



**HAL**  
open science

## MUSE: Multimodal Separators for Efficient Route Planning in Transportation Networks

Mohamed Amine Falek, Cristel Pelsser, Sébastien Julien, Fabrice Theoleyre

► **To cite this version:**

Mohamed Amine Falek, Cristel Pelsser, Sébastien Julien, Fabrice Theoleyre. MUSE: Multimodal Separators for Efficient Route Planning in Transportation Networks. *Transportation Science*, 2022, 56 (2), 10.1287/trsc.2021.1104 . hal-03402845

**HAL Id: hal-03402845**

**<https://hal.science/hal-03402845v1>**

Submitted on 25 Oct 2021

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

# MUSE: Multimodal Separators for Efficient Route Planning in Transportation Networks

Amine Mohamed Falek

Technology & Strategy group, 4 rue de Dublin, 67300 Schiltigheim, France.  
falek@unistra.fr - a.falek@technologyandstrategy.com

Cristel Pelsser

ICube Lab, CNRS/University of Strasbourg, Pole API, Boulevard Sebastien Brant, 67412 Illkirch, France.  
pelsser@unistra.fr

Sebastien Julien

Technology & Strategy group, 4 rue de Dublin, 67300 Schiltigheim, France.  
s.julien@technologyandstrategy.com

Fabrice Theoleyre

ICube Lab, CNRS/University of Strasbourg, Pole API, Boulevard Sebastien Brant, 67412 Illkirch, France.  
theoleyre@unistra.fr <http://www.theoleyre.eu>

Many algorithms compute shortest-path queries in mere microseconds on continental-scale networks. Most solutions are, however, tailored to either road or public transit networks in isolation. To fully exploit the transportation infrastructure, multimodal algorithms are sought to compute shortest-paths combining various modes of transportation. Nonetheless, current solutions still lack performance to efficiently handle interactive queries under realistic network conditions where traffic jams, public transit cancelations, or delays often occur. We present MUSE, a new multimodal algorithm based on graph separators to compute shortest travel time paths. It partitions the network into independent, smaller regions, enabling fast and scalable preprocessing. The partition is common to all modes and independent of traffic conditions so that the preprocessing is only executed once. MUSE relies on a state automaton that describes the sequence of modes to constrain the shortest path during the preprocessing and the online phase. The support of new sequences of mobility modes only requires the preprocessing of the cliques, independently for each partition. We also augment our algorithm with heuristics during the query phase to achieve further speedups with minimal effect on correctness. We provide experimental results on France’s multimodal network containing the pedestrian, road, bicycle, and public transit networks.

*Key words:* multimodal shortest path; graph separators, route planning; time-dependent graph

## 1. Introduction

The well-being of the economy and even social infrastructure is highly reliant on an efficient transportation network. According to the U.S Department of Transportation (DOT 2015), 40% of all traveled passenger-kilometers consists of commuting. While commuting distances are increasing, average speeds are, however, decreasing due to steadily rising congestion. By 2015, the average commuter typically spent 40 hours stuck in traffic, costing \$121 billion annually. For decades, personal-vehicle travel has been, and remains, the dominating trend: 105 million American commuters rely mostly on driving while the remaining 32 million depend on all other transportation

28 modes. Nevertheless, although accounting for only 5% of the overall commute trips, public transit  
29 is vital to alleviating congestion. The same report indicates that if public transit users in major  
30 metropolitan areas of the united states suddenly reverted to driving, congestion is estimated to  
31 increase by 24% and cost an additional \$17 billion annually. Fortunately, over the past two decades,  
32 public transportation thrived with a registered 25% increase in public transit ridership. Enhanced  
33 access to information significantly improved transit networks and steered urban populations to  
34 consider better alternatives to their private cars.

35 Multimodal algorithms are thereby of increasing importance for route planning to exploit the  
36 diversity of transportation infrastructure fully. A central problem to route planning is the ability to  
37 compute *shortest paths* efficiently. Most commonly, the shortest path denotes that which requires  
38 minimal travel time to reach the destination. In that regard, researchers accomplished significant  
39 progress, and as a result, the literature is abundant (Bast et al. 2016). Early algorithms such as  
40 Dijkstra’s (Dijkstra 1959) solve the *one-to-many* shortest path problem by greedily exploring the  
41 graph data structure that abstracts the transportation network. However, for large networks with  
42 up to several millions of vertices, Dijkstra’s algorithm becomes too slow for practical applications.  
43 Therefore, newer algorithms tend to split the problem into two parts (i) in a first *offline* phase,  
44 known as preprocessing, additional data is extracted from the graph based on some of its unique  
45 features, (ii) in a second *online* phase, known as the query, the preprocessed data is used to speed  
46 up a variant of Dijkstra’s algorithm. Ultimately, efficiency is measured as a tradeoff between query  
47 time, preprocessing time, and memory requirements.

48 The 9<sup>th</sup> Dimacs Challenge (Demetrescu, Goldberg, and Johnson 2009) was a competition on the  
49 shortest path problem and renowned for the significant contribution it brought to the field of route  
50 planning in road networks. Algorithms such as *Arc Flags* (Hilger et al. 2009), *Reach* (Goldberg,  
51 Kaplan, and Werneck 2006), *Transit Node Routing* (Bast, Funke, and Matijevic 2006) and *Highway*  
52 *Hierarchies* (Delling et al. 2006) laid the foundations for modern algorithms.

53 For instance, the commercial solution *TomTom* uses a variant of Arc Flags (Schilling 2012) and  
54 the Open Source Route Machine (OSRM) (Huber and Rust 2016) runs a faster version of Highway  
55 Hierarchies known as *Contraction Hierarchies* (Geisberger et al. 2008). Bast (2009) summarizes the  
56 key *ingredients* behind many algorithms that achieved speedups of several orders of magnitude.  
57 Techniques such as goal direction, contraction, and hierarchy stem from fundamental properties  
58 observed in road networks. For that reason, the same techniques proved to be much less effective  
59 when applied to public transit networks, which are time-dependent and lack hierarchy at the  
60 intracity level (Bast 2009).

61 The *time-expanded* model (Pyrga et al. 2004, 2008) is an event-based graph used to model transit  
62 networks. Mainly, a single station is abstracted with many vertices, one for each event (arrival,  
63 departure, and transfer). Thereby, techniques such as bidirectional search become challenging;  
64 although the target station is known, the specific target vertex is not, and the backward search is  
65 not trivial. Another approach, the *time-dependent* model (Brodal and Jacob 2004) is more compact  
66 and treats each station as a unique vertex in the graph, consequently providing better performance.  
67 Nonetheless, the same speedup techniques applied to road networks are orders of magnitude lower  
68 when applied to transit networks, as Bauer et al. (Bauer, Delling, and Wagner 2011) experimentally  
69 show. Unlike road networks, transit networks are modeled according to timetable information  
70 (Müller-Hannemann et al. 2007). Instead of a graph structure, some of the fastest algorithms  
71 such as *Round-Based Public Transit Routing* (RAPTOR) (Delling, Pajor, and Werneck 2014) and  
72 *Connection Scan Algorithm* (CSA) (Dibbelt et al. 2013) operate directly on the timetable data.

73 The structural differences between road and public transit networks represent the main challenge  
74 for efficient multimodal algorithms. Moreover, modal constraints, which are typically imposed by  
75 users, must be dealt with to compute the shortest path with a valid sequence of transportation  
76 types. We tackle this problem aiming for a practical solution with fast queries (few milliseconds)  
77 for interactive applications and fast preprocessing times (few minutes) to account for traffic, unex-  
78 pected congestion, and transit delays.

79 *Customizable Route Planning* (CRP) (Delling et al. 2011a) is a road network algorithm based  
80 on multilevel separators. Initially, the graph is partitioned, aiming to minimize the number of  
81 boundary vertices. Then, an overlay is constructed by computing full-cliques across all boundary  
82 vertices within each cell in the partition. Finally, during a query, bidirectional-Dijkstra is run on  
83 the query graph combining the source and target cells and the overlay, allowing to skip most of  
84 the vertices in the underlying graph.

85 We detail here MUSE, a MULTimodal SEparators-based algorithm, extending CRP to handle  
86 the multimodal travel computation. By combining label constraints with a profile label correcting  
87 algorithm during preprocessing, we can efficiently handle modal constraints and time-dependency.  
88 Our main contributions are the following:

- 89 • We present a multimodal graph partitioning approach, with a label correcting algorithm to  
90 compute multimodal time-dependent cliques efficiently;
- 91 • We associate the multimodal graph to a labeled automaton to reduce the number of vertices in  
92 the product graph. It reduces the memory footprint and achieves faster preprocessing times;
- 93 • We experimentally evaluate our algorithm on a country-scale network with different heuristics  
94 for faster queries.
- 95 • We provide an open-source tool to construct multimodal networks and a unified dataset to serve  
96 as a benchmark for multimodal route planning algorithms.

97 We first review in section 2 the relevant multimodal route planning algorithms. We then detail  
98 in section 3 the model we used to represent a realistic multimodal network. In section 4, we explain  
99 each stage of our solution. Mainly, we discuss partitioning in section 4.1, the process of computing  
100 a multimodal overlay in section 4.2, and the algorithm and heuristics used to answer queries in  
101 section 4.3. The experimental setup, the results, and associated discussion are available in section  
102 5. Finally, in section 6, we detail our conclusion and possible future work.

## 103 2. Related Work

104 Multimodality involves considering multiple transit networks. The public transit network is typi-  
105 cally modeled with a time-dependent graph (Brodal and Jacob 2004): an edge exists in the graph  
106 when a shuttle connects both vertices. Inversely, the road and other unrestricted networks (pedes-  
107 trian and bicycle) are time-independent: the graph structure remains unchanged for very long  
108 timescales.

### 109 2.1. Road Networks

110 Route planning in road networks involves computing a path that provides the shortest travel time.  
111 However, road congestion is very frequent, and travel duration may vary in time. Thus, Baum  
112 et al. (2016) compute a profile search, to compute the minimum travel time for all departure times  
113 of the day. They compute a multilevel partition of the graph, constructing an overlay. Then, the  
114 approach stores the clique matrix to compute the paths between the different border vertices in  
115 the overlay. The preprocessing has to be re-executed only for the affected cells in the partition to  
116 handle live traffic. MUSE extends this work to deal with multimodal networks.

### 117 2.2. Public Transit networks

118 *Transit* (Antsfeld and Walsh 2012) operates on a graph combining different public transit networks  
119 and evaluates the shortest path based on the associated risk of transfers (probability to miss  
120 one transfer and its impact on travel time) on a combination of transit networks. We believe,  
121 however, that *true* multimodal networks must combine both unrestricted networks such as road  
122 and pedestrian networks and schedule-based networks such as trains, trams, and buses. Otherwise,  
123 a classical time-dependent variant of Dijkstra’s algorithm (Bauer, Delling, and Wagner 2011) would  
124 suffice to solve the problem.

125 Bast, Hertel, and Storandt (2016) exploit the natural hierarchy of local and long-distance trans-  
126 ports by partitioning the graph. Each stop is assigned to a cluster so that the number of inter-cluster

edges is minimized. Then, an existing algorithm is executed to precompute the paths inside each cluster. Only the local transport connections are considered for clustering since long-distance transportation should rather interconnect the clusters. Delling et al. (2017) also partition the transit network to speed up the computation. They consider a hyper-graph where the vertices correspond to routes (bus/subway) and edges to stops. Typically, multiple public transport agencies are sparsely interconnected and would correspond to different partitions. Some routes are then precomputed in the overlay so that the query can only process the source and target clusters and these overlay routes. In MUSE, we also propose to partition the graph, but we consider rather multimodal networks.

Cionini et al. (2017) exploit a Dynamic Timetable Model (DTM): the graph is updated when a delay is observed in the public transit network. In particular, DTM reduces the number of changes to apply to the graph when a timetable has to be modified. The authors only modify the graph structure to model the public transit network. Thus, a classical algorithm such as ALT (Goldberg and Harrelson 2005) may be directly applied to the modified graph to compute the shortest routes.

### 2.3. Multimodal route planning

Multimodal networks are usually modeled with labeled graph structures. Vertices belonging to the same modal network are marked with a similar label, and *link* edges (aka. connecting edges) are added to account for transfers from one mode to another. In that context, the shortest path problem comes with additional constraints that restrict the sequence of transportation during query time. Furthermore, path *feasibility* must be ensured: even if the private car is chosen as a valid means, the path is not feasible if it requires using the car after a bus ride.

Minimizing the travel time may not be sufficient in multimodal networks: the number of transfers/modes or the price may also be considered. Three main approaches exist in the literature (Gianakopoulou, Paraskevopoulos, and Zaroliagis 2019):

1. Scalar approaches: multiple metrics are transformed into a single objective, for instance, with a weighted sum. A transfer may typically be transformed into a time penalty to optimize the travel time uniquely. While such an approach reduces the complexity, combining different metrics is often reductive;
2. Pareto sets: the algorithm keeps all the optimal paths for at least one criteria. This set may be quite large, and some heuristics exist to prune the less interesting elements;
3. Label constrained shortest paths: the graph is explored while respecting a sequence of possible modes. These modes are often represented with a Finite State Automata.

**2.3.1. Scalar approaches:** A few approaches handle only a single objective. Ulloa, Lehoux-Lebacque, and Roulland (2018) propose XTP, an iterative approach to merge the static road network and the public transit network. The shortest path is computed at each iteration from the previously reached bus stops to the next ones. While this approach works well for small networks (*e.g.*, on a metropolitan scale), the complexity becomes intractable for large networks.

The road network subgraph in a multimodal network is usually far bigger (and denser) than the public transit network. *Access-Node Routing* (ANR) (Delling, Pajor, and Wagner 2009) was designed to exploit this property. It borrows ideas from TNR (Bast, Funke, and Matijevic 2006) to perform table lookups on the road network and restrict the search on the public transit network. ANR pre-computes *access-nodes*, a set of vertices in the multimodal graph forming a boundary between the road and the public transit network. For each vertex in the road network, a profile-graph consisting of the shortest paths to all access-nodes is precomputed, requiring significant additional memory. An alternative technique, core-based ANR, hierarchically contracts the road network first. Subsequently, access-nodes are only evaluated for the *core*, *i.e.*, the contracted graph, with significantly fewer vertices.

174 **2.3.2. Pareto sets: User Constrained Contraction Hierarchies (UCCH)** (Dibbelt, Pajor, and  
175 Wagner 2015) is a hierarchical technique originally designed for road networks. The main idea is  
176 to contract while preprocessing the graph vertices ordered by their measured *importance*. For each  
177 contracted vertex, a shortcut edge is inserted between its neighboring vertices and whose weight  
178 preserves the cost of the shortest path containing the contracted vertex. The algorithm solves  
179 queries by running a bidirectional Dijkstra, only scanning vertices with higher importance until  
180 the forward and the backward search meet. To separate modal constraints from preprocessing, the  
181 authors initially split all vertices into different sets based on their labels. They contract each set of  
182 vertices separately to make sure that all shortcut edges have a unique label. Thus, UCCH is scalable  
183 since it handles each mode of transport separately. However, travel times are often unpredictable  
184 in road networks. Thus, a weight change implies that the whole hierarchy has to be recomputed:  
185 routes may become suboptimal. We propose instead in MUSE to recompute the paths only for the  
186 geographical areas where the weight change occurred.

187 Sauer, Wagner, and Zündorf (2020) propose to compute journeys in multimodal networks, with  
188 several bike-sharing operators. The operator-dependent variant introduces a new label identifying  
189 each operator which can only be changed at the specific bike-sharing stations. Handling labels  
190 allow the authors to compute Pareto sets, the bike operator being a criterion. The second variant  
191 expands the graph to execute any public transit algorithm: each operator corresponds to a layer.  
192 A pruning step helps to reduce the preprocessing time when different operators are far from each  
193 other.

194 ULTRA (Baum et al. 2019) constructs the routes in a public transit network, with transfers  
195 using an unrestricted mode (*e.g.*, walking or taxi). More precisely, ULTRA computes shortcuts  
196 in the transfer graph (with the unrestricted mode) to connect pairs of stops in the public transit  
197 network. ULTRA minimizes the number of shortcuts while still guaranteeing that Pareto-optimal  
198 journeys can be extracted. Giannakopoulou, Paraskevopoulos, and Zaroliagis (2019) also adopt a  
199 similar approach, extending the Dynamic Timetable Model: they update timetables when a delay is  
200 observed. The public transit network is modeled with a time expanded graph, and some additional  
201 arcs are inserted to connect two stations through an unrestricted mode (*i.e.*, walking or driving).  
202 The authors group vertices in the graph to accelerate the computation. However, the unrestricted  
203 mode graph is considered static, while congestion may be frequent in urban environments. More-  
204 over, inserting additional arcs may be expensive in large graphs (*e.g.*, for a whole country).

205 **2.3.3. Label constrained: State-Dependent ALT (SDALT)** (Kirchler et al. 2011) adapts  
206 ALT (Goldberg and Harrelson 2005), a popular goal directed algorithm, for multimodal queries.  
207 The intuition is to speed up Dijkstra with a heuristic based on the triangle inequality observed in  
208 transportation networks. In essence, it computes during preprocessing a set of landmark vertices.  
209 Then, for each landmark, it constructs a constrained shortest-path tree to (and from) all other  
210 vertices in the graph. During query time, the preprocessed shortest path costs are used to assign a  
211 potential to each vertex representing the tentative distance to reach the target vertex. It relies on  
212  $D_{RegLC}$  (Barrett, Jacob, and Marathe 2000), a multimodal version of Dijkstra, constrained with  
213 predefined automata. The complexity of the automaton impacts both preprocessing and query  
214 times. The main drawback of this approach corresponds to its scalability: computing landmark  
215 distances becomes too costly for large graph instances.

