



Constrained Backward Time Travel Planning is in P

Quentin Bramas, Jean-Romain Luttringer, Sébastien Tixeuil

► To cite this version:

Quentin Bramas, Jean-Romain Luttringer, Sébastien Tixeuil. Constrained Backward Time Travel Planning is in P. 2022. hal-03655239v1

HAL Id: hal-03655239

<https://hal.science/hal-03655239v1>

Preprint submitted on 3 May 2022 (v1), last revised 14 Dec 2023 (v2)

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.

Constrained Backward Time Travel Planning is in P

Quentin Bramas¹

Jean-Romain Luttringer¹

Sébastien Tixeuil²

¹ICUBE, Strasbourg University, Strasbourg, France

²Sorbonne University, CNRS, LIP6, Paris, France

Abstract

We consider transportation networks that are modeled by dynamic graphs, and introduce the possibility for traveling agents to use Backward Time-Travel (BTT) devices at any node to go back in time (to some extent, and with some appropriate fee) before resuming their trip. We focus on dynamic line graphs.

In more detail, we propose exact algorithms to compute travel plans with constraints on the BTT cost or on how far back in time you can go, while minimizing travel delay (that is, the time difference between the arrival instant and the starting instant), in polynomial time. We study the impact of the BTT devices pricing policies on the computation process of those plans considering travel delay and cost, and provide necessary properties that pricing policies should satisfy to enable to compute such plans. Finally, we provide an optimal online algorithm for the unconstrained problem when the cost function is the identity.

1 Introduction

Backward time travel (BTT) (that is, the ability to go back in time) enables various applications, such as increasing the ratio of sunny days when visiting Brittany for vacation, and allowing researchers to avoid submitting papers to conferences that will not accept them (therefore boosting their productivity and self-esteem). In this paper, we consider the application of BTT devices to transportation networks. In particular we focus on the ability to travel from point A to point B with minimal delay (that is, minimizing the time difference between arrival and start instants), taking into account meaningful constraints, such as the cost induced by BTT devices, or their span (how far back in time you are allowed to go).

To this paper, BTT was mostly envisioned in unrealistic settings (with respect to the cost associated to time travel or its span). For example, the AE model [11] considers that a single cost unit permits to travel arbitrarily in both space and time, trivializing the space-time travel problem entirely. Slightly more realistic models such as TM [10] and BTTF [13] consider devices that either: (i) only permit time travel [10] (but remain at the same position), or (ii) permit either time travel or space travel, but not both at the same time [13]. However, the cost involved is either null [10], or single cost unit per time travel [13]. This paper is the first to discuss BTT in a realistic, cost-aware, span-aware context, that implies efficiently using BTT devices within a transportation system (from a simultaneous delay and cost point of view), and the computation of the corresponding multi-modal paths.

More precisely, in this paper, we address the problem of space-time travel planning, taking into account both the travel delay of the itinerary and the cost policy of BTT device providers. The context we consider is that of transportation systems, where BTT devices are always available to the agents traveling. Using each BTT device has nevertheless a cost, decided by the BTT device provider, and may depend on the span of the backward time jump. Although BTT devices are always active, the ability to go from one location to another (that is, from one BTT device to another) varies across time. We consider that this ability is conveniently modeled by a dynamic graph, whose nodes represent BTT devices, and whose edges represent the possibility to instantly go from one BTT device to another. We consider that agents first fix their itinerary (that is, the list of BTT devices they plan to visit to reach their destination) and that

their itinerary is simple (a given BTT device occurs only once in the list). Given an itinerary, we aim at computing travel plans, from one BTT device to another (the closest to the agent’s actual destination), considering not only travel delay and induced cost, but also schedule availability and common limitations of BTT devices.

