansteenwegen. A fast solution method for the time-dependent orienteering problem. *European Journal of Operational Research*, 236(2):419–432, 2014.

[48] C. Verbeeck, P. Vansteenwegen, and E.-H. Aghezzaf. An extension of the arc orienteering problem and its application to cycle trip planning. *Transportation Research Part E: Logistics and Transportation Review*, 68:64 – 78, 2014.