### 216 3. Model and Assumptions

217 Transportation networks are usually modeled with *graph* structures for their intuitiveness and the  
218 extensive algorithmic toolbox of graph theory (Thomson and Richardson 1995). Mainly, we model  
219 a multimodal network using a labeled directed graph with time-dependent edge costs. It consists  
220 of multiple *layers* of unimodal networks that are interconnected via *link* edges.

221 We detail in sections 3.1 through 3.4 each unimodal network followed by the process of computing  
222 link edges in section 3.5 to obtain the multimodal graph. For better readability, we summarize  
223 recurring notation in table 1. Following are definitions consistently used throughout the paper:

Table 1 keywords and symbol definitions

notation	definition
$G(V, E)$ or $G$	Directed graph with $V$ vertices and $E$ edges.
$G(V, E, \Sigma)$ or $G^\Sigma$	Directed graph, labeled with alphabet set $\Sigma$ .
$E^{road} \subset E$	routes/streets that cars may take
$E^{foot} \subset E$	paths that pedestrians may take
$\mathcal{C}$	set of elementary connections between two public stations (with their time schedule)
$V^{road} \subset V$	crossroads that cars may take
$V^{rentcar} \subset V$	renting stations for shared vehicles
$V^{park} \subset V$	parking spots, where parking a car is authorized
$V^{foot} \subset V$	crossroads that pedestrians may take
$V^{bike} \subset V$	crossroads authorized for bikes
$V^{rentbike} \subset V$	renting stations for shared bikes
$V^{stations} \subset V$	stations of the public transit network
$V^{platform} \subset V$	platforms, each of them attached to a given station
$speed_{walk} \in \mathbb{R}$	average speed of pedestrians
$speed_{bicycle} \in \mathbb{R}$	average speed of bicycles
$c, f, b, p$	modes: <i>car</i> , <i>foot</i> , ( <i>private or rental</i> ) <i>bicycle</i> , and <i>public</i>
$c_r$ and $c_s$	to denote when a car is <i>rented</i> and <i>restored</i>
$b_r$ and $b_s$	to denote when a bike is <i>rented</i> and <i>restored</i>
$(v, w) \in E$	Directed edge from vertex $v$ to $w$ .
$c(v, w)$	The cost, represents travel time [s] of $(v, w)$ .
$c(v, w, \tau)$	Time dependent cost of $(v, w)$ at time $\tau$ .
$len(v, w)$	Physical length [m] of $(v, w)$ .
$lab(v, w)$	Label attached to $(v, w)$ .
$P = \{v_0, v_1, \dots, v_k\}$ or $P_{v_0 v_k}$	A path is an ordered set of vertices $v_i \in V(G)$ .
$c(P, \tau)$	The cost of path $P$ when departing at time $\tau$ .
$d(r, t, \tau)$	The cost of the shortest path $P_{rt}$ departing at $\tau$ .
$word(P) = \cup_{i=0}^{k-1} lab(v_i, v_{i+1})$	The sequence of edge labels associated to path $P$ .
$\mathcal{A} = \{S, \Sigma, \delta, s_0, F\}$	Deterministic Finite Automaton (DFA) consists of a set of states $S$ and alphabet $\Sigma$ , a transition function $\delta$ , an initial state $s_0$ , and a set of final states $F$ . See section 4.2.
$G(V^\times, E^\times)$ or $G^\times = G^\Sigma \times G^{\mathcal{A}}$	Product graph merging graph $G^\Sigma$ and graph automaton $G^{\mathcal{A}}$ .
$\langle v, s \rangle \in V^\times$	Product vertex is a pair of vertex $v \in V(G^\Sigma)$ and a state $s \in S(\mathcal{A})$ .
$(\langle v, s \rangle, \langle w, s' \rangle) \in E^\times$	Product edge, requires a valid transition $s' \in \delta(s, lab(v, w))$ .

224 **A directed Graph**  $G(V, E)$  consists of a set of vertices  $v \in V$ , and directed edges  $(v, w) \in E$   
225 connecting vertices  $v, w \in V$ . A vertex is an abstraction of a physical entity in the network: an  
226 intersection of road segments in the road network or a public transit network station. Throughout  
227 the paper, all graphs are considered directed unless otherwise specified.

228 **The Edge Cost Function**  $c(v, w, \tau)$  (also written  $f_{vw}(\tau)$  in the literature) represents the time  
229 required to travel from vertex  $v$  to vertex  $w$ , when departing at time  $\tau$ . This variable is referred  
230 to as the travel-time. This cost is a periodic positive piece-wise linear function  $f: \Pi \rightarrow \mathbb{R}^+$  where  
231  $\Pi = [0, p] \subset \mathbb{R}$  with a period  $p \in \mathbb{N}$ . If  $f$  is constant, then the edge belongs to a time-independent  
232 network such as the foot or bicycle networks; otherwise, the edge is said to be time-dependent.  
233 In public transportation, the period is typically one week, and the cost function is typically a  
234 non-continuous time function: the travel time increases suddenly when a bus/train leaves a specific

235 station. In road networks, the cost depends on the congestion and is rather a continuous time  
 236 function. However, this cost is unknown a priori since making mid or long-term traffic forecasting  
 237 is still an open problem (Chen et al. 2019).

238 In any case, the *FIFO* property is maintained to ensure polynomial complexity (Kaufman and  
 239 Smith 1993). Also known as the non-overtaking property, it holds that for all  $\tau_1, \tau_2 \in \Pi$  such that  
 240  $\tau_1 \leq \tau_2$  then  $\tau_1 + f(\tau_1) \leq \tau_2 + f(\tau_2)$ . In other words, waiting at a vertex never pays-off.

241 **A Path**  $P = \{v_0, v_1, \dots, v_k\}$ , also written  $P_{v_0 v_k}$ , is an ordered sequence of vertices  $v_i \in V$ . Its asso-  
 242 ciated cost is recursively evaluated with the edge cost of its edges by  $c(v_0, v_1, \tau) + c(\{v_1, \dots, v_k\}, \tau +$   
 243  $c(v_0, v_1, \tau))$  when departing from  $v_0$  at time  $\tau$ . For a query  $q(r, t, \tau)$  with  $r, t \in V$ , our goal is to  
 244 compute the *shortest path*  $P$  with the smallest cost denoted by  $d(r, t, \tau)$ .

### 245 3.1. Road Network (private cars, taxis, and rental vehicles)

246 Structurally, road networks consist of intersecting road segments. Each segment is characterized  
 247 by its length and a specific speed-limit, which can be used to derive the cost function. In its graph  
 248 representation, each edge  $(v, w) \in E^{road}$  represents a road segment, and the vertices  $v, w \in V^{road}$   
 249 mark the junction of two or more road segments.

250 For a realistic model, though, dynamic traffic conditions must be taken into account as speed  
 251 can unpredictably vary. Thus, we rely on speed measurements to construct the cost function. For  
 252 each segment, we collect a set of speed values over a time window  $\Pi$  sampled at a fine-grained rate  
 253  $\Delta t$  (we detail our dataset in the experimental evaluation of section 5). Then, for each edge  $(v, w)$   
 254 we construct its speed profile as a piece-wise linear function  $f_{vw}$ . During query time, we compute  
 255 the cost  $c(v, w, \tau_1)$  of departing from  $v$  at time  $\tau_1$  by evaluating the area under the speed-curve,  
 256 adjusting the arrival time  $\tau_2$  such that:  $\int_{\tau_1}^{\tau_2} f_{vw} dt = len(v, w)$  where  $len(v, w)$  is the length of the  
 257 edge. This is a trivial geometric computation considering the function is piecewise linear. Solving  
 258 the integral for  $\tau_2$ , travel-time is given by  $c(v, w, \tau_1) = \tau_2 - \tau_1$ .

259 Moreover, getting onto or off of a car is only permissible where parking is possible. Hence, for  
 260 each vertex belonging to a road segment where parking is authorized (typically excludes high-  
 261 ways, tunnels, bridges, and sidewalks), we label it as a suitable parking spot  $v \in V^{park} \subset V^{road}$ .  
 262 Furthermore, instead of restricting the road network to private driving only, we consider alter-  
 263 native options, including rental vehicles and on-demand services such as Uber and taxis. While  
 264 on-demand vehicles are typically accessible at every vertex  $v \in V^{park}$  (we discuss the additional  
 265 incurred waiting cost in section 5), rental vehicles are only available at rental stations. Thus, we  
 266 add a vertex  $v \in V^{rentcar} \subset V^{road}$  for each rental station and insert an edge  $(v, w) \in E$  to connect  
 267 the station to the closest junction  $w$  in the road network.

### 268 3.2. Foot Network

269 The foot network is represented by a time-independent graph  $(V^{foot}, E^{foot})$ . Vertices  $V^{foot}$  represent  
 270 junctions and edges  $E^{foot}$  are added for each footpath including sidewalks, bridges, and stairs. The  
 271 cost of an edge  $(v, w)$  is thereby a constant given by  $c(v, w) = len(v, w) / speed_{walk}$  where  $speed_{walk}$   
 272 represents the average pedestrian walking speed.

### 273 3.3. Bicycle Network

274 Similar to a foot network, the bicycle network is based on a time-independent graph  $G(V, E)$  where  
 275 vertices  $V$  represent junctions and edges  $E$  depict either cycling lanes or road segments where  
 276 biking is allowed (typically non-motorway road segments). The cost  $c(v, w) = len(v, w) / speed_{bicycle}$   
 277 is evaluated based on average cycling speed  $speed_{bicycle}$ .

278 Besides private bicycles, bicycle-sharing systems are widespread in urban cities and are efficient  
 279 for fast transfers between nearby public transit stations. Rental bicycles must, however, often be  
 280 picked up and returned at specific locations (bicycle stations). Thus, we add a vertex  $v \in V^{rentbike} \subset V$   
 281 for each rental station and an edge  $(v, w)$  between the station and its closest junction in the bicycle  
 282 network.

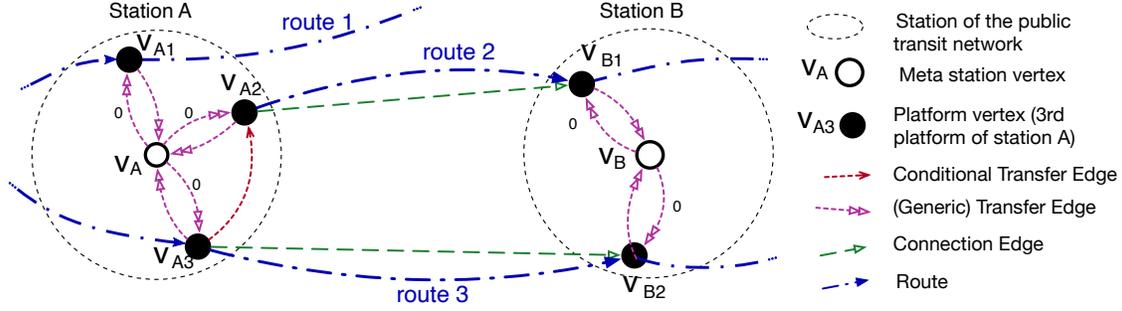


Figure 1 Time-dependent graph representing the public transit network

### 3.4. Public Transit Network

The public transit network is based on a timetable  $\mathcal{T} = (\mathcal{Z}, V^{stations}, \mathcal{C})$  which consists of a set of shuttle vehicles  $\mathcal{Z}$ , a set of stations  $V^{stations}$ , and a set of *elementary connections*  $\mathcal{C}$ . An elementary connection is a 5-tuple  $c = \{z, v_d, v_a, t_d, t_a\} \in \mathcal{C}$  and represents a unique shuttle  $z \in \mathcal{Z}$  departing from a station  $v_d \in V^{stations}$  at time  $t_d$  and arriving at another station  $v_a \in V^{stations}$  at time  $t_a$  without a stop at any intermediate station.

We define a *trip* as a sequence of elementary connections  $\{c_0, c_1, \dots, c_k\}$  fulfilled by a single vehicle such that  $\forall i = 1, 2, \dots, k, c_i = \{z \in \mathcal{Z}, v_d(i) \in V^{stations}, v_a(i) \in V^{stations}, t_d(i), t_a(i)\}$ , where  $v_d(i) = v_a(i-1)$ . A trip typically denotes the itinerary of a single-vehicle scheduled at a particular time. We group multiple trips into a single *route*, which is time-independent and denotes a fixed ordered sequence of stations. To obtain the set of routes  $\mathcal{R}$ , we iterate over all trips and extract for each of them, a route  $r = \{v_d(i)\}_{i \in [0, k]} \cup \{v_a(k)\}$  that we add to  $\mathcal{R}$  iff  $r \notin \mathcal{R}$ . We denote by  $R_{sub}(v) \subset \mathcal{R}$  the subset of routes passing through station  $v \in V^{stations}$ , that is  $\forall r \in R_{sub}(v), v \in r$ .

We have to model the transfers between different trips and routes to construct a realistic time-dependent graph  $G(V, E)$  from the timetable  $\mathcal{T}$ . We proceed as follows:

- $\forall v \in V^{stations}, \forall r \in R_{sub}(v)$ , we add a *platform vertex*  $v_p \in V^{platforms}$ , modeling the *platform* in the station for the corresponding route. Let us denote by  $f_{station} : (V^{stations}, R_{sub}(v)) \rightarrow V^{platforms}$  the bijective function connecting each pair of station and route to its *platform vertex*.
- if we know exactly the transfer time between two platforms,  $\forall v \in V^{stations}$ , we add *conditional transfer edges*  $(v_p, v_{p'})$  between platform vertices  $v_p, v_{p'} \in V^{platforms}$  of the same station, i.e.,  $f_{station}(v_p) = f_{station}(v_{p'})$ . Their cost is the transfer time (walking time) from one platform to another;
- If the transfer time is not known exactly, we insert *generic transfer edges*. We compute the average walking time to go to a central point in the station, denoting a fixed transfer cost  $c(v_p, f_{station}(v_p))$  toward any other platform of the same station  $f_{station}(v_p)$ . As  $c(v_p, f_{station}(v_p))$  includes the full transfer cost,  $\forall v_p' \in V^{platforms} | f_{station}(v_p) = f_{station}(v_p')$ , we add an edge  $(f_{station}(v_p), v_p')$  with  $c(f_{station}(v_p), v_p') = 0$  to preserve the connectivity: we consider that the transfer time has already been considered when the passenger arrives at the station vertex.

These transfer edges are also used when a passenger exits the transit network.

- $\forall c \in \mathcal{C}$ , we add a *connection edge* representing one "edge" for each edge of a route followed by a collection of vehicles. A *connection edge*  $(v_p, v_{p'})$  is inserted between the platform vertices  $v_p \in V^{platforms}$  and  $v_{p'} \in V^{platforms}$  where  $\exists c_i \in \mathcal{R} | v_p = v_d(i)$  and  $v_{p'} = v_a(i)$  (a vehicle leaves  $v_d(i)$  and arrives at  $v_a(i)$ ). Its cost  $c(v_p, v_{p'}, \tau) = t_a(i) - t_d(i)$  is time dependent, and depends on the timetable. It is worth noting that we consider non-multi-edge graphs: a connection edge is inserted between two consecutive platforms of a route. Since different routes correspond to different platforms, we have at most one edge between any pair of platforms (either a (conditional) transfer edge or a connection edge depending on whether the platforms correspond to the same station or not).

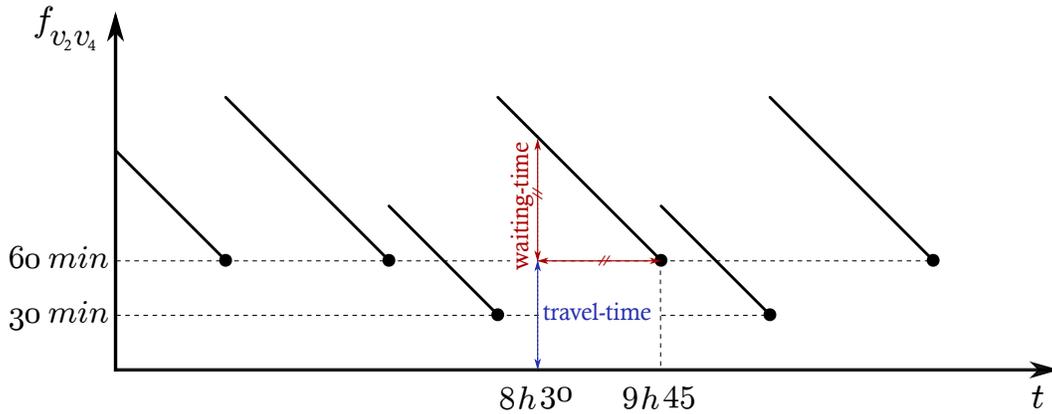


Figure 2 Piecewise linear function describing the cost of edge  $(v_{A2}, v_{B1})$  from figure 1. Six trains are scheduled: four regular trains of 60 min travel-time and two fast trains of 30 min travel-time. Upon arriving to the station at 8:30, the next shuttle is a regular train scheduled to depart at 9:45.

321 The resulting graph is illustrated in figure 1. A hypothetical journey may begin at station  
 322  $v_A$  at  $\tau = 8:30$ . There, we wait for the next train departing from the second platform (vertex  
 323  $v_{A2}$ ) towards station  $v_B$  via route 2, reaching the first platform  $v_{B1}$ . The cost function of edge  
 324  $(v_{A2}, v_{B1})$  denoted  $f_{v_{A2}, v_{B1}}$  is depicted in figure 2. The total cost for the traversal of  $(v_{A2}, v_{B1})$  at  
 325  $t = \tau = 8:30$  is  $c(v_{A2}, v_{B1}, \tau) = 135min$ , comprising both the waiting and the travel time. Because  
 326 the slope  $df_{v_{A2}, v_{B1}}(t)/dt = -1$ , waiting time is given by the elapsed time between arrival and the  
 327 next departure (8:30  $\rightarrow$  9:45). Upon arriving to station  $v_B$  at 10:45, we disembark on the platform  
 328 denoted by vertex  $v_{B2}$  and proceed to transfer to another train traveling along route 3, costing us  
 329  $c(v_{B1}, v_B, \tau + c(v_{A2}, v_{B1}, \tau))$  additional time.