In the following, we study the feasibility of finding such travel plans, depending on the pricing policy. It turns out that when the schedule of connections is available, very loose conditions on the pricing policy enable to devise optimal algorithms (with respect to the travel delay and induced cost) in polynomial time, given a cost constraint for the agents, or a span constraint for the BTT devices. Furthermore, when the schedule of the transportation system is unavailable, we propose an online algorithm solving the same problem with an optimal competitive ratio. We hope our study will help future BTT device providers to clarify their price policy, and space-time travelers to maximize their travel experience without compromising their wealth.

Related Work. Space-Time routing has been studied, but assuming only forward time travel, *i.e.*, waiting, is available. The idea of using dynamic graphs to model transportation network was used by many studies (see *e.g.* Casteigts et al. [2] and references herein), leading to recently revisit popular problems previously studied in static graphs [1, 4, 9]. The closest work in this research path is due to Casteigts et al [3], who study the possibility of discovering a temporal path between two nodes in a dynamic network with a waiting time constraint: at each step, the traveling agent cannot wait more than c time instants, where c is a given constant. As previously mentioned, this line of work only considers *forward* time travel: a temporal path cannot go back in time.

Constrained-shortest-paths computation problems have been extensively studied in the context of static graphs [5]. Although these problems may be NP-Hard even when considering a single metric, the ones considering two additive metrics (commonly, the delay and a cost) gained a lot of traction over the years due to their practical relevance, the most common use-case being computer networks [7, 8]. In this context, each edge is characterized by a weight vector, comprising both cost and delay. Path computation algorithms thus have to maintain and explore all non-comparable paths, whose number may grow exponentially with respect to the size of the network. To avoid a worst-case exponential complexity, most practical algorithms rely on either approximation schemes [12] or heuristics. However, these contributions do not study multi-criteria path computation problems within a time travel context. Conversely, we study and provide results regarding the most relevant time-traveling problems while considering the peculiarities of this context (in particular, the properties of the cost function). In addition, we show that most of these problems can be solved optimally in polynomial-time.

Contributions. In this paper, we provide the following contributions:

- An in-depth analysis of the impact of the BTT device providers pricing policies on the computation of low-latency and low-cost paths. In particular, we show that few features are required to ensure that the efficient computation of such paths remains possible.
- Two exact polynomial algorithms able to compute travels with smallest delay to a given destination and minimizing the cost of traveling back in time. The first algorithm also supports the addition of a constraint on the backward cost of the solution. The other one supports a constraint on how far back in the past one can go at each given time instant.
- An online algorithm with optimal competitive ratio able to compute a travel with lowest latency and having a cost of at most two times the optimal cost.

Our algorithms assume the footprint of the graph is a line.

2 Model

In this section, we define the models and notations used throughout this paper, before formalizing the aforementioned problems.

We represent the network as an evolving graph, as introduced by Ferreira [6]: a graph-centric view of the network that maps a dynamic graph as a sequence of static graphs. The *footprint* of the dynamic graph (that includes all nodes and edges that appear at least once during the lifetime of the dynamic graph), is fixed. Furthermore, we assume that the set of nodes is fixed over time, while the set of edges evolves.

More precisely, an evolving graph G is a pair $(V, (E_t)_{t \in \mathbb{N}})$, where V denotes the set of vertices, \mathbb{N} is the set of time instants, and for each $t \in \mathbb{N}$, E_t denotes the set of edges that appears at time t . The *snapshot* of G at time t is the static graph $G(t) = (V, E_t)$, which corresponds to the state, supposedly fixed, of the network in the time interval $[t, t+1)$. The *footprint* $\mathcal{F}(G)$ of G is the static graph corresponding the union of all its snapshots, $\mathcal{F}(G) = (V, \bigcup_{t \in \mathbb{N}} E_t)$. We say $((u, v), t)$ is a temporal edge of graph G if $(u, v) \in E_t$.

Space-time Travel. We assume that at each time instant, an agent can travel along any number of adjacent consecutive communication links. However, the graph may not be connected at each time instant, hence it may be that the only way to reach a particular destination node is to travel forward (*i.e.*, wait) or backward in time, to reach a time instant where an adjacent communication links exists. In more details, an agent travels from a node s to a node d using a *space-time travel* (or simply travel when it is clear from the context).