### 3.5. Assembling the Multi-modal Network

331 The multimodal network combines all of the road, foot, bicycle, and public transit networks within  
 332 a single data structure: a labeled directed graph  $G(V, E, \Sigma)$ . To distinguish the networks, we attach  
 333 a unique label  $\sigma \in \Sigma = \{c, f, b, rb, p\}$  to each edge, where  $c, f, b, rb,$  and  $p$  stand for *car, foot,*  
 334 *(private) bicycle, rental bicycles,* and *public* respectively. Let us denote by  $G_\sigma(V_\sigma, E_\sigma)$  the uni-modal  
 335 graph labeled  $\sigma \in \Sigma$ . The vertex-set of the multi-modal graph  $G$  is given by  $V = \cup_{\sigma \in \Sigma} V_\sigma$  and its  
 336 edge-set  $E = \cup_{\sigma \in \Sigma} E^\sigma \cup E_{link}$  where  $E_{link}$  contains the set of *link* edges allowing modal changes in  
 337  $G$ , such as taking a bus after a short walk, by linking the different uni-modal networks together.

338 Intuitively, any transition from one network to another should be mediated via the foot network,  
 339 as *some* walking is usually required for any modal change. Depending on the network however,  
 340 transitions to and from the foot network are only allowed at specific locations, and thus, we select  
 341 a subset of *link* vertices  $V_{link}^\sigma \subseteq V^\sigma$  from each graph for the mode  $\sigma$ :

342 **Foot  $\leftrightarrow$  Road:** The road network is accessible everywhere a car is allowed to park. Furthermore,  
 343 rental vehicles are accessible at rental stations. Thus, all parking spots and rental stations are  
 344 marked as link vertices  $V_{link}^{road} = V^{park} \cup V^{rentcar}$ .

345 **Foot  $\leftarrow$  Private Bicycle:** Considering that bicycles can be used almost everywhere walking is  
 346 possible (except on stairways, for instance), every vertex  $v \in V^{bike}$  is a link vertex from which we  
 347 can access the foot network. Thus,  $V_{link}^{bike} = V^{bike}$ .

348 **Foot  $\leftrightarrow$  Rental Bicycle:** when renting a bicycle, we can 'enter' and 'exit' the bicycle mode only  
 349 when entering a rental station. Thus,  $V_{link}^{rentbike} = V^{rentbike}$ , where  $V^{rentbike}$  represents the rental  
 350 stations only.

351 We assume here that a bicycle can be returned to any rental station. If different companies are  
 352 colocated, we should force the user to return its bike to the right rental station. To support such a

feature, we would need to use a different label for each company, similarly to (Sauer, Wagner, and Zündorf 2020). The automaton would, in that case, introduce a new state denoting the company id to force to return the bicycle to the right location.

**Foot  $\leftrightarrow$  Public Transit:** Public transit stations are accessible via station vertices. Hence,  $V_{\text{link}}^{\text{stations}} = V^{\text{stations}}$ .

Then, for each vertex  $v \in V_{\text{link}}^{\sigma \in \Sigma}$ , we must compute the closest vertex  $w \in V^{\text{foot}}$  in the foot network and add the link edges  $(v, w), (w, v) \in E_{\text{link}}$ . Additional labels are added to  $\Sigma$  as we label link edges according to the type of transfer:

- link edges  $(v, w), (w, v) | v \in V^{\text{foot}} \wedge w \in V^{\text{rentbike}}$  are labeled with  $\text{lab}(v, w) = b_r$  and  $\text{lab}(w, v) = b_s$  which imply *renting* and *restoring* the bike respectively.
- link edges  $(v, w) | v \in V^{\text{bike}}$  and  $w \in V^{\text{foot}}$  are labeled with  $\text{lab}(v, w) = b$ . Indeed, we can only transition from the bike to the foot network (i.e., the bike is only available at the point of departure).
- link edges  $(v, w), (w, v) | v \in V^{\text{foot}}$  and  $w \in V^{\text{rentcar}}$  denote transfers to car-rental stations, and are labeled with the labels  $c_r$  and  $c_s$  similarly to the bike-rental case.

Relying on a brute force approach to compute link edges is costly: we have to scan the whole foot network to identify the closest foot-vertex for each vertex in the bicycle network, leading to a quadratic complexity of  $\mathcal{O}(V_b \times V_f)$ . A better approach relies on clustering the foot vertices using a 2d-tree (Bhatia et al. 2010) based on latitude and longitude, which is a suitable data structure for solving the *nearest neighbor* problem in logarithmic time.

It is worth noting that this preprocessing is executed once, even if the traffic congestion evolves later: multimodal links are not time-dependent. Thus, the cost of a link edge  $c(v, w)$  is fixed and depends on the transfer type. It includes the required walking-time to transfer to or from the foot network and an additional cost to consider either parking-time, processing at a rental station, or for instance, the time it takes to secure a bicycle. Nonetheless, we must also ensure path *feasibility*. That is, if the private car (resp. bicycle) is left behind at some point during the trip in favor of using the bus, we would not be able to use our private car (resp. bicycle) again. Similarly, a scenario in which the private car is used after taking a train is not valid. However, the road network remains accessible via other means such as a taxi or from a rental station. Such constraints are not embedded within the graph but instead dealt with using an automaton, as detailed in the upcoming section.

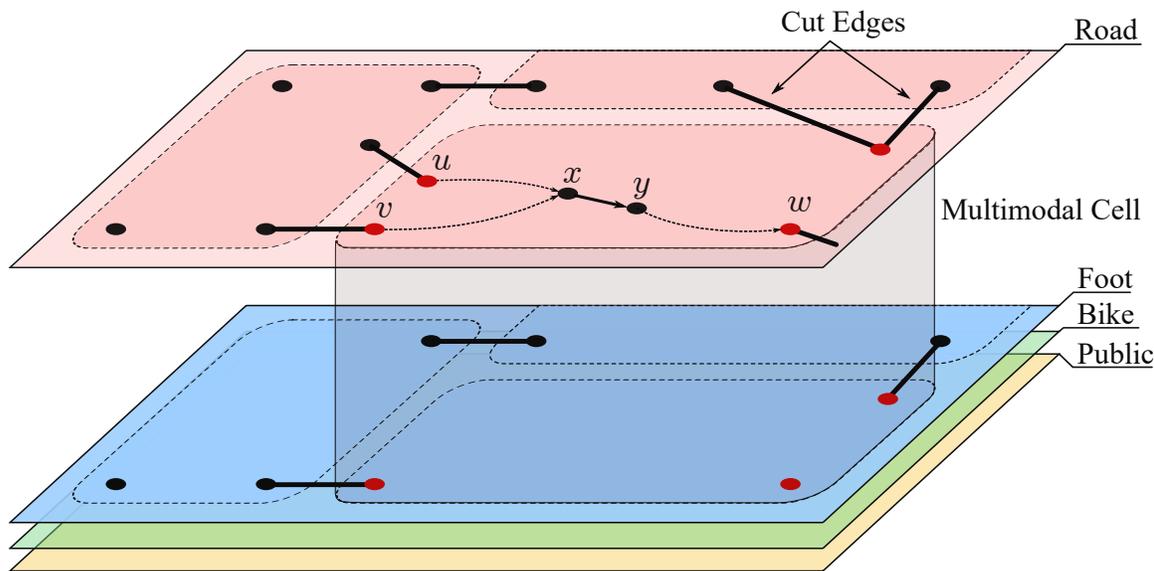
We may extend our multimodal graph to handle passengers carrying their private bikes on public transport. In that case, we would accept a direct transition between the bicycle and public networks. The link vertices correspond to the subway/bus stations that accept bicycles. The label associated with the link edges are different when transitioning from bicycles to the public network, and inversely. That way, we can force in the automaton the sequence *bicycle*  $\rightarrow$  *public*  $\rightarrow$  *bicycle* (cf. section 4.3).

## 4. MUSE: The Algorithm

We now detail MUSE, a speedup technique to *Regular Language Constrained Dijkstra* ( $D_{\text{RegLC}}$ ) based on graph separators to solve the multimodal shortest path problem. By dividing the multimodal graph  $G^\Sigma$  into multiple smaller regions, we can precompute an overlay graph  $H$  significantly smaller than  $G^\Sigma$  and use it to compute queries much faster.

MUSE runs in three stages:

1. **partitioning the graph:** we partition the graph  $G^\Sigma$  to split it into  $k$  cells  $\{C_0, C_1, \dots, C_k\}$  so that we minimize the number of boundary vertices (section 4.1). Regardless of the time-dependent nature of  $G^\Sigma$ , partitioning is run only once, as it only depends on the topology of the graph.
2. **computing the overlay:** we compute for each cell a *full clique*, essentially adding a directed edge for each pair of boundary vertices (section 4.2). It is worth noting that the full clique has to



**Figure 3** Multimodal graph partition with cut edges shown in bold segments. The multimodal cell spans across all unimodal networks. The vertices  $u, v$ , and  $w$  are boundary vertices and the shortest paths  $P_{uv}$  and  $P_{vw}$  share the subpath  $\{x, y, ..w\}$ .

401 be recomputed when the weights change. The overlay graph consists of all cliques and *cut edges*,  
 402 i.e., edges connecting two different cells. The cliques are constrained by an automaton that  
 403 defines the sequences of modes allowed by the user. Furthermore, considering that the graph  
 404 contains time-dependent edges, the weight of clique edges can also, possibly, be time-dependent.  
 405 Therefore, we run *profile queries* to compute the cost function of each clique edge. Even though  
 406 profile queries are costly, cliques are independent, and the computations are parallelized, allowing  
 407 to compute the whole overlay  $H$  for large instances in a few minutes only.

408 3. **solving queries:** we compute the shortest path  $P = q(r, t, \tau)$  by running  $D_{RegLC}$  on the query  
 409 graph  $G^q = G_r \cup H \cup G_t$  which consists of the subgraphs  $G_r$  and  $G_t$ , induced by the cells  
 410 containing the source and target vertices  $r$  and  $t$ , respectively and the overlay  $H$ . We can improve  
 411 query times with a goal-directed version of the algorithm called MUSE\*. Better speed-ups are  
 412 possible using heuristics (MUSE<sup>SV</sup>) with a tradeoff on correctness (section 4.3).

#### 413 4.1. Stage 1: Partitioning The Graph

414 Planar graphs (i.e., graphs that can be drawn on a flat surface without any intersecting edges)  
 415 belong to a class of sparse graphs with valuable topological properties. Most importantly, planar  
 416 graphs can be partitioned in linear time with small separators (Djidjev 1982).

417 Given an undirected graph  $G(V, E)$ , a partition on  $G$  is a collection  $\{C_1, C_2, \dots, C_k\}$  where each  
 418 element  $C_i \subseteq V$  is referred to as a *cell*. Partitioning breaks apart the vertex-set  $V$  into  $k$  cells  
 419  $V = \cup_{i=1}^k C_i$  with no overlap  $C_i \cap C_j = \emptyset | i \neq j$ . Most importantly, the goal is to compute a partition  
 420 such that the number of boundary vertices is minimized. An edge  $(v, w)$  is a *cut edge* if both  $v$  and  
 421  $w$  are boundary vertices belonging to two different cells.

422 Although not planar (due to overpasses and tunnels), road networks can also be efficiently  
 423 partitioned (Eppstein and Goodrich 2008, Sanders and Schulz 2012). In a multimodal graph, the  
 424 vast majority of the vertices belong to the road network. In fact, most foot and bicycle vertices  
 425 are duplicates of road vertices (think of sidewalks along a road segment, for instance); therefore,  
 426 partitioning remains viable.

427 Computing ideal partitions is, however, NP-hard (Garey and Johnson 2002); thus, we recourse  
 428 to approximations that have exhibited good results for transportation networks. METIS (Karypis

and Kumar 1998) is a Multilevel Graph Partitioning algorithm that runs in three stages. The first, *coarsening* stage, aims to reduce the graph’s size by repeatedly merging neighboring vertices and collapsing edges. Each iteration  $i$ , produces a coarser graph  $G_i$  than its predecessor  $G_{i-1}$ . The second stage, *partitioning*, is run on the coarsest graph. Finally, during the *uncoarsening and refining* stage, the graph is expanded at each iteration  $i$ , and the partition is refined until the graph reaches its initial size.

**Setup:** our objective is to minimize the overall number of boundary vertices, regardless of the number of incoming and outgoing cut edges at each cell. Therefore, we transform the directed labeled graph  $G^\Sigma$  into an undirected graph  $G(V, E)$  without labels. Furthermore, considering that the public transit network is modeled as a time-dependent graph, the partitioning process might break apart the platform vertices of a single station across multiple cells. This would mean that the different platforms of a station would all be border vertices with many transfer edges, which would be suboptimal. To speed-up the partitioning, we contract all the platforms ( $V^{platforms}$ ) attached to the same station into a single vertex  $v_s \in V^{station}$ , corresponding to a station. We also keep an edge  $(v_s, v_{s'}) \in E$  if there exists at least one connection edge between any substation of  $v_s$  and any substation of  $v_{s'}$  (i.e., the two stations are connected by at least one route).

**Coarsening** is an iterative process where pairs of vertices are contracted together to reduce the size of the graph. At each iteration  $i$  we obtain a coarser graph  $G_i(V_i, E_i)$ , so that  $|V_i| < |V_{i-1}|$ . Initially,  $\forall v \in G$  we attribute a size label  $s(v) = 1$ . Contracting two vertices  $v, w$  means replacing them with a new vertex  $x$  such that  $s(x) = s(v) + s(w)$ . To contract multiple vertices simultaneously, a set of edges  $M \subset E_i$  is selected, such that no two edges share the same vertex, also known as a *matching*. Thus, at each iteration, we compute the maximal matching, i.e., a matching such that no additional edge can be added, and contract all vertices  $v, w | (v, w) \in M$ . Coarsening is halted when the size of the graph is sufficiently small. Most importantly, we must ensure that  $|V_i| \geq k$  where  $k$  represents the number of desired cells in the partition.

**Partitioning:** the coarsest graph  $G_c(V_c, E_c)$  can be partitioned using Breadth-First-Search (BFS) starting from a random vertex  $v \in V_c$  and growing a tree  $T \subset V_c$  until  $|T| \simeq 1/2|V_c|$ . To obtain  $k$  partitions, the initial partitions are then recursively partitioned  $\log_2(k)$  times. Kernighan and Lin (Kernighan and Lin 1970) use a heuristic to polish the partitioning by exchanging vertices between cells producing a smaller cut-size.

**Uncoarsening and Refining:** at each iteration  $i$ , a less coarse graph  $G_i$  is obtained by expanding  $G_{i-1}$ . For all vertices  $x \in V_{i-1}$ , we expand  $x$  to obtain vertices  $v, w \in V_i$ , and assign them to the same cell  $x$  belonged to. Kernighan Lin’s heuristic is then run again to refine the partition. After the last iteration, we expand the station vertices to retrieve the public transit network and transform the undirected partitioned graph  $G$  back to the multimodal graph  $G^\Sigma$ .

We obtain a partitioned multimodal graph, as illustrated in figure 3, where each layer corresponds to a specific type of transportation. The cells of the partition are, therefore, multimodal cells with boundary vertices located in different layers.

#### 4.2. Stage 2: Computing The Overlay

The partition produces a set of cells, where each cell  $C_i$  contains a set of boundary vertices  $V_i^b \subset C_i$ . The overlay graph  $H(V^H, E^H)$  of  $G^\Sigma$  has the vertex-set  $V^H = \cup_{i=1}^k V_i^b$  which consists of all boundary vertices. The goal is to compute for each pair of boundary vertices  $v, w \in V_i^b$  belonging to the same cell  $C_i$ , a shortcut edge  $(v, w)$  denoting the shortest path  $P_{vw}$ .

Computing all shortcuts within a cell produces a *clique*. It is worth noting that MUSE is naturally parallelizable. Indeed, each cell is independent, and the computation of the shortest paths for each clique can be executed on a different processor/core. After computing all cliques, the edge-set of the overlay  $E^H = E_{cut} \cup_{i=1}^k (v, w) | v, w \in V_i^b$  consists of all edge cuts  $E_{cut}$  together with the clique edges. This explains why we minimize the number of boundary vertices when partitioning the graph. Computing a clique edge  $(v, w)$  requires, however, solving the two following problems:

- 478 • The shortest path  $P_{vw}$  between  $v$  and  $w$  is multimodal. Therefore, we must solve the Label  
 479 Constrained Shortest Path Problem (Barrett, Jacob, and Marathe 2000) (LCSP) to restrict  
 480 the modal sequence of  $P_{vw}$ .
- 481 • Edges can be time-dependent; hence, the cost of  $P_{vw}$  varies based on departure time. Therefore,  
 482 we must compute the *profile* of  $(v, w)$  denoting its associated cost for all departure times.

483 We solve both problems simultaneously using a label correcting algorithm, which produces a  
 484 Constrained Profile Clique (CPC) for each cell in the partition. We first detail each problem  
 485 separately, then delve into the details of the algorithm.

486 **Solving the LCSP:** Using formal languages, we can define modal constraints using *regular*  
 487 *expressions* to prune prohibited edge transitions while computing clique edges. A regular expression  
 488 defines a rule for combining a set of letters of an alphabet using the *or*, *and*, and *kleene* opera-  
 489 tors. Moreover, any regular expression can be written in the form of a Deterministic finite-state  
 490 automaton (DFA) (Brüggemann-Klein 1993). Formally, an DFA is a 5-tuple  $\mathcal{A} = \{S, \Sigma, \delta, s_0, F\}$   
 491 which consists of a finite number of states  $S$ , an alphabet  $\Sigma$ , a transition function  $\delta : S \times \Sigma \rightarrow 2^S$ ,  
 492 an initial state  $s_0$ , and a set of accepting states  $F \subseteq S$ . Therefore, given a directed labeled graph  
 493  $G^\Sigma$  and an DFA  $\mathcal{A}$ , MUSE solves the LCSP by computing a path in  $G^\Sigma$  such that (i) its cost  
 494 is minimum and (ii) the *word* obtained by concatenating its edge labels (sequence of modes) is  
 495 accepted by  $\mathcal{A}$  (*i.e.*, it corresponds to an acceptable sequence of modes).

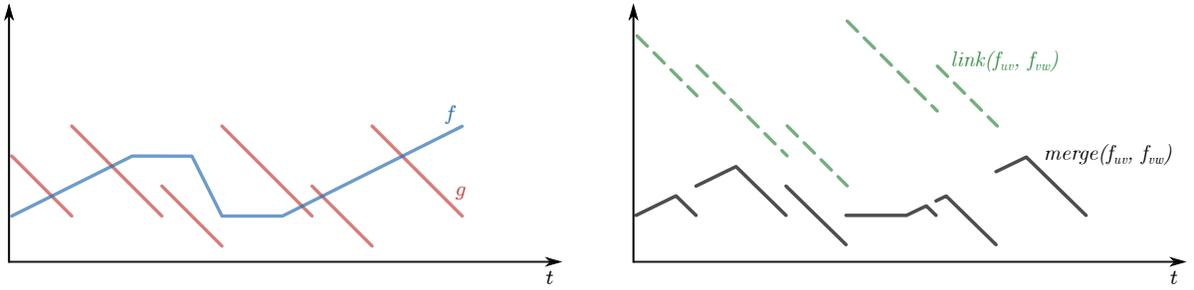
496 Conveniently,  $\mathcal{A}$  can be implemented as a directed labeled graph  $G^{\mathcal{A}}$ . Thus, we can solve the  
 497 LCSP with  $D_{RegLC}$  (Barrett, Jacob, and Marathe 2000), a regular language constrained Dijkstra  
 498 algorithm, deployed on the *product graph*  $G(V^\times, E^\times)$  where  $V^\times$  are vertices and  $E^\times$  edges. The  
 499 product graph  $G^\times = G^\Sigma \times G^{\mathcal{A}}$  is a composition of the underlying graph  $G^\Sigma$  and the automaton  
 500 graph  $G^{\mathcal{A}}$ . A product vertex  $\langle v, s \rangle \in V^\times$  is a pair of a vertex  $v \in V(G^\Sigma)$  and a state  $s \in S(\mathcal{A})$ .  
 501 A product edge  $(\langle v, s_i \rangle, \langle w, s_j \rangle) \in E^\times$  is added *iff* there exists an edge  $(v, w) \in E(G^\Sigma)$  such that  
 502  $s_i \times lab(v, w) \rightarrow s_j$  is a valid transition of  $\delta \in \mathcal{A}$ .

503 **Computing a shortest path profile ( $D_p$ ):** In this section we use the term *distance label* of a  
 504 vertex  $v$ , written  $l(v)$ , to designate the tentative cost to reach  $v$  from a source vertex  $r$  (not to  
 505 be confused with multimodal labels). In a unimodal time-dependent graph  $G$ , we call  $D_p$  a *label*  
 506 *correcting* version of Dijkstra’s algorithm (Orda and Rom 1990) allowing to compute the *profile*  
 507 function  $d_*(r, t)$  denoting the minimum travel time of the shortest path  $P_{rt}$  for all departure times  
 508  $\tau$ . The algorithm does not necessarily settle a vertex  $v$  after it is extracted from the queue because  
 509 of time-dependency. Rather, it propagates its profile  $d_*(r, v)$  to its neighbors and potentially, re-  
 510 inserts  $v$  into the queue (label correction) if  $d_*(r, v)$  is improved for some departure time  $\tau$ .

511 In that context, performing edge relaxations requires additional operations to manipulate func-  
 512 tions rather than scalars. Therefore, we define the following operations:

- 513 • **evaluate:** given a piece-wise linear function  $f$  and a time  $\tau$ , evaluation returns  $f(\tau)$  in  $\mathcal{O}(\log|f|)$ .
- 514 • **link:** given two edges  $(v, w)$  and  $(w, x)$  with cost functions  $f$  and  $g$ , respectively, linking returns  
 515 the cost function  $f * g = f + g \circ (f + \tau)$  of the path  $(v, w, x)$  with complexity  $\mathcal{O}(|f| + |g|)$ . The  
 516 operator  $\circ$  designates the composition of two functions.
- 517 • **merge:** given two parallel edges  $(v, w)$  and  $(v, w)'$ , with cost functions  $f$  and  $g$ , respectively,  
 518 merging returns the cost function  $h = \min(f, g)$ , with minimal travel time from  $u$  to  $v$  for any  
 519 time  $\tau$ . That is,  $merge(f, g) = \min\{f(\tau), g(\tau) \mid \forall \tau \in \Pi\}$  with complexity  $\mathcal{O}(|f| + |g|)$ .

520 Performing the merge operation, in particular, is computationally expensive because we have  
 521 to compute all segment intersections of the input functions  $f$  and  $g$ . We use a line sweep algo-  
 522 rithm (Bentley and Ottmann 1979) over the breakpoints of  $f$  and  $g$ . Computing intersections is a  
 523 requirement because the merged functions are not necessarily homogeneous; that is, the slope of  
 524 the piecewise linear functions  $f$  and  $g$  can be different, rendering the merge operation non-trivial.  
 525 Figure 4 illustrates the outcome of the link and merge operations. In this example, we are merging  
 526 (and linking) the cost function  $f$  of a road network edge with the cost function  $g$  of a public transit



(a) cost functions  $f_{uv}$  of a road network edge, and (b) the resulting functions after linking or merging  $f_{uv}$  and  $f_{vw}$ .

**Figure 4** Illustrating the outcome of the linking and merging non-homogeneous edge cost functions  $f$  and  $g$ .

edge. As illustrated, the merged function in figure 4b contains breakpoints belonging to either  $f$  or  $g$  (for instance the first and last breakpoints) but also breakpoints corresponding to the segment-intersections of  $f$  and  $g$ . Furthermore, the resulting function has more breakpoints than either  $f$  or  $g$ , and thus, memory requirement becomes a concern after successive merge-operations.

Nonetheless, merging can be avoided if one of the functions dominates the other; that is, if  $f \geq \bar{g}$ , then  $g$  is kept as the result of the merge. Similarly, explicit linking can be avoided if either function is constant. If both  $f$  and  $g$  are constant, linking is a trivial addition. If only one function is constant, we distinguish two cases: if  $f$  is constant, we simply translate all  $g$  segments upwards by the constant. If otherwise,  $g$  is constant, we additionally translate all  $f$  segments left by the constant. Only if  $f$  and  $g$  are non-constant functions, we scan their breakpoints and perform the link operation.

Algorithm 1 details the steps of  $D_p$ , which is very analogous to the classic label setting Dijkstra algorithm with a few adaptations to handle vertex labels which are piecewise linear functions. At each iteration, the vertex  $v$  with the smallest key  $key(v) = \underline{d}_*(r, v)$  (i.e., it has the highest priority) is extracted from the queue (line 4) and all of its neighboring vertices are scanned (line 7). For each outgoing edge  $(v, w)$ , we compute the distance label  $l(w) = \underline{d}_*(r, v) * f_{vw}$  (link operation) and check if it improves  $\underline{d}_*(r, w)$  for some departure time. In that case, we update the distance label, with the merge operation:  $\underline{d}_*(r, w) = merge(\underline{d}_*(r, w), l(w))$  and compute  $key(w) = \underline{d}_*(r, w)$  (line 8-10). If  $w$  is scanned for the first time, we insert it into the queue; otherwise, we just update its key inside the queue (11-15). This whole process is repeated until a vertex  $v$  is extracted from the queue such that  $\underline{d}_*(r, v) > \bar{d}_*(r, t)$  holds (line 5). Therefore, the target's distance label cannot be improved anymore, the algorithm is stopped, and  $\underline{d}_*(r, t)$  is correct. Similarly, for a set of targets  $T$ , the same condition must hold for each  $t \in T$  for the algorithm to stop.

**Constrained Profile Clique Algorithm ( $lcD_{pRegLC}$ ):** Given a cell  $C_i$ , we have to compute the shortest path profiles from each boundary vertex  $v \in V_i^b$  toward all other boundary vertices. Since  $G^\Sigma$  is a labeled graph, we must also ensure that all shortest paths abide by the regular language of automaton  $\mathcal{A}$ . Let us call  $D_{pRegLC}$  the regular language constrained version of  $D_p$ , which computes profiles of constrained shortest paths. Hence, one approach is to run  $D_{pRegLC}$ ,  $|V_i^b| \times |S|$  times, from each product vertex  $\langle v, s \rangle$  where  $v \in V_i^b$  and  $s \in S(\mathcal{A})$ . In other words, the algorithm would compute a path from any boundary vertex: the path is possibly different for any state of the DFA. However, this approach is wasteful in terms of memory usage since all pairs of vertex  $v$  and state  $s$  of the DFA are always possible.

Consider  $\mathcal{A}_5$ , the automaton shown in figure 5e. A vertex  $v$  whose label  $lab(v) = c$  (corresponds to the road network) is not compliant with the state  $s_2$  that corresponds to the public transit network. Besides, this state is only reachable from either itself or  $s_1$  (the foot network). Therefore, we design the automaton such that each state is attributed a unique label  $\{f, c, b, p\}$  designating a specific type of network. Hence, the product vertex  $\langle v, s \rangle$  is *valid* only if  $lab(v) = lab(s)$ . This way, we obtain a smaller set  $V_i^{x_b}$  of boundary vertices in the product graph.

---

**Algorithm 1:**  $D_p$  (Profile Dijkstra)

---

**Input** : graph  $G(V, E)$ , source vertex  $r$ , set of target vertices  $T$

**Output:** profile function  $d_*(r, t)$  for all  $t \in T$

```

1  $d_*(r, v) \equiv \infty \mid \forall v \in V$  and  $scanned(v) = \text{false}$ ;
2  $d_*(r, r) \equiv 0$ , set  $key(r) = 0$  and insert  $r$  into priority queue  $Q$ ;
3 while  $Q$  not empty do
4     extract vertex  $v$  with smallest key from  $Q$ ;
5     if  $\underline{d}_*(r, v) > \overline{d}_*(r, t) \forall t \in T$  then // stop when  $v$  cannot improve any target  $t \in T$ 
6         return
7     foreach outgoing edge  $(v, w)$  do
8         if  $d_*(r, v) * f_{vw} < d_*(r, w)$  then //  $(u, w)$  yields an improvement for some departure  $\tau$ 
9              $d_*(r, w) = \text{merge}(d_*(r, w), d_*(r, v) * f_{vw})$ ; // update profile of  $w$ 
10             $key(w) = \underline{d}_*(r, v)$ ; // set key to global minimum of the profile function
11            if not scanned $(w)$  then
12                insert  $w$  into  $Q$ ;
13                 $scanned(w) = \text{true}$ ;
14            else
15                update  $key(w)$  inside  $Q$ ;
16        end
17 end

```

---

565 Moreover, instead of running  $D_{pRegLC}$  a number of  $|V_i^{x^b}|$  times (once for each valid product  
 566 boundary vertex), we are better off using a *multiple source* approach to compute the whole clique  
 567 in a single run. To understand the appeal of using a multiple-source approach over multiple calls to  
 568 the single-source shortest path algorithm, consider the diagram of figure 3. Vertices  $u, v$ , and  $w$  are  
 569 boundary vertices belonging to the same cell and the edge  $(x, y)$  is common to both shortest paths  
 570  $P_{uw}$  and  $P_{vw}$ . With a multiple source approach, each vertex is attributed one distance label for each  
 571 source vertex. Hence, when scanning vertex  $y$ , we can update both  $d_*(u, y)$  and  $d_*(v, y)$  at once.  
 572 This approach is favored when the set of source vertices are close to one another, increasing the  
 573 likelihood of overlapping shortest paths. Thereby, particularly effective to compute our constrained  
 574 cliques considering that boundary vertices are tightly packed in practice, outlining the boundary  
 575 of the cell they belong to, as illustrated in figure 6b.

576 We call our multiple-source algorithm  $lcD_{pRegLC}$ , which stands for *label-correcting Regular Lan-*  
 577 *guage Constrained profile Dijkstra*. As shown in Algorithm 2,  $lcD_{pRegLC}$  takes as input the mul-  
 578 timodal graph  $G^\Sigma$ , the automaton  $\mathcal{A}$  and a set of source vertices  $R$ . Since we are manipulating  
 579 product vertices, the distance label of vertex  $\langle v, s' \rangle$  with respect to a source vertex  $\langle r, s \rangle$  is written  
 580  $d_*^{r,s}(v, s')$ .

581 The first block (line 1-5) initializes, for each source vertex  $\langle r, s \rangle$ , its own distance label to  
 582  $d_*^{r,s}(r, s) \equiv 0$  and that of all other vertices  $d_*^{r,s}(v, s') \equiv \infty$  (with respect to the source  $\langle r, s \rangle$ ). Then,  
 583 all source vertices are added to the queue. At each iteration, the vertex  $\langle v, s' \rangle$  with the smallest  
 584 key is extracted from the queue (line 7). Then, for each outgoing product edge  $(\langle v, s' \rangle, \langle w, s'' \rangle)$   
 585 (line 8-9), we check if the distance label of  $\langle w, s'' \rangle$  can be improved with respect to each source  
 586 vertex  $\langle r, s \rangle$ . That is, if  $d_*^{r,s}(v, s') * f_{vw} < d_*^{r,s}(w, s'')$  holds for any departure time  $\tau$ , then the edge  
 587 is relaxed (line 13). In that case, we update the key  $key(w, s'')$  to the smallest  $\underline{d}_*^{r,s}(w, s'')$  among  
 588 all source vertices  $\langle r, s \rangle$  and add  $\langle w, s'' \rangle$  to the queue (line 16-19).

589 To compute the clique of a cell  $C_i$ , we define the set of source vertices as the boundary vertices  
 590  $R = V_i^b$ . We also define a set of targets  $T = V_i^b$  and stop the algorithm when the next extracted  
 591 vertex  $\langle v, s' \rangle$  respects rule  $key(v, s') > \overline{d}_*^{r,s}(t, s'')$  for all targets  $\langle t, s'' \rangle \in T \times S$ . This means that we  
 592 cannot find a shorter path for the most distant boundary vertex.

**Algorithm 2:**  $lcD_{pRegLC}$  (label correcting  $D_{pRegLC}$ )**Input :** Graph  $G(V, E, \Sigma)$ , Automaton  $\mathcal{A}(S, \Sigma, \delta, s_0, F)$ , and Source set  $R \subseteq V$ **Output:** Constrained shortest path profiles  $d_*^{r,s}(v, s') \mid \forall \langle r, s \rangle \in R \times S$  and  $\forall \langle v, s' \rangle \in V \times S$ 

```

1 foreach source  $\langle r, s \rangle \in R \times S$  do
2    $d_*^{r,s}(r, s) \equiv 0$ ; // set distance label of  $\langle r, s \rangle$  to 0
3    $d_*^{r,s}(v, s') \equiv \infty \mid \forall \langle v, s' \rangle \neq \langle r, s \rangle$ ; // set distance label of all other vertices to infinity
4   set  $key(r, s) = 0$  and add  $\langle r, s \rangle$  to priority queue  $Q$ ;
5 end
6 while  $Q$  not empty do
7   extract product vertex  $\langle v, s' \rangle$  with smallest key from  $Q$ ;
8   foreach outgoing edge  $(v, w)$  do
9     foreach transition  $s' \times lab(v, w) \rightarrow s''$  do
10       $updated = false$ ;
11      foreach source  $\langle r, s \rangle$  do // compute distance labels of  $\langle w, s'' \rangle$  for each source  $\langle r, s \rangle$ 
12        if  $d_*^{r,s}(v, s') * f_{vw} < d_*^{r,s}(w, s'')$  then // relax product edge  $(\langle v, s' \rangle, \langle w, s'' \rangle)$ 
13           $d_*^{r,s}(w, s'') = merge(d_*^{r,s}(w, s''), d_*^{r,s}(v, s') * f_{vw})$ ;
14           $updated = true$ ;
15        end
16      if  $updated$  then
17        // set key to the global minimum among all distance labels of  $\langle w, s'' \rangle$ 
18         $key(w, s'') = \min\{d_*^{r,s}(w, s'') \mid \forall \langle r, s \rangle \in R \times S\}$ ;
19        add  $\langle w, s'' \rangle$  to  $Q$  if not in  $Q$ , otherwise update its key;
20      end
21 end

```

**4.3. Stage 3: Solving Queries**

During the online phase, MUSE solves a query  $q(r, t, \tau)$  by running  $D_{RegLC}$  on the query graph  $G^q = G_r \cup H \cup G_t$  which consists of the overlay  $H$  (precomputed in the previous stage) and the subgraphs  $G_r$  and  $G_t$ , induced by the cells containing the source and target vertices  $r$  and  $t$ , respectively. As  $G^q$  is significantly smaller than the original multimodal graph, MUSE computes the shortest path  $P_{rt}(\tau)$  fast enough for interactive queries. Algorithm 3 details the flow of execution.