Definition 1. A space-time travel is a sequence $((u_0, t_0), (u_1, t_1), \dots, (u_k, t_k))$ such that

- $\forall i \in \{0, \dots, k\}$, $u_i \in V$ is a node and $t_i \in \mathbb{N}$ is a time instant,
- $\forall i \in \{0, \dots, k-1\}$, if $u_i \neq u_{i+1}$, then $t_i = t_{i+1}$ and $(u_i, u_{i+1}) \in E_{t_i}$ *i.e.*, there is a temporal edge between u_i and u_{i+1} at time t_i .

By extension, the *footprint* of a travel is the static graph containing all edges (and their adjacent nodes) appearing in the travel. Now, the *itinerary* of a travel $((u_0, t_0), (u_1, t_1), \dots, (u_k, t_k))$ is its projection (u_0, u_1, \dots, u_k) on nodes, while its *schedule* is its projection (t_0, t_1, \dots, t_k) on time instants.

Definition 2. A travel $((u_0, t_0), (u_1, t_1), \dots, (u_k, t_k))$ is simple if for all $i \in \{2, \dots, k\}$ and $j \in \{0, \dots, i-2\}$, we have $u_i \neq u_j$.

Intuitively, a travel is simple if its footprint is a line (*i.e.*, a simple path) and contains at most one time travel per node (as a consequence, no node appears three times consecutively in a simple travel).

Definition 3. The delay of a travel $T = ((u_0, t_0), (u_1, t_1), \dots, (u_k, t_k))$, denoted $\text{delay}(T)$ is defined as $t_k - t_0$.

The Backward cost of a travel.

Definition 4. The backward-cost is the cost of going to the past. The backward-cost function $\mathfrak{f} : \mathbb{N}^* \rightarrow \mathbb{R}^+$ returns, for each $\delta \in \mathbb{N}$, the backward-cost $\mathfrak{f}(\delta)$ of traveling δ time instants to the past. As we assume that there is no cost associated to forward time travel (that is, waiting), we extend \mathfrak{f} to \mathbb{Z} by setting $\mathfrak{f}(-\delta) = 0$, for all $\delta \in \mathbb{N}$. In particular, the backward-cost of traveling 0 time instants in the past is zero. When it is clear from context, the backward-cost function is simply called the cost function.

Definition 5. The backward-cost (or simply cost) of a travel $T = ((u_0, t_0), (u_1, t_1), \dots, (u_k, t_k))$, denoted $\text{cost}(T)$ is defined as follows:

$$\text{cost}(T) = \sum_{i=0}^{k-1} \mathfrak{f}(t_i - t_{i+1})$$

Definition 6. Let $T_1 = ((u_0, t_0), (u_1, t_1), \dots, (u_k, t_k))$ and $T_2 = ((u'_0, t'_0), (u'_1, t'_1), \dots, (u'_k, t'_k))$ be two travels. If $(u_k, t_k) = (u'_0, t'_0)$, then the concatenated travel $T_1 \oplus T_2$ as follows:

$$T_1 \oplus T_2 = ((u_0, t_0), (u_1, t_1), \dots, (u_k, t_k), (u'_1, t'_1), \dots, (u'_k, t'_k))$$

Remark 1. One can easily prove that $\text{cost}(T_1 \oplus T_2) = \text{cost}(T_1) + \text{cost}(T_2)$. In the following, we sometimes decompose a travel highlighting an intermediate node: $T = T_1 \oplus ((u_i, t_i)) \oplus T_2$. Following the definition, this means that T_1 ends with (u_i, t_i) , and T_2 starts with (u_i, t_i) , so we also have $T = T_1 \oplus T_2$ and $\text{cost}(T) = \text{cost}(T_1) + \text{cost}(T_2)$.

Our notion of space-time travel differs from the classical notion of *journey* found in literature related to dynamic graphs [6] as we do *not* assume time instants monotonically increase along a travel. As a consequence, some evolving graphs may not allow a journey from A to B yet allows one or several travels from A to B (See Figure 2).

Problem specification. We now present the problems that we aim to solve in this paper. First, we want to arrive at the destination as early as possible, *i.e.*, finding a time travel that minimizes the delay. Among such travels, we want to find one that minimizes the backward cost.

Definition 7. The Optimal Delay Optimal Cost space-time travel planning (ODOC) problem consists, given a cost function \mathfrak{f} , an evolving graph $G = (V, (E_t)_{t \in \mathbb{N}})$, and two nodes x and y in V , in computing, among all travels starting at node x at time 0, arriving at node y , and minimizing the travel delay, a travel with minimum cost. A solution to the ODOC problem is called an ODOC travel.

Definition 8. The \mathcal{C} -constrained ODOC problem consists in finding a ODOC travel with cost at most $\mathcal{C} \geq 0$.

Definition 9. The \mathcal{H} -history-constrained ODOC problem consists in finding a ODOC travel T satisfying, $\forall u, u', t, t'$, if $T = T_1 \oplus ((u, t)) \oplus T_2 \oplus ((u', t')) \oplus T_3$, then $t' \geq t - \mathcal{H}$.

In the remaining, we consider evolving graphs whose footprint is a line graph. Arbitrary topologies are left for “future” work. Thus, we assume that the graph has $n + 1$ nodes denoted x_0, x_1, \dots, x_n , and that the footprint of the graph is the line graph connecting x_i with x_{i+1} for all $i \in [0, n - 1]$. The problem we want to solve is to find a time travel from x_0 to x_n starting at time 0 that minimizes the travel delay and, among such travels, that has minimal backward cost (either \mathcal{C} -constrained or \mathcal{H} -history-constrained).

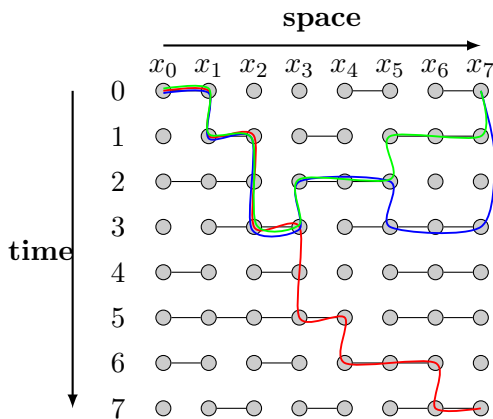


Figure 1: Possible representation of an evolving graph. Possible travels from x_0 to x_7 are shown in red, green and blue. Note that the blue and green travels require to send an agent to the past (to a previous time instant).




Figure 2: Example of an evolving graph for which there exists no journey, yet there exists several travels from x_0 to x_7 (in blue and green).




Figure 3: Example of an evolving graph for which there exist at least three travels from x_0 to x_7 with a cost constraint of 1 (assuming $f : d \mapsto d$). The blue travel has optimal delay.




Figure 4: Example of an evolving graph for which there exist at least three travels from x_0 to x_7 with a cost constraint of 3 (assuming $f : d \mapsto d^2$). The blue and green travels have optimal delay. Among the travels with optimal delay, the blue travel has optimal cost (2).