Nonetheless, knowing the target’s location allows for making informed decisions while searching for the shortest path. Instead of ”blindly” expanding from the source, the intuition behind goal direction is to use heuristics to guide the search by prioritizing vertices closer to the target. We call MUSE\* (pronounced ‘MUSE-star’) the goal directed version of the algorithm (analogous to  $A^*$  (Hart, Nilsson, and Raphael 1968)) which uses Euclidean distances to compute the *potential function*. The potential function  $\pi(v)$  of a vertex  $v$  represents an estimate of the remaining cost to reach the target vertex  $t$ . To guarantee correctness however, the potential function must be admissible, which means that it must underestimate the true cost, that is,  $\pi(v) \leq cost(P_{vt})$  where  $P_{vt}$  denotes the shortest path between  $v$  and  $t$ .

Another variant called MUSE<sup>SV</sup> applies the same strategy but relies on the *Sedgewick-Vitter* heuristic (Sedgewick and Vitter 1986), which trades off correctness for the sake of speed. The potential function is evaluated as  $\pi(v) = \alpha \times d_{Euc}(v, t) / Speed_{max}$  where  $d_{Euc}(v, t)$  denotes the Euclidean distance between  $v$  and the target  $t$ ,  $Speed_{max}$  the highest speed in the network and  $\alpha$  a tuning parameter. Setting  $\alpha > 1$  may potentially overestimate the shortest path’s cost, which impacts correctness but results in a smaller search space overall.

---

**Algorithm 3:** MUSE (query algorithm)

---

**Input** : graph  $G(V, E, \Sigma)$ , partition  $\{C_1, C_2, \dots, C_k\}$ , overlay  $H$ , DFA  $\mathcal{A} = \{S, \Sigma, \delta, s_0, F\}$ ,  
 source  $r \in V$ , target  $t \in V$ , departure time  $\tau$

**Output:** cost of shortest path  $P_{rt}(\tau)$

```

1 Let  $G^q(V^q, E^q) = C_r \cup H \cup C_t$ ; //  $G^q$  is the query graph and  $C_r$  and  $C_t$  are the source and
   target subgraphs induced by their respective cells
2  $d(v, s) = \infty \mid \forall \langle v, s \rangle \in V^q \times S$ ; // set distance label of all vertices to infinity
3  $d(r, s_0) = 0$ ; // set distance label of source vertex to 0
4 add source  $\langle r, s_0 \rangle$  to priority queue  $Q$ 
5 while  $Q$  not empty do
6   extract product vertex  $\langle v, s \rangle$  with smallest key  $d(v, s)$  from  $Q$ ;
7   if  $v = t$  and  $s \in F(\mathcal{A})$  then
8     return  $d(v, s)$ ; //  $v$  is the target and  $s$  is a final state
9   foreach outgoing edge  $(v, w) \in E^q$  do
10    foreach transition  $s \times lab(v, w) \rightarrow s'$  do
11      if  $d(v, s) + f_{vw}(\tau + d(v, s)) < d(w, s')$  then
12         $d(w, s') = d(v, s) + f_{vw}(\tau + d(v, s))$ ; // relax the product edge  $(\langle v, s \rangle, \langle w, s' \rangle)$ 
13         $key(w, s') = d(w, s') + \pi(w)$ ;
14        add  $\langle w, s' \rangle$  to  $Q$  if not in  $Q$ , otherwise update its key;
15      end
16    end
17 end
    
```

---

614 It is worth noting that different modes may be used inside the same cell. Indeed, Algorithm 3 does  
 615 not restrict the number of time a specific re-enters a specific cell. Only the Finite State Automaton  
 616 constrains the exploration and will possibly prevent re-using the same mode several times.

617 **4.4. Handling time-dependent weights and unpredictable evolutions**

618 In multimodal transportation, we cannot assume weights are static: the user decides according to  
 619 the current network conditions. However, dynamic weights do not have the same meaning and are  
 620 not handled similarly for all transportation modes.

621 For **Public Transit Networks**, the weight of each edge is a non-continuous function: the travel  
 622 time increases discretely after a station has been deserved, including the waiting time before the  
 623 next vehicle's arrival. Thus, merging two paths corresponds to merging the corresponding end-  
 624 to-end cost functions (section 4.2). These linking and merging operations are expensive, but the  
 625 weights are predictive and known a priori.

626 For **Foot and Bike Networks**, we can safely neglect the variations. Thus, these modes don't  
 627 impact significantly the computation time.

628 On the contrary, **Road Networks** also have time-dependent weights. However, these travel  
 629 times for each road segment are highly complex to estimate accurately (Yang, Vereshchaka, and  
 630 Dong 2018). Thus, we cannot update the routes with the linking - merging functions. We must  
 631 rather recompute the routes for the road network, taking into account the new weights.

632 Typically, UCCH (Dibbelt, Pajor, and Wagner 2015) considers each mode independently. Thus,  
 633 only the road network mode has to be reprocessed. However, the routes have to be recomputed as  
 634 soon as one of the routes becomes sub-optimal. MUSE has been rather designed to cope efficiently  
 635 with this kind of situation: each cell is independent, and the routes have to be recomputed only  
 636 for the cells which undergo a change.

**Table 2** Graph characteristics. For the road and bicycle networks, the stations column designates the number of rental stations.

Transportation Network	Vertices		edges		stations	
	$G_{idf}$	$G_{fr}$	$G_{idf}$	$G_{fr}$	$G_{idf}$	$G_{fr}$
Foot ( $G_f$ )	519,558	8,048,695	1,363,995	20,157,922	-	-
Road ( $G_r$ )	457,406	7,252,489	1,018,814	16,712,036	249	899
Bicycle ( $G_b$ )	406,711	7,354,519	1,000,591	18,292,964	800	3,324
Public Transit ( $G_p$ )	60,959	255,872	351,551	1,089,911	18,836	251,685
Foot-Road edges ( $E_r^{\text{link}}$ )	-	-	44,270	431,470	-	-
Foot-Bicycle edges ( $E_b^{\text{link}}$ )	-	-	813,422	14,665,980	-	-
Foot-Public edges ( $E_p^{\text{link}}$ )	-	-	37,672	192,886	-	-
Multimodal ( $G^\Sigma$ )	1,444,634	22,911,575	4,630,315	71,543,169	-	-

## 5. Experimental Evaluation

We report in this section the performance evaluation of MUSE and its associated heuristic-based variants, namely MUSE\* and MUSE<sup>SV</sup>. Furthermore, we evaluate State-Dependent ALT (SDALT) (Kirchler et al. 2011) on the same dataset to compare both algorithms. SDALT is a landmark-based approach applied to multimodal networks. Its preprocessing consists of computing all shortest paths in the network from and to a set of landmarks vertices, as described in section 2.3.3.

Our experimental setup is fully available:

1. a tool to transform OSM data into a directed graph with adjacency lists (Falek 2019);
2. our dataset (GTFS, map, NFA, graph overlay) is available online (Falek 2021a). In particular, it can be reused to execute other algorithms in the same conditions;
3. our implementation of our algorithms (Falek 2021b).

### 5.1. Evaluation Setup

We run all algorithms on two graph instances to assess scalability: the Ile-de-France region denoted  $G_{idf}$  and a country-size graph  $G_{fr}$ , representing France. Table 2 summarizes the graph characteristics based on each transportation layer. The graphs were modeled by combining:

1. a topological dataset from **OpenstreetMap** (Mooney, Minghini et al. 2017), an open-access dataset built through the effort of crowd-sensing and accessible via the Geofabric (Geofabrik 2019) online platform. We use it to construct the road, cycling, and pedestrian graphs. It consists of latitude and longitude coordinates denoting roads, parking spots, and even bicycle and car rental service stations.
2. a *General Transit Feed Specification* (**GTFS**) dataset that represents the standard format used to encode public transit schedule information. For the Ile-de-France graph instance, we rely on the *Ile-de-France mobilités* (IdFm) dataset (Ile de France Mobilités 2020), an online platform providing up-to-date data in GTFS format combining train, RER, subway, tramways, and bus networks in the region. For France’s public transit network, we combine four GTFS datasets covering the whole country and obtained via the open-source platform Navitia (Navitia Open Data 2020).

To model common use case multimodal trips, we define several automata (Fig. 5):

$\mathcal{A}_1$  depicts the combination of walking and all types of public transportation.

$\mathcal{A}_2$  depicts trips that rely on the private car only, for the whole trip;

$\mathcal{A}_3$  extends  $\mathcal{A}_1$  with faster transfers, using rental bicycles, mostly available in the city center at specific locations. An additional state  $s_2$  is reachable via link edges labeled  $b_r$  and  $b_s$  to either rent or return the rental bicycle before pursuing the trip;

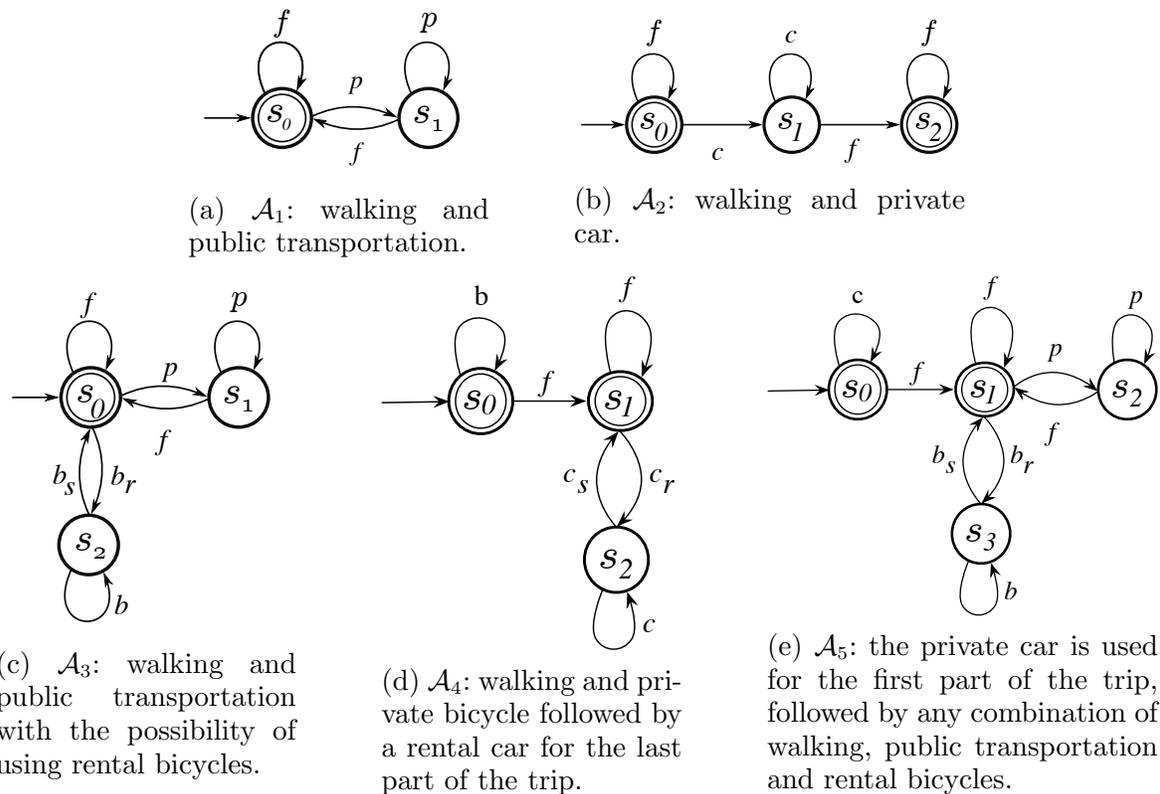


Figure 5 Set of Automata depicting four scenarios used to constrain preprocessing and queries.

670  $\mathcal{A}_4$  excludes the public transit network. It allows using the private bicycle initially followed by a  
 671 rental car for the remaining part of the trip;

672  $\mathcal{A}_5$  combines all means of transportation. The private car is used only initially, followed by any  
 673 combination of walking, public transportation, and rental bicycles.

674 We ran all our experiments on the High-Performance Computing (HPC) of the University of  
 675 Strasbourg. We generated a unique *job* for each partition and automaton that was executed on an  
 676 Intel Haswell node totaling 24 cores and 32 GB of RAM.

## 677 5.2. Preprocessing

678 **5.2.1. MUSE:** Graph partitioning results are reported in table 3. We tested various partitions  
 679 ranging from 100 - 500 cells for  $G_{idf}$  and 800 - 6,000 cells for  $G_{fr}$ . The partition size, i.e., the  
 680 number of cells, impacts both preprocessing and query times. Increasing the number of cells yields  
 681 smaller cells with fewer border vertices but produces a larger overlay graph, which slows the queries.  
 682 In contrast, a smaller partitioning results in large cells, which increases the preprocessing times.  
 683 Consequently, selecting the *adequate* partition is a tradeoff between preprocessing and query times.

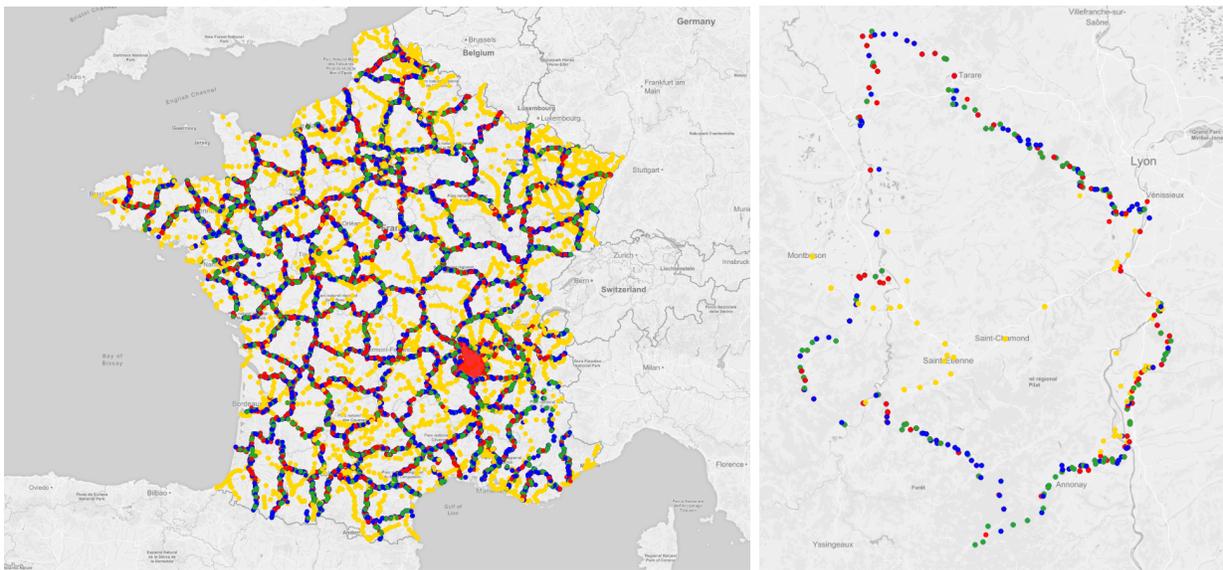
684 Once the graph is partitioned, the cells are independent of each other, and therefore preprocessing  
 685 benefits from parallelism. For instance, figure 6a depicts a partitioning of the  $G_{fr}$  graph into 100  
 686 cells. As illustrated, the number of nodes of each cell varies according to the density of the graph.  
 687 Given a cell with  $N$  border product vertices, we evaluate two strategies to compute its clique:

688 **One-to-Many Strategy:** runs a profile regular language constrained Dijkstra  $D_{pRegLC}$  algorithm  
 689 from each border product vertex. The algorithm is run  $N$  times to compute the clique.

690 **Many-to-Many Strategy:** runs  $lcD_{pRegLC}$  only once. It simultaneously computes all the shortest  
 691 path profiles of the clique in a single call.

**Table 3** Computational time for partitioning the graph with different number of cells and the associated number of border vertices.

Graph	Number of cells	# Border vertices				Time [s]
		min	med	max	tot	
$G_{idf}$	100	124	239	428	24,518	1.023
	200	65	180	344	36,975	1.278
	300	52	159	407	49,102	1.474
	400	56	134	263	55,303	1.673
	500	29	122	308	63,081	1.874
$G_{fr}$	800	13	91	317	94,897	16.471
	1,000	13	91	317	94,897	18.351
	2,000	3	67	331	141,104	19.742
	4,000	4	47	191	194,352	24.598
	6,000	2	38	183	239,639	33.366



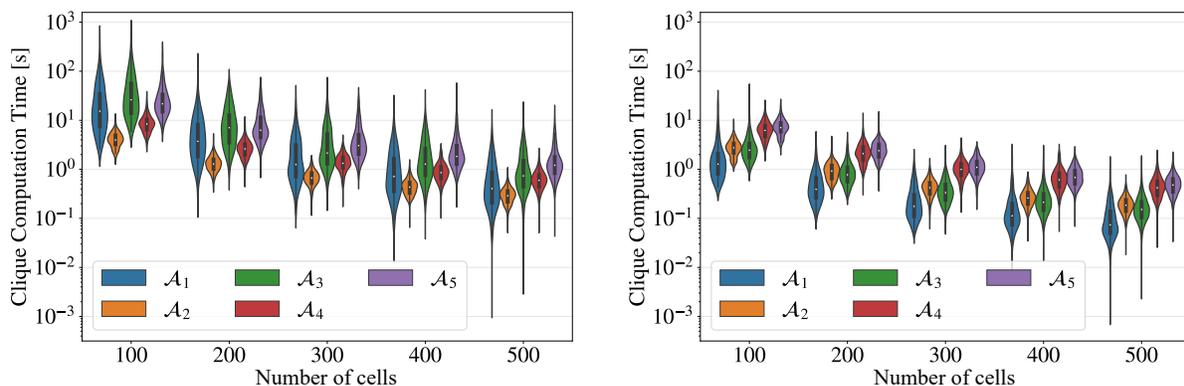
(a) France multimodal network partition (b) boundary vertices of a single cell

**Figure 6** A multimodal partition of France separating the underlying graph into 100 cells. (a) Boundary vertices belonging to the foot, bicycle, road and public transit networks are displayed in blue, green, red, and yellow dots, respectively. (b) boundary vertices of the area highlighted in red in figure (a)

692 Given that the graph's density is not uniform, preprocessing time varies significantly (few mil-  
 693 liseconds to seconds) from one cell to another within the same partition. Using the one-to-many  
 694 strategy, we make the computation in parallel for each border vertex, *i.e.*, we do not need to allo-  
 695 cate all the border vertices of a cell to the same CPU. However, the one-to-many strategy does not  
 696 leverage the fact that the shortest paths of a clique are likely to share common subpaths. Therefore,  
 697 vertices that are shared by several shortest paths are processed several times.