Visual representation of space-time travels.. Since the footprint of the graph is a line, the associated evolving graph can be seen as a partial grid mentioning for each time instant which edges of the line are present. A possible visual representation of an evolving graph can be seen in Fig. 1. One can see the evolution of the line topology (consisting of the nodes x_0 to x_7) over time through eight snapshots performed from time instants 0 to 7. Several possible travels are shown in red, green and blue. The red travel only makes use of forward time travel (that is, waiting) and is the earliest arriving travel in this class (arriving at time 7). The green and blue travels both make use of backward time travel and arrive at time 0, so they have minimal travel delay. Similarly, the red travel concatenated with $((x_7, 7), (x_7, 0))$ (i.e., a backward travel to reach x_7 at time 0) also has minimal travel delay. However, if we assume that the cost function is the identity ($f : d \mapsto d$) then the green travel has a backward cost of 3, the blue travel has a backward cost of 4, and the concatenated red travel has a backward cost of 7. Adding constraints yields more challenging issues: assuming $f : d \mapsto d$ and a maximal cost C of 1, at least three travels can be envisioned for the evolving graph depicted in Figure 3, but finding the 1-constrained travel that minimizes the delay (that is, the blue travel) is not as straightforward in this case. Similarly, in Figure 2 we show two \mathcal{H} -history-constrained travels, with $\mathcal{H} = 1$ (assuming $f : d \mapsto d$). Here, clearly, the green travel is optimal with a cost of 2 (the blue travel has cost 3). The choice made by the green travel to wait at node x_2 one time instant was a good one because there is no need to go back at time 0 (that would be impossible after node x_2 for the green travel because of the 1-history-constraint). If we add more nodes to the graph and repeat this kind of choice, we can create a graph with an exponential number of 1-history-constrained travels and finding one that minimizes the cost is challenging. Surprisingly, we show that it remains polynomial in the number of nodes and edges.

3 Backward-cost Function Classes

The cost function f represents the cost of going back to the past, according to the policy of the BTT device provider. Intuitively, it seems reasonable that the function is non-decreasing (so that BTT providers charge travelers more if they go further back in time), however we demonstrate that such a pricing policy is not necessary to enable travelers to derive optimal cost space-time travel plans. As a matter of fact, the two necessary conditions we identify to optimally solve the space-time travel planning problem are f to be non-negative (which also allows to prevent the bankruptcy of BTT device providers) and that it attains its minimum (not just converge to it). Due to space constraints, some proofs are moved to the appendix.

Definition 10. A cost function f is user optimizable if it is non-negative, and it attains its minimum when restricted to any interval $[C, \infty)$, with $C > 0$. Let \mathcal{UO} be the set of user optimizable cost functions.

Theorem 1. *If the cost function \mathbf{f} is not in \mathcal{UO} , then there exist evolving graphs where no solution exists for the optimal cost space-time travel planning problem.*

Proof. First, it is clear that if $\mathbf{f}(d) < 0$ for some $d \in \mathbb{N}^*$, then we can construct travels with arbitrarily small cost by repeatedly appending $((y, t), (y, t + d), (y, t))$ to any travel arriving at node y at time t (i.e., by waiting for d rounds and going back in time d rounds), rendering the problem unsolvable.

Now, let $C \in \mathbb{N}^*$ and \mathbf{f} be a non-negative function that does not attain its minimum when restricted to $[C, \infty)$. This implies that there exists an increasing sequence $(w_i)_{i \in \mathbb{N}}$ of integers $w_i \geq C$, such that the sequence $(\mathbf{f}(w_i))_{i \in \mathbb{N}}$ is decreasing and converges towards the lower bound $m_C = \min_{t \geq C} (\mathbf{f}(t))$ of \mathbf{f} . Consider a graph with two nodes x_0 and x_1 that are connected by a temporal edge at time C and disconnected at other time instants. Since a travel from x_0 to x_1 arriving at time 0 must contain a backward travel to the past of amplitude at least C , its cost is at least equal to m_C . Since m_C is not attained, there is no travel with cost exactly m_C . Now, assume for the sake of contradiction that a cost-optimal travel T to x_1 arriving at time 0 has cost $m_C + \varepsilon$ with $\varepsilon > 0$. Then, we can construct a travel with a smaller cost. Let i_ε such that $\mathbf{f}(w_{i_\varepsilon}) < m_C + \varepsilon$ (this index exists because the sequence $(\mathbf{f}(w_i))_{i \in \mathbb{N}}$ converges to m_C).