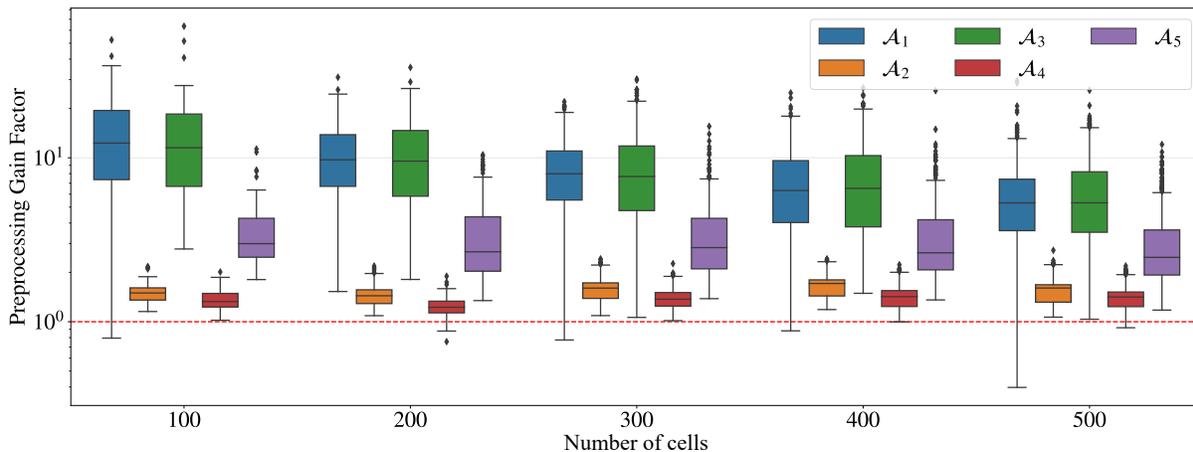
698 In contrast, the many-to-many strategy is more efficient as it simultaneously processes all the  
 699 clique's shortest paths. However, it prohibits load balancing because it allocates a whole cell to a  
 700 single CPU. In terms of memory, the largest overlays we obtained are 289 MB and 3062 MB for  
 701  $G_{idf}$  and  $G_{fr}$  respectively.

702 Figure 7 depicts the clique computation time distribution in the Ile-de-France region using the  
 703 one-to-many (left) and many-to-many (right) strategies for each partition and automaton. As



(a) one-to-many algorithm ( $D_{pRegLC}$ )      (b) many-to-many algorithm ( $lcD_{pRegLC}$ )

**Figure 7** Computational time for the profile cliques of  $G_{idf}$  as a function of number of cells and automaton.



**Figure 8** The preprocessing Gain Factor denoting the speedup of the many-to-many ( $lcD_{pRegLC}$ ) over the one-to-many ( $D_{pRegLC}$ ) strategy for each automaton and number of cells. The horizontal red line corresponds to a gain of 1 and imply no speedup.

704 expected, preprocessing time decreases when the number of cells increases. Furthermore, larger  
 705 automata require more processing time (for instance,  $\mathcal{A}_5$  compared to  $\mathcal{A}_1$ ) because it increases the  
 706 size of the product graph in addition to computing one shortest path for each available state. An  
 707 interesting result is that  $lcD_{pRegLC}$  is an order of magnitude faster compared to running  $D_{pRegLC}$   
 708 multiple times. This confirms our previous hypothesis and shows that within a cell, there is indeed  
 709 a significant number of shortest paths of the clique that do share common subpaths.

710 For a clearer comparison, we compute the *Preprocessing Gain Factor* (PGF) for each clique,  
 711 denoting the ratio of computation time using the one-to-many strategy to that of using the many-  
 712 to-many strategy. The results are shown in figure 8 and suggest that the gain depends mainly on  
 713 the type of automaton used to constrain the preprocessing. The gain is smaller on automata that  
 714 rely, partially, on using the road network ( $\mathcal{A}_5$ ) and smallest for automata that exclude the public  
 715 transit network ( $\mathcal{A}_2$  and  $\mathcal{A}_4$ ). In our setup, the road and public transit networks are the dominating  
 716 means of transportation; that is, if all means are allowed, the shortest path between two locations is  
 717 comprised, mostly, of the road network and public transit edges (fastest alternatives). Furthermore,  
 718 since the road network is much denser and larger than the public transit network, the shortest

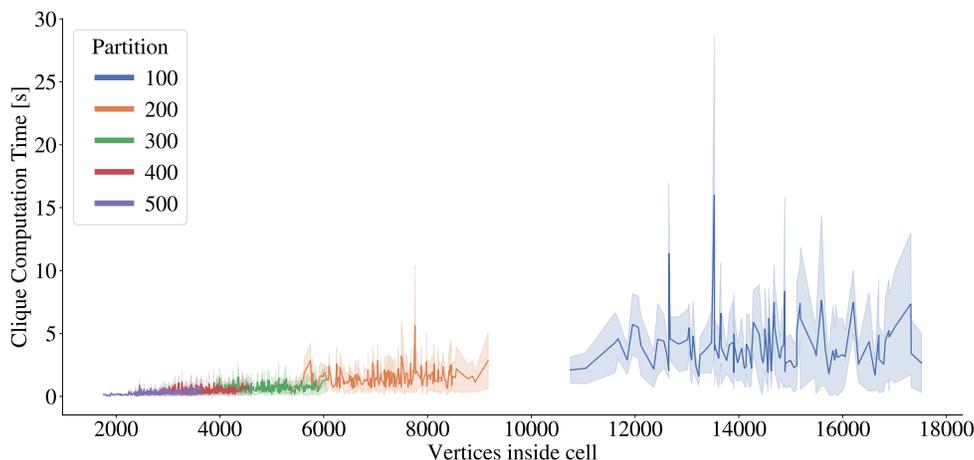


Figure 9 MUSE clique computation time for Ile-de-France based on the number of vertices contained in each cell of the partition.



Figure 10 Distribution of 32 landmarks on the France graph based on the Avoid algorithm (Goldberg and Harrelson 2005).

719 paths that make up a clique overlap less when constrained by the road network compared to the  
720 public transit network. Hence, the gain factor is significant for automata  $\mathcal{A}_1$  and  $\mathcal{A}_3$ .

721 For a better insight, figure 9 shows the preprocessing time (using the many-to-many strategy)  
722 of each cell in the partition as a function of the cell size, i.e., the number of vertices inside the cell.  
723 Even for the worst-case scenario, the computation time is less than 30 sec and corresponds to a  
724 large cell of the 100-cell partition with 14000 vertices. The overall processing time is usually less  
725 than 5 min, even without considering the benefit of parallel computing.

726 For the France region, we experimentally selected the partitions such that the cell sizes are  
727 equivalent to those of the Ile-de-France region, and therefore, preprocessing times are similar when  
728 one CPU is dedicated to each cell in both cases.

729 **5.2.2. SDALT:** The multimodal goal-directed algorithm known as SDALT improves upon the  
730 A\* algorithm heuristics to achieve fast queries. SDALT relies on the *triangle inequality* principle  
731 (Goldberg and Harrelson 2005) to produce better bounds than Euclidean distances. To do so, it  
732 meticulously precomputes a small set of vertices  $L$  called *landmarks* spread across the network.  
733 Then, for each landmark  $l \in L$ , it precomputes a forward and reverse *constrained shortest path tree*  
734 rooted at  $l$ . We implement SDALT as described in (Kirchler et al. 2011). Preprocessing consists of  
735 computing the landmark-set and their associated shortest-path trees.

**Table 4** Preprocessing space and time of SDALT (computing landmark distances) based on 32 landmarks for each automaton in the Ile-de-France region.

	$\mathcal{A}_1$	$\mathcal{A}_2$	$\mathcal{A}_3$	$\mathcal{A}_4$	$\mathcal{A}_5$
space [MB]	634	450	611	438	423
runtime [s]	64.92	83.54	79.92	108.04	106.07

**Landmarks selection:** we select 32 landmarks using the *avoid* algorithm (Goldberg and Harrelson 2005). The first landmark  $l_0 \in L$  is chosen uniformly at random from the vertex-set  $V_f$  of the foot network. We compute all shortest paths from  $l_0$  to every other vertex  $v \in V_f$ . Note that the reverse shortest paths have the same cost as the foot network is bidirectional. Then, we pick a random vertex  $r \in V_f$  and compute the shortest path tree rooted at  $r$ , denoted  $T_r$ . For every other vertex  $v \in V_f$ , we compute its *weight* given by  $weight(v) = d(r, v) - \max\{d(l, v) \mid \forall l \in L\}$ . The weight of a vertex  $v$  consequently represents the error in the distance estimation for  $r$  when using the set of landmarks  $L$ , including  $v$ .

To select a new landmark, we compute the *size* of every vertex  $v \in V_f$ . First we define the subtree  $T_v \subset T_r$  rooted at  $v$ . In particular, we set  $size(v) = 0$  if  $T_v$  contains a landmark  $l \in L$ :  $v$  is useless since another landmark exists that is already well located. Otherwise,  $size(v)$  is evaluated as the sum of all the weights of the vertices in  $T_v$ . Finally, we select the vertex  $w$  greedily with the largest size, and we walk in its subtree  $T_w$ , traversing recursively, the child with the largest subtree until a leaf  $l$  is reached. That leaf  $l$  is then added to the set of landmarks  $L$ . This strategy produces good coverage landmarks that yield good lower bounds. This preprocessing phase is run only once.

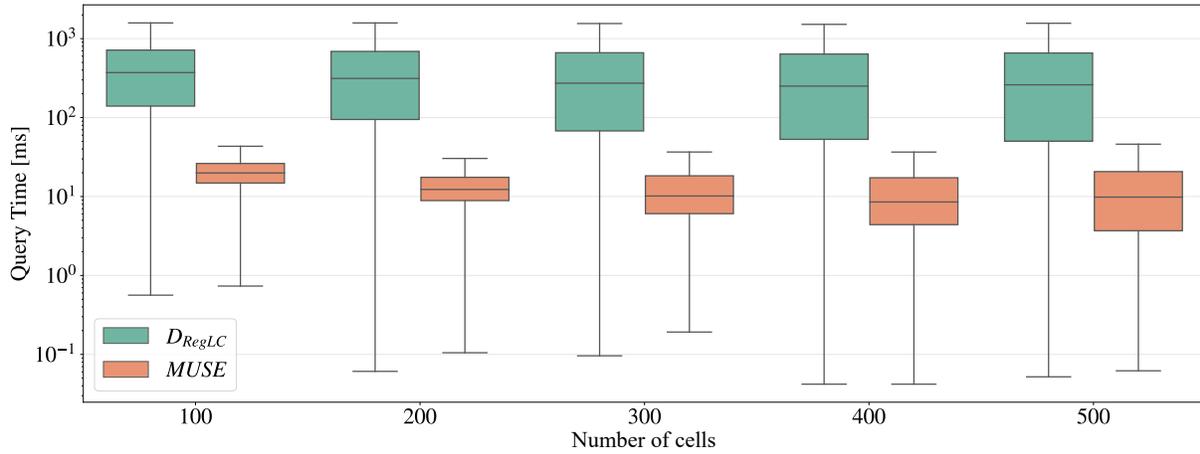
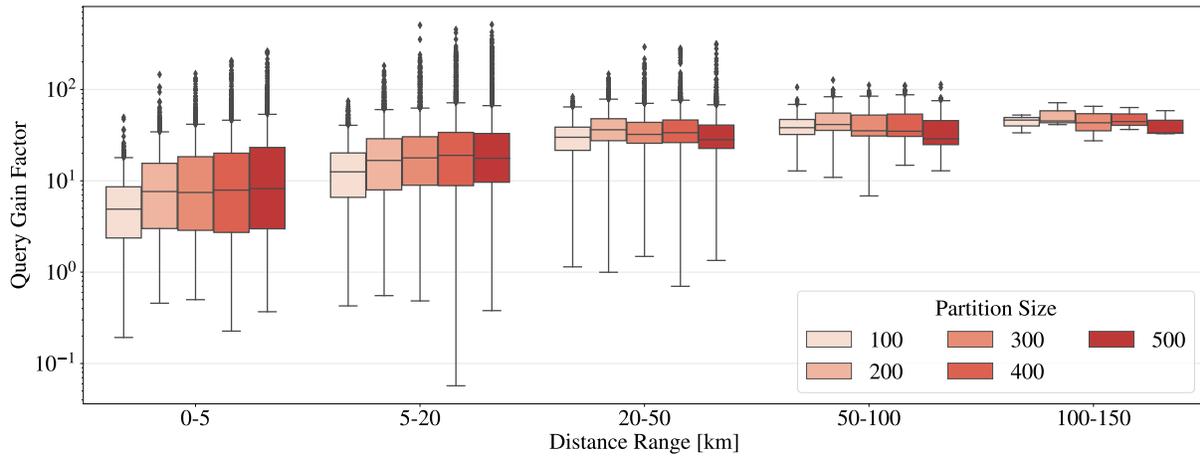
**Landmark Distances:** From each landmark  $l \in L$ , we run  $D_{RegLC}$  on the regular and the reverse multimodal graphs (we reverse the graph by inverting the direction of all its edges). Therefore, we obtain  $|L|$  forward and backward constrained shortest-path trees rooted at each landmark  $l$ , referred to as the *landmark distances*. In the original SDALT implementation (Kirchler et al. 2011), the authors defined four strategies to precompute landmark distances: standard, basic, advanced, and specific. Based on the author’s conclusions, the advanced strategy provides only mild speedups compared to the basic strategy and requires more memory as the number of calculated bounds is a factor  $|L|$  higher in the worst case. Hence, we favor the basic implementation, which uses  $D_{RegLC}$  in conjunction with a single-state automaton with a loop transition labeled with all the allowed transportation modes.

The computational time for selecting 32 landmarks using the Avoid algorithm was 29 seconds for  $G_{idf}$  and 8.75 minutes for  $G_{fr}$ . The additional preprocessing time and space required to compute landmark distances is summarized in table 4 for  $G_{idf}$ . However, SDALT is not scalable: computing all shortest paths between the 32 landmarks and all the vertices of the France graph proved to be impractical, even with the resources available on the HPC.

### 5.3. Queries

After precomputing an overlay graph with all the partitions and automata, we generate 10,000 random queries  $q(r, t, \tau)$  where  $r, t$  are the source and target vertices respectively, and  $\tau$  is the departure time. We compare MUSE with both  $D_{RegLC}$  (multimodal Regular Language Constrained Dijkstra (Barrett, Jacob, and Marathe 2000)) and SDALT (State-Dependent ALT (Kirchler et al. 2011)) (Table 5). Figure 11a depicts the query times of MUSE compared to  $D_{RegLC}$  on the Ile-de-France region, for all queries, regardless of the constraining automaton. MUSE achieves a speedup of almost two orders of magnitude over  $D_{RegLC}$ .

By organizing the queries according to their range (Euclidean distance between the source and the target), we observe that the speedup increases with the distance range (figure 11b). Indeed, the query time not only depends on the number of cells but also the location of the source and target vertices. Recalling that the query graph  $G^q = G_r \cup H \cup G_t$  combines the subgraphs  $G_r$  and  $G_t$ , induced by the cells containing the source and target vertices  $r$  and  $t$ , respectively and the

(a) Query times distribution of  $D_{RegLC}$  and MUSE for all automata.(b) The Query Gain Factor represents the achieved speedup of MUSE over  $D_{RegLC}$  for different partitions and query distance ranges (Euclidean distance).**Figure 11** Performance comparison of  $D_{RegLC}$  and MUSE on the Ile-de-France region graph.

779 overlay  $H$ , MUSE always explores  $G_r$  and  $G_t$ . Therefore, if the cells are large or if the source and  
 780 target are either located inside the same cell or within adjacent cells, MUSE does not benefit from  
 781 the overlay to skip over the rest of the graph. Hence, MUSE performs best for long-range queries,  
 782 as shown in figure 11b.

783 In contrast, figure 12 shows the query times distribution of all of  $D_{RegLC}$ , SDALT, and MUSE for  
 784 each automaton in the Ile-de-France region. We notice that the speedup of our own implementation  
 785 of SDALT over  $D_{RegLC}$  is very close to the original results published by Kirchler et al. MUSE  
 786 is overall an order of magnitude faster than SDALT and achieves better stability across all the  
 787 automata. By comparing the query times distribution of SDALT constrained by  $\mathcal{A}_1$  and  $\mathcal{A}_5$ , we  
 788 observe that the automaton's complexity has a significant impact on the performance of SDALT.  
 789 In contrast, MUSE is more robust to the modal constraints.