Let $T' = ((x_0, 0), (x_0, C), (x_1, C), (x_1, w_{i_\varepsilon}), (x_1, 0))$. Then we have

$$\text{cost}(T') = \mathbf{f}(w_{i_\varepsilon}) < m_C + \varepsilon = \text{cost}(T),$$

which contradicts the optimality of T . \square

We now present the set of *user friendly* cost functions that we use in the sequel to ease proving optimization algorithms, as they allow *simple* solutions to the ODOC problem (Lemma 1). We prove in Theorem 2 that we do not lose generality since an algorithm solving the ODOC problem with user friendly cost functions can be transformed easily to work with any user optimizable ones.

Definition 11. *A cost function \mathbf{f} is user friendly if it is user optimizable, non-decreasing, and sub-additive¹. Let \mathcal{UF} be the set of user friendly cost functions.*

Lemma 1. *If the cost function \mathbf{f} is in \mathcal{UF} , then there exists a simple travel solution to the optimal cost space-time travel planning problem.*

Proof. Let T be a solution to the optimal cost space-time travel planning problem (such a solution exists since \mathbf{f} is in \mathcal{UF} , hence in \mathcal{UO}). If there exists a node x_i and two time instants t_1 and t_2 , such that $T = T_1 \oplus ((x_i, t_1)) \oplus T_2 \oplus ((x_i, t_2)) \oplus T_3$, then we construct T' as follows

$$T' = T_1 \oplus ((x_i, t_1), (x_i, t_2)) \oplus T_3$$

and we show that $\text{cost}(T') \leq \text{cost}(T)$. Indeed, it is enough to show (thanks to Remark 1) that

$$\text{cost}(((x_i, t_1), (x_i, t_2)))) \leq \text{cost}(T_2).$$

By definition $\text{cost}(((x_i, t_1), (x_i, t_2)))) = \mathbf{f}(t_1 - t_2)$. If $t_1 < t_2$, then the cost is null by convention and the Lemma is proved. Otherwise $t_1 > t_2$. On the right hand side, we have:

$$\text{cost}(T_2) = \sum_{i=1}^k \mathbf{f}(d_i)$$

where d_1, d_2, \dots, d_k is the sequence of differences between the times appearing in T_2 . Since T_2 starts at time t_1 and ends at time t_2 , then $\sum_{i=1}^k d_i = t_1 - t_2$. Since the function is sub-additive and increasing, we obtain:

$$\mathbf{f}(t_1 - t_2) < \sum_{i=1}^k \mathbf{f}(d_i)$$

¹sub-additive means that for all $a, b \in \mathbb{N}$, $\mathbf{f}(a + b) \leq \mathbf{f}(a) + \mathbf{f}(b)$

By repeating the same procedure, we construct a time-travel with the same destination and same backward-cost as T but that does not contain two occurrences of the same node, except if they are consecutive. \square

Theorem 2. *If an algorithm A solves the optimal cost space-time travel planning problem for any cost function in \mathcal{UF} , then there exists an algorithm A' solving the same problem with any f in \mathcal{UO} .*

Proof. We consider an algorithm A as stated. Let f be an arbitrary cost function in \mathcal{UO} , that is, f is non-negative, and always attains its minimum.

From f , we now construct a cost function f_{inc} as follows:

$$f_{inc}(t) = \min_{j \geq t} (f(j))$$

By construction, f_{inc} is non-decreasing. Moreover, since f is in \mathcal{UO} , it always attains its minimum, and we have:

$$\forall d, \exists d_m \text{ such that } f_{inc}(d) = f(d_m). \quad (1)$$

Then, we construct \tilde{f} as follows:

$$\tilde{f}(t) = \min_{a \in \alpha(t)} \left(\sum_{a_i \in a} f_{inc}(a_i) \right)$$

where $\alpha(t)$ is the set of all the non-negative sequences that sum to t . Now, \tilde{f} is sub-additive by construction, hence $\tilde{f} \in \mathcal{UF}$. Since $\alpha(t)$ is finite, the minimum is attained.