790 MUSE can be augmented using any literature technique that does not require preprocessing  
 791 to achieve further speedups during the query phase. Bi-directional search has been proved to be  
 792 inefficient (and often results in speed-downs) on time-dependent graphs (Nannicini et al. 2008).  
 793 Instead, we relied on goal direction to implement MUSE\* and MUSE<sup>SV</sup> (Table 5), as already

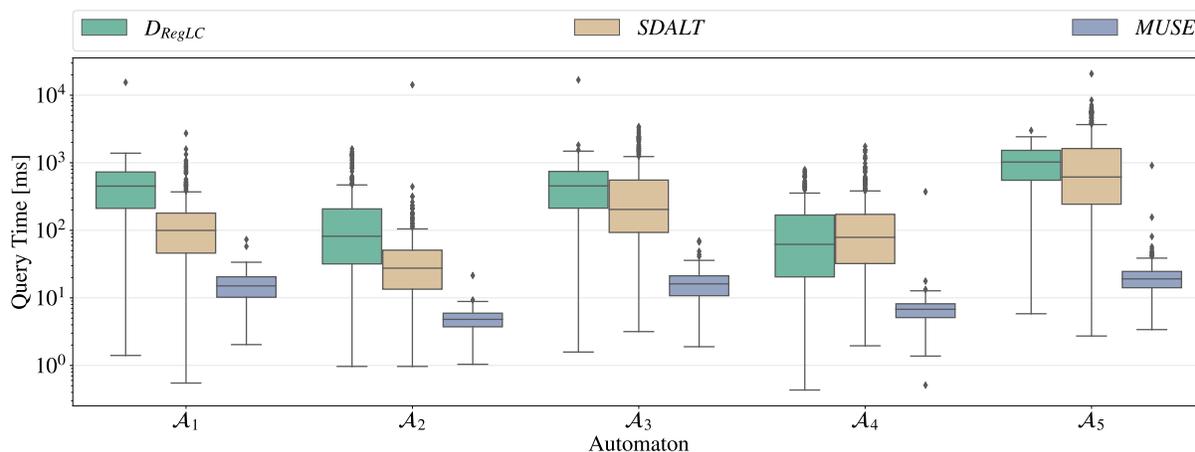


Figure 12 Query times distribution of  $D_{RegLC}$ , SDALT, and the 300-based partition size MUSE for each automaton on the Ile-de-France region graph.

Table 5 List of algorithms used in the query phase of the experimental evaluation.

$D_{RegLC}$	Regular Language Constrained Dijkstra (Barrett, Jacob, and Marathe 2000).
SDALT	State-Dependent ALT (Kirchler et al. 2011).
MUSE	similar to $D_{RegLC}$ but runs on the query graph $G^q = G_r \cup H \cup G_t$ , formed by the overlay $H$ and the source and target induced subgraphs $G_r$ and $G_t$ , respectively.
MUSE*	augments MUSE to obey Goal-Direction principles, similar to $A^*$ (Hart, Nilsson, and Raphael 1968).
MUSE <sup>SV</sup>	similar to MUSE* but implements the Sedgewick-Vitter heuristic (Sedgewick and Vitter 1986) to tradeoff correctness for the sake of speed.

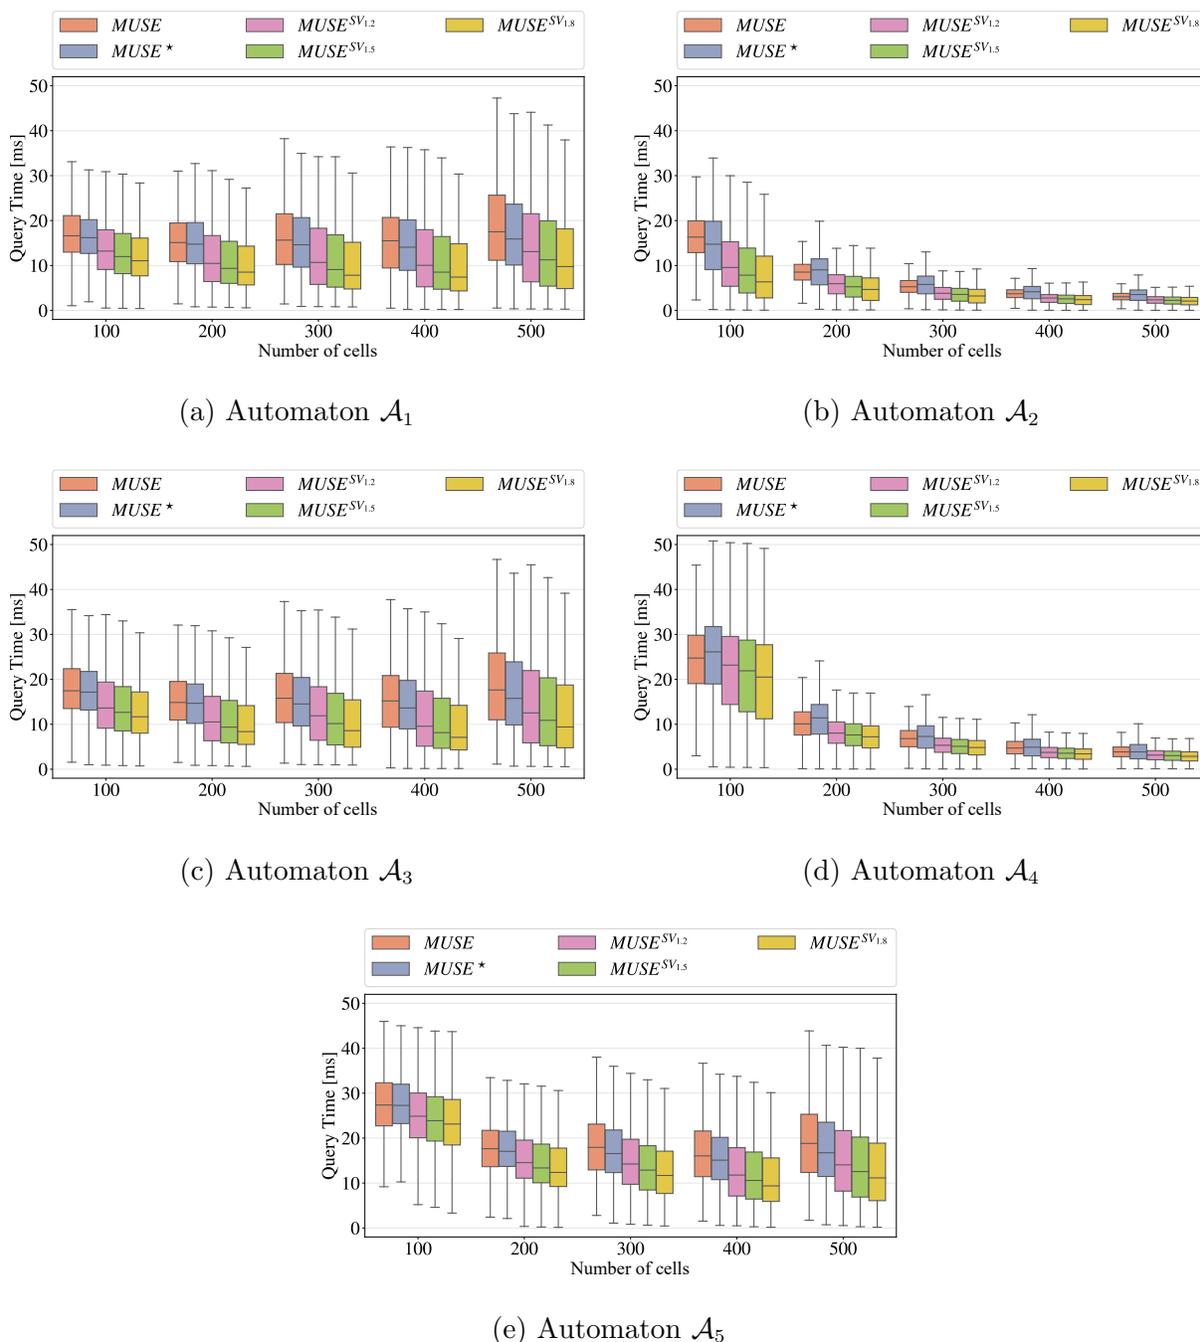
discussed in section 4.3. For the MUSE<sup>SV</sup> variant, we tested three values for the heuristic tuning parameter  $\alpha_1 = 1.2$ ,  $\alpha_2 = 1.5$ , and  $\alpha_3 = 1.8$ . We provide the results in figure 13, depicting the query time distribution of all MUSE variants for each automaton and number of cells.

We notice that query times on  $\mathcal{A}_2$  and  $\mathcal{A}_4$  are faster and keep decreasing when the number of cells increases. In contrast, an optimal number of cells (200 cells) exists for the automata  $\mathcal{A}_1$ ,  $\mathcal{A}_3$ , and  $\mathcal{A}_5$ . Ultimately, query times are bound to increase again beyond a certain threshold of the number of cells. The largest possible partition is, in fact,  $|V|$  (the size of the graph) with each cell containing a single vertex. Hence, the query graph is the whole graph, and MUSE degrades to  $D_{RegLC}$ . Considering that  $\mathcal{A}_2$  and  $\mathcal{A}_4$  exclude the public transit network, the link and merge operations, which are the bottleneck of the algorithm, are trivial, which explains why the queries are significantly faster.

Goal direction does not seem to be a good strategy to accelerate queries in the Ile-de-France region. In fact, Delling et al. (2011a) seem to obtain similar results for small partitions on unimodal road networks using Precomputed Cluster Distances (PCD) (Maue, Sanders, and Matijevic 2010), another goal directed algorithm.

In the France region, we fix the number of cells to 1,000 and provide the query times in figure 14. As query the distance range increases significantly (up to 1,000 km), the computational gain of MUSE over  $D_{RegLC}$  reaches up to 3 orders of magnitude, especially on automata that include the public transit network ( $\mathcal{A}_1$ ,  $\mathcal{A}_3$ , and  $\mathcal{A}_5$ ). Furthermore, on the same automata, MUSE\* provides a significant speedup.

Using the Sedgewick-Vitter heuristic allows us to reduce query times by half in Ile-de-France and up to an order of magnitude in France depending on the automaton and number of cells (Fig. 14).



**Figure 13** Query times distribution for the Ile-de-France region for each automaton of figure 5. Each plot depicts the running times of MUSE, MUSE\*, and MUSE<sup>SV</sup> (with a heuristic factor of 1.2, 1.5, and 1.8) for each number of cell of the preprocessing.

816 This is, however, at the cost of sacrificing the travel time correctness. We report in figure 15 the  
 817 Travel Time Error (in minutes) for each MUSE<sup>SV</sup> variant across all automata. MUSE<sup>SV1.2</sup>, the  
 818 less aggressive overestimating variant, provides paths that are 99% of the time only a few seconds  
 819 longer than the shortest paths. In the worst case, we measure a travel time error of 5 minutes.  
 820 Interestingly, even MUSE<sup>SV1.5</sup> and MUSE<sup>SV1.8</sup> behave well if the automata do not allow using the  
 821 public transit network ( $\mathcal{A}_2$  and  $\mathcal{A}_4$ ). Since the travel time function of public transit graph edges is

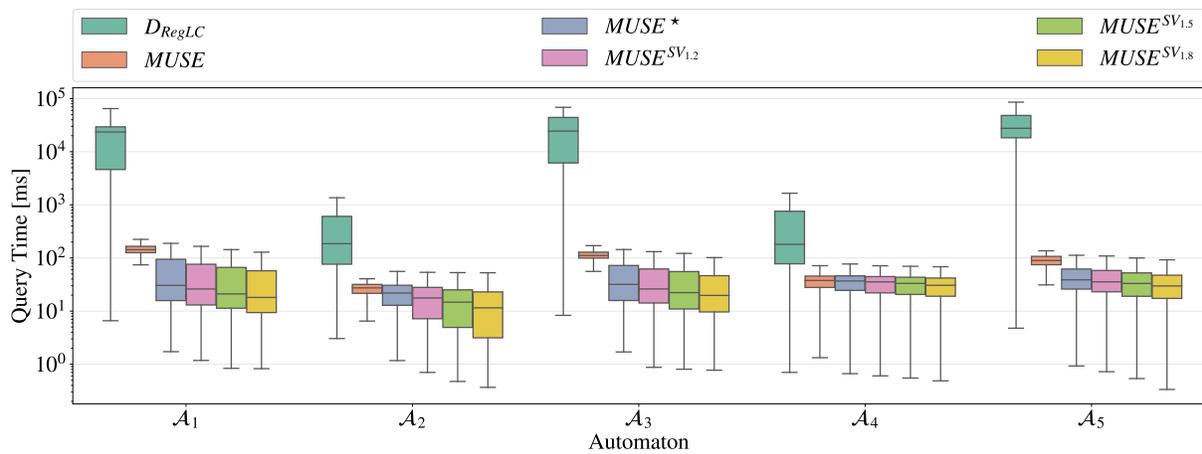
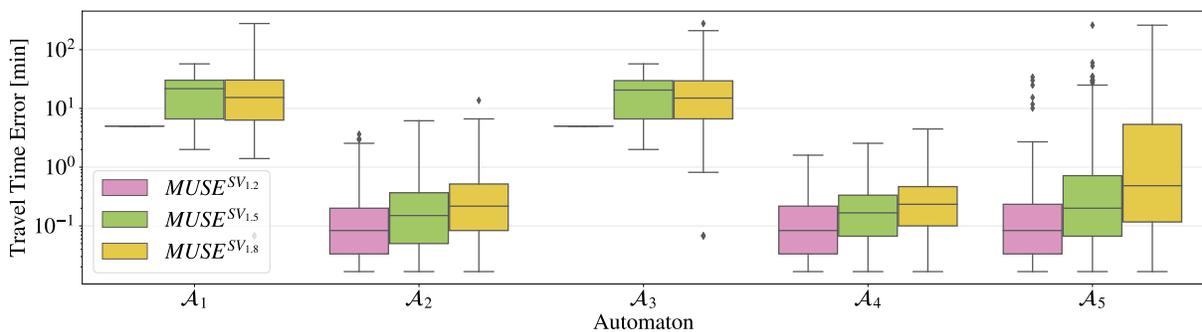
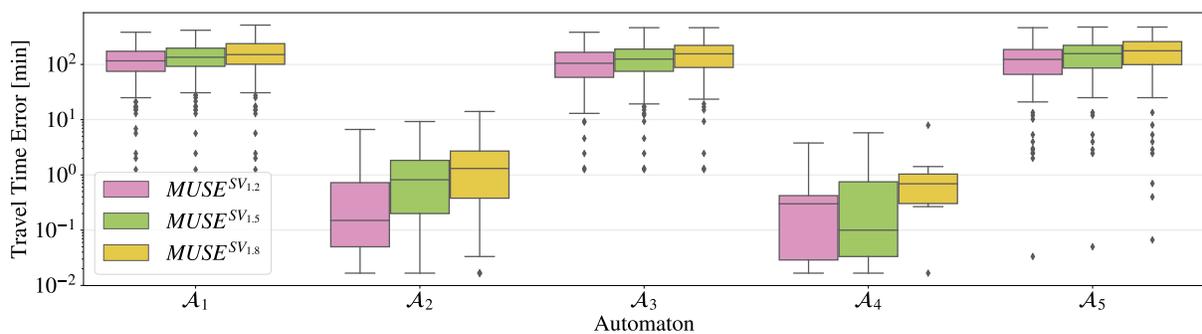


Figure 14 On the France region, we fix the number of cells to 1,000 cells and report the query times distribution of  $D_{RegLC}$  versus MUSE and all of its variants, for each automaton.



(a) Ile-de-France region



(b) France region

Figure 15 Travel Time Error distribution on the Ile-de-France and France regions for each automaton based on the heuristic factor setting of  $MUSE^{SV}$ .

822 non-continuous (figure 2), missing a connection results in a sudden increase of travel time, which  
 823 explains the large errors observed on the remaining automata. Thus, this variant should be carefully  
 824 exploited.

## 6. Conclusion and Future Work

We presented MUSE, a new approach to solve the multimodal shortest path problem via graph partitioning and finite state automaton. Its main advantage is scalability, as it splits large graph instances into several independent cells that can be processed in parallel. We also provide different strategies to improve the performance of both the preprocessing and the query phases. We achieve preprocessing and query times that are fast enough to process a cell and resolve a query in only a few milliseconds. Hence, MUSE is a viable solution for real-time multimodal route planning. We experimentally tested our implementation and achieved a speedup of up to two orders of magnitude compared to  $D_{RegLC}$ , the label constrained version of Dijkstra’s algorithm for multimodal networks. Using goal direction and heuristics, we further accelerate the queries and show that the obtained paths are, most of the time, only a few seconds longer than the shortest paths. We also validated the performance of MUSE against SDALT, a landmark-based goal-directed multimodal algorithm. MUSE is an order of magnitude faster than SDALT and proved more robust to varying modal constraints. Furthermore, SDALT could not be scaled to handle large graph instances such as France’s multimodal transportation network. A problem that MUSE circumvents using graph partitioning.

The partitioning is critical for the performance of both the preprocessing and the query. We first tested our implementation of the PUNCH algorithm (Delling et al. 2011b) but switched later on to METIS, the general graph partitioning algorithm, which was more reliable and provided better results. Therefore, a formal analysis of graph partitioning algorithms tailored to transportation networks is essential to improve our current results. Furthermore, to improve query times, multilevel partitioning for multimodal graphs is an interesting next step. It allows skipping over most of the network using the overlay edges on the partition’s low level while reducing the search space in the source and target cells by exploring the high levels of the partition. During preprocessing, we also need to compute multimodal profile cliques. We observed that the *merge* and *link* operations are the bottleneck of  $lcD_{pRegLC}$  due to a rapidly increasing number of breakpoints in the travel time functions. An interesting work perspective would be to rely on heuristics (Strasser 2017) to simplify these operations and evaluate the tradeoff between computational speed and impact on correctness. Finally, we considered here only a route computation that minimizes the travel time. Handling multiple criteria will require computing the cliques in each cell by maintaining Pareto sets. Then, the different Pareto sets will be explored during the query stage. While we conjecture that keeping a small number of solutions would be sufficient, we have to investigate how MUSE can support multiple criteria with an acceptable computation time and memory.

## References

- Antsfeld L, Walsh T, 2012 *Finding multi-criteria optimal paths in multi-modal public transportation networks using the transit algorithm*. *ITS World Congress*.
- Barrett C, Jacob R, Marathe M, 2000 *Formal-language-constrained path problems*. *SIAM Journal on Computing* 30(3):809–837, URL <http://dx.doi.org/10.1137/S0097539798337716>.
- Bast H, 2009 *Car or public transport—two worlds*. *Efficient Algorithms*, 355–367 (Springer), URL [http://dx.doi.org/10.1007/978-3-642-03456-5\\_24](http://dx.doi.org/10.1007/978-3-642-03456-5_24).
- Bast H, Delling D, Goldberg A, Müller-Hannemann M, Pajor T, Sanders P, Wagner D, Werneck RF, 2016 *Route planning in transportation networks*. *Algorithm engineering*, 19–80 (Springer), URL [http://dx.doi.org/10.1007/978-3-319-49487-6\\_2](http://dx.doi.org/10.1007/978-3-319-49487-6_2).
- Bast H, Funke S, Matijevic D, 2006 *Ultrafast shortest-path queries via transit nodes*. *Ninth DIMACS Implementation Challenge*, 175–192, URL <http://dx.doi.org/10.1090/dimacs/074/07>.
- Bast H, Hertel M, Storandt S, 2016 *Scalable transfer patterns*. *Proceedings of the Meeting on Algorithm Engineering and Experiments (ALENEX)*, 15–29, URL <http://dx.doi.org/10.1137/1.9781611974317.2>.
- Bauer R, Delling D, Wagner D, 2011 *Experimental study of speed up techniques for timetable information systems*. *Networks* 57(1):38–52, URL <http://dx.doi.org/10.1002/net.20382>.

- 875 Baum M, Buchhold V, Sauer J, Wagner D, Zündorf T, 2019 *UnLimited TRAnsfers for Multi-Modal Route*  
876 *Planning: An Efficient Solution. Annual European Symposium on Algorithms (ESA)*, volume 144, 14:1–  
877 14:16, URL <http://dx.doi.org/10.4230/LIPIcs.ESA.2019.14>.
- 878 Baum M, Dibbelt J, Pajor T, Wagner D, 2016 *Dynamic time-dependent route planning in road net-*  
879 *works with user preferences. Experimental Algorithms*, 33–49, URL [http://dx.doi.org/10.1007/](http://dx.doi.org/10.1007/978-3-319-38851-9_3)  
880 [978-3-319-38851-9\\_3](http://dx.doi.org/10.1007/978-3-319-38851-9_3).
- 881 Bentley JL, Ottmann TA, 1979 *Algorithms for reporting and counting geometric intersections. IEEE Trans-*  
882 *actions on computers* C-28(9):643–647, URL <http://dx.doi.org/10.1109/TC.1979.1675432>.
- 883 Bhatia N, et al., 2010 *Survey of nearest neighbor techniques. International Journal of Computer Science and*  
884 *Information Security* 8(2).
- 885 Brodal GS, Jacob R, 2004 *Time-dependent networks as models to achieve fast exact time-table queries. Elec-*  
886 *tronic Notes in Theoretical Computer Science* 92:3–15, URL [http://dx.doi.org/10.1016/j.entcs.](http://dx.doi.org/10.1016/j.entcs.2003.12.019)  
887 [2003.12.019](http://dx.doi.org/10.1016/j.entcs.2003.12.019).
- 888 Brüggemann-Klein A, 1993 *Regular expressions into finite automata. Theoretical Computer Science*  
889 120(2):197–213, URL [http://dx.doi.org/10.1016/0304-3975\(93\)90287-4](http://dx.doi.org/10.1016/0304-3975(93)90287-4).
- 890 Chen Y, Guizani M, Zhang Y, Wang L, Crespi N, Lee GM, Wu T, 2019 *When traffic flow prediction and*  
891 *wireless big data analytics meet. IEEE Network* 33(3):161–167, URL [http://dx.doi.org/10.1109/](http://dx.doi.org/10.1109/MNET.2018.1800134)  
892 [MNET.2018.1800134](http://dx.doi.org/10.1109/MNET.2018.1800134).
- 893 Cionini A, D’Angelo G, D’Emidio M, Frigioni D, Giannakopoulou K, Paraskevopoulos A, Zaroliagis C, 2017  
894 *Engineering graph-based models for dynamic timetable information systems. Journal of Discrete Algo-*  
895 *rithms* 46-47:40 – 58, URL <http://dx.doi.org/https://doi.org/10.1016/j.jda.2017.09.001>.
- 896 Delling D, Dibbelt J, Pajor T, Zündorf T, 2017 *Faster Transit Routing by Hyper Partitioning. D’Angelo G,*  
897 *Dollevoet T, eds., 17th Workshop on Algorithmic Approaches for Transportation Modelling, Optimiza-*  
898 *tion, and Systems (ATMOS 2017)*, volume 59, 8:1–8:14, URL [http://dx.doi.org/10.4230/OASlcs.](http://dx.doi.org/10.4230/OASlcs.ATMOS.2017.8)  
899 [ATMOS.2017.8](http://dx.doi.org/10.4230/OASlcs.ATMOS.2017.8).
- 900 Delling D, Goldberg AV, Pajor T, Werneck RF, 2011a *Customizable route planning. International*  
901 *Symposium on Experimental Algorithms (SEA)*, 376–387, URL [http://dx.doi.org/10.1007/](http://dx.doi.org/10.1007/978-3-642-20662-7_32)  
902 [978-3-642-20662-7\\_32](http://dx.doi.org/10.1007/978-3-642-20662-7_32).
- 903 Delling D, Goldberg AV, Razenshteyn I, Werneck RF, 2011b *Graph partitioning with natural cuts. Inter-*  
904 *national Parallel & Distributed Processing Symposium (IPDPS)*, 1135–1146 (IEEE), URL [http:](http://dx.doi.org/10.1109/IPDPS.2011.108)  
905 [//dx.doi.org/10.1109/IPDPS.2011.108](http://dx.doi.org/10.1109/IPDPS.2011.108).
- 906 Delling D, Pajor T, Wagner D, 2009 *Accelerating multi-modal route planning by access-nodes. European*  
907 *Symposium on Algorithms*, 587–598, URL [http://dx.doi.org/10.1007/978-3-642-04128-0\\_53](http://dx.doi.org/10.1007/978-3-642-04128-0_53).
- 908 Delling D, Pajor T, Werneck RF, 2014 *Round-based public transit routing. Transportation Science* 49(3):591–  
909 604, URL <http://dx.doi.org/10.1287/trsc.2014.0534>.
- 910 Delling D, Sanders P, Schultes D, Wagner D, 2006 *Highway hierarchies star. Ninth DIMACS Implementation*  
911 *Challenge*, 141–174, URL <http://dx.doi.org/10.1090/dimacs/074/06>.
- 912 Demetrescu C, Goldberg AV, Johnson DS, 2009 *The Shortest Path Problem: Ninth Dimacs Implementation*  
913 *Challenge* (American Mathematical Society), URL <http://dx.doi.org/10.1090/dimacs/074>.
- 914 Dibbelt J, Pajor T, Strasser B, Wagner D, 2013 *Intriguingly simple and fast transit routing. Interna-*  
915 *tional Symposium on Experimental Algorithms (SEA)*, 43–54, URL [http://dx.doi.org/10.1007/](http://dx.doi.org/10.1007/978-3-642-38527-8_6)  
916 [978-3-642-38527-8\\_6](http://dx.doi.org/10.1007/978-3-642-38527-8_6).
- 917 Dibbelt J, Pajor T, Wagner D, 2015 *User-constrained multimodal route planning. Journal of Experimental*  
918 *Algorithmics (JEA)* 19, URL <http://dx.doi.org/10.1145/2699886>.
- 919 Dijkstra EW, 1959 *A note on two problems in connexion with graphs. Numerische mathematik* 1(1):269–271,  
920 URL <http://dx.doi.org/10.1007/BF01386390>.
- 921 Djidjev HN, 1982 *On the problem of partitioning planar graphs. SIAM Journal on Algebraic Discrete Methods*  
922 3(2):229–240, URL <http://dx.doi.org/10.1137/0603022>.

- 923 DOT U, 2015 *Beyond traffic 2045: Trends and choices*. Technical Report [https://www.transportation.gov/sites/dot.gov/files/docs/BeyondTraffic\\_tagged\\_508\\_final.pdf](https://www.transportation.gov/sites/dot.gov/files/docs/BeyondTraffic_tagged_508_final.pdf), US Department of Transportation. 924  
925
- 926 Eppstein D, Goodrich MT, 2008 *Studying (non-planar) road networks through an algorithmic lens*. *ACM SIGSPATIAL international conference on Advances in geographic information systems* (ACM), URL <http://dx.doi.org/10.1145/1463434.1463455>. 927  
928
- 929 Falek A, 2019 *osm-to-graph*. <https://github.com/aminefalek/osm-to-graph>.
- 930 Falek A, 2021a *Muse: Multimodal separators for efficient route planning in transportation networks*. Zenodo dataset, URL <http://dx.doi.org/10.5281/zenodo.5276749>. 931
- 932 Falek A, 2021b *Muse: Multimodal separators for efficient route planning in transportation networks*. <https://github.com/aminefalek/muse>. 933
- 934 Garey MR, Johnson DS, 2002 *Computers and intractability*, volume 29 (W. H. Freeman & Co.).
- 935 Geisberger R, Sanders P, Schultes D, Delling D, 2008 *Contraction hierarchies: Faster and simpler hierarchical routing in road networks*. *International Workshop on Experimental and Efficient Algorithms (WEA)*, 319–333, URL [http://dx.doi.org/10.1007/978-3-540-68552-4\\_24](http://dx.doi.org/10.1007/978-3-540-68552-4_24). 936  
937
- 938 Geofabrik O, 2019 *Openstreetmap data extract*. URL <https://download.geofabrik.de/europe/france/ile-de-france.html>. 939
- 940 Giannakopoulou K, Paraskevopoulos A, Zaroliagis C, 2019 *Multimodal dynamic journey-planning*. *Algorithms* 12(10):213, URL <http://dx.doi.org/10.3390/a12100213>. 941
- 942 Goldberg AV, Harrelson C, 2005 *Computing the shortest path: A\* search meets graph theory*. *ACM-SIAM symposium on Discrete algorithms (SODA)*, 156–165. 943
- 944 Goldberg AV, Kaplan H, Werneck RF, 2006 *Reach for a\*: Shortest path algorithms with preprocessing*. *Ninth DIMACS Implementation Challenge*, 93–140, URL <http://dx.doi.org/10.1090/dimacs/074/05>. 945
- 946 Hart PE, Nilsson NJ, Raphael B, 1968 *A formal basis for the heuristic determination of minimum cost paths*. *IEEE Transactions on Systems Science and Cybernetics* 4(2):100–107, URL <http://dx.doi.org/10.1109/TSSC.1968.300136>. 947  
948
- 949 Hilger M, Köhler E, Möhring RH, Schilling H, 2009 *Fast point-to-point shortest path computations with arc-flags*. *Ninth DIMACS Implementation Challenge*, volume 74, 41–72, URL <http://dx.doi.org/10.1090/dimacs/074/03>. 950  
951
- 952 Huber S, Rust C, 2016 *Calculate travel time and distance with openstreetmap data using the open source routing machine (osrm)*. *The Stata Journal* 16(2):416–423, URL <http://dx.doi.org/10.1177/1536867X1601600209>. 953  
954
- 955 Ile de France Mobilités, 2020 *Données offre de transport ile-de-france mobilités au format gtfs*. <https://data.iledefrance-mobilites.fr/explore/dataset/offre-horaires-tc-gtfs-idf/information/>. 956  
957
- 958 Karypis G, Kumar V, 1998 *A fast and high quality multilevel scheme for partitioning irregular graphs*. *SIAM Journal on scientific Computing* 20(1):359–392, URL <http://dx.doi.org/10.1137/S1064827595287997>. 959  
960
- 961 Kaufman DE, Smith RL, 1993 *Fastest paths in time-dependent networks for intelligent vehicle-highway systems application*. *Journal of Intelligent Transportation Systems* 1(1), URL <http://dx.doi.org/10.1080/10248079308903779>. 962  
963
- 964 Kernighan BW, Lin S, 1970 *An efficient heuristic procedure for partitioning graphs*. *Bell system technical journal* 49(2):291–307, URL <http://dx.doi.org/10.1002/j.1538-7305.1970.tb01770.x>. 965
- 966 Kirchler D, Liberti L, Pajor T, Wolfler Calvo R, 2011 *Unialt for regular language constrained shortest paths on a multi-modal transportation network*. *Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS)*, URL <http://dx.doi.org/10.4230/OASICS.ATMOS.2011.64>. 967  
968  
969
- 970 Maue J, Sanders P, Matijevic D, 2010 *Goal-directed shortest-path queries using precomputed cluster distances*. *Journal of Experimental Algorithmics (JEA)* 14:3–2, URL [http://dx.doi.org/10.1007/11764298\\_29](http://dx.doi.org/10.1007/11764298_29). 971  
972

- 973 Mooney P, Minghini M, et al., 2017 *Mapping and the Citizen Sensor*, chapter A review of OpenStreetMap  
974 data, 37–59 (Ubiquity Press), URL <http://dx.doi.org/10.5334/bbf.c>.
- 975 Müller-Hannemann M, Schulz F, Wagner D, Zaroliagis C, 2007 *Timetable information: Models and algo-*  
976 *rithms. Algorithmic Methods for Railway Optimization*, 67–90, URL [http://dx.doi.org/10.1007/](http://dx.doi.org/10.1007/978-3-540-74247-0_3)  
977 [978-3-540-74247-0\\_3](http://dx.doi.org/10.1007/978-3-540-74247-0_3).
- 978 Nannicini G, Dellling D, Liberti L, Schultes D, 2008 *Bidirectional a\* search for time-dependent fast paths.*  
979 *International Workshop on Experimental and Efficient Algorithms*, 334–346, URL [http://dx.doi.](http://dx.doi.org/10.1007/978-3-540-68552-4_25)  
980 [org/10.1007/978-3-540-68552-4\\_25](http://dx.doi.org/10.1007/978-3-540-68552-4_25).
- 981 Navitia Open Data, 2020 *France public transport data.* [https://navitia.opendatasoft.com/explore/](https://navitia.opendatasoft.com/explore/?refine.geographicarea=France)  
982 [?refine.geographicarea=France](https://navitia.opendatasoft.com/explore/?refine.geographicarea=France).
- 983 Orda A, Rom R, 1990 *Shortest-path and minimum-delay algorithms in networks with time-dependent*  
984 *edge-length. Journal of the ACM (JACM)* 37(3):607–625, URL [http://dx.doi.org/10.1145/79147.](http://dx.doi.org/10.1145/79147.214078)  
985 [214078](http://dx.doi.org/10.1145/79147.214078).
- 986 Pyrga E, Schulz F, Wagner D, Zaroliagis C, 2008 *Efficient models for timetable information in public trans-*  
987 *portation systems. Journal of Experimental Algorithmics (JEA)* 12, URL [http://dx.doi.org/10.](http://dx.doi.org/10.1145/1227161.1227166)  
988 [1145/1227161.1227166](http://dx.doi.org/10.1145/1227161.1227166).
- 989 Pyrga E, Schulz F, Wagner D, Zaroliagis CD, 2004 *Experimental comparison of shortest path approaches for*  
990 *timetable information. ALENEX/ANALC*, 88–99 (Citeseer).
- 991 Sanders P, Schulz C, 2012 *Distributed evolutionary graph partitioning. Workshop on Algorithm Engineering*  
992 *and Experiments (ALENEX)*, 16–29 (SIAM), URL <http://dx.doi.org/10.5555/2790265.2790267>.
- 993 Sauer J, Wagner D, Zündorf T, 2020 *Faster Multi-Modal Route Planning With Bike Sharing Using ULTRA.*  
994 *18th International Symposium on Experimental Algorithms (SEA)*, volume 160, 16:1–16:14, URL [http:](http://dx.doi.org/10.4230/LIPIcs.SEA.2020.16)  
995 [//dx.doi.org/10.4230/LIPIcs.SEA.2020.16](http://dx.doi.org/10.4230/LIPIcs.SEA.2020.16).
- 996 Schilling H, 2012 *Tomtom navigation—how mathematics help getting through traffic faster. International*  
997 *Symposium on Mathematical Programming (ISMP)*.
- 998 Sedgewick R, Vitter JS, 1986 *Shortest paths in euclidean graphs. Algorithmica* 1:31–48, URL [http://dx.](http://dx.doi.org/10.1007/BF01840435)  
999 [doi.org/10.1007/BF01840435](http://dx.doi.org/10.1007/BF01840435).
- 1000 Strasser B, 2017 *Dynamic time-dependent routing in road networks through sampling. Workshop on Algo-*  
1001 *rithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS)*, URL [http:](http://dx.doi.org/10.4230/OASIcs.ATMOS.2017.3)  
1002 [//dx.doi.org/10.4230/OASIcs.ATMOS.2017.3](http://dx.doi.org/10.4230/OASIcs.ATMOS.2017.3).
- 1003 Thomson RC, Richardson DE, 1995 *A graph theory approach to road network generalisation. International*  
1004 *cartographic conference*, 1871–1880.
- 1005 Ulloa L, Lehoux-Lebacque V, Roulland F, 2018 *Trip planning within a multimodal urban mobility. IET Intel-*  
1006 *ligent Transport Systems* 12(2):87–92(5), URL <http://dx.doi.org/10.1049/iet-its.2016.0265>.
- 1007 Yang F, Vereshchaka A, Dong W, 2018 *Predicting and optimizing city-scale road traffic dynamics using*  
1008 *trajectories of individual vehicles. IEEE International Conference on Big Data (Big Data)*, 173–180,  
1009 URL <http://dx.doi.org/10.1109/BigData.2018.8622356>.