Also, $\forall t \geq 1, \tilde{f}(t) \leq f(t)$, so that for any travel, its backward cost with respect to f is at least equal to its backward cost with respect to \tilde{f} .

Let G be a dynamic graph. Our goal is to find an optimal cost (with respect to f) space-time travel plan in G . Let \tilde{T} be an optimal solution found by algorithm A on G assuming function \tilde{f} is used. We now construct, from \tilde{T} , a time-travel T that is a cost optimal (with respect to f) solution on G .

The travel T is constructed from \tilde{T} by replacing any sub-space-time travel $((x_i, t_i), (x_i, t_i - t))$, with $t \geq 0$, by the following sub space-time travel: $((x_i, t_i - a_1), (x_i, t_i - a_1 - a_2), \dots, (x_i, t_i - \sum_{j=1}^k a_j))$ satisfying:

$$a \in \alpha(t) \quad \wedge \quad \tilde{f}(t) = \sum_{j=1}^{\text{length of } a} f_{inc}(a_j)$$

Then, each $((u, t), (u, t - d))$, with $d \geq 0$, is replaced by $((u, t), (u, t - d + d_m), (u, t - d))$ such that:

$$d_m \geq d \quad \wedge \quad f_{inc}(d_m) = f(d)$$

We know that d_m exists thanks to Equation 1. The space-time travel T uses the same temporal edges as \tilde{T} , so it is well defined. Moreover, by construction $f(T) = \tilde{f}(\tilde{T})$, and T is optimal with respect to f because the backward-cost of a travel with respect to f is at least equal to its backward-cost with respect to \tilde{f} , as observed earlier. Hence, if a better solution exists for f , it is also a solution with the same, or smaller, cost with \tilde{f} , contradicting the optimality of \tilde{T} . The above procedure defines an algorithm, based on A , that solves the ODOC problem with function f . \square

4 Offline \mathcal{C} -constrained ODOC Algorithm

In this section, we present Algorithm 1 that solves the \mathcal{C} -constrained ODOC problem in time polynomial in the number of edges. More precisely, since the number of edges can be infinite, we only consider edges occurring before a certain travel (see the end of the section for a more precise description of the

complexity). Algorithm 1 has some similarities with the Dijkstra algorithm, but in our case, we need to take into account the cost and the delay of travels. At each iteration, we extract the minimum cost to reach a particular node at a particular time and we extend travels from there by updating the best-known cost of the next node. We reach the next nodes either by using the next temporal edge that exists in the future (we prove that considering only the next future edge is enough) or using one of the past temporal edges. As a convention, we consider that graph G has $n + 1$ nodes denoted x_0, x_1, \dots, x_n .

Algorithm 1: Offline \mathcal{C} -constrained ODOC Algorithm

```

1  /* nodeCost[i,t] stores the current best cost of travels to node  $x_i$  arriving at time  $t$  */
2   $\forall i, \forall t, \text{ nodeCost}[i, t] = \infty$ ;
3   $\text{nodeCost}[0, 0] \leftarrow 0$ ;  $\text{done} \leftarrow \emptyset$ ;
4  while  $\exists (i, t) \notin \text{done}$  such that  $\text{nodeCost}[i, t] < \infty$  do
5       $(i, t) \leftarrow \text{argmin}_{(i, t) \notin \text{done}} (\text{nodeCost}[i, t])$ ;
6       $\text{done} \leftarrow \text{done} \cup \{(i, t)\}$ ;
7       $c \leftarrow \text{nodeCost}[i, t]$ ;
8      let  $t_{\text{future}}$  the smallest time after (or equal to)  $t$  where edge  $((x_i, x_{i+1}), t_{\text{future}})$  exists;
9      if  $\text{nodeCost}[i + 1, t_{\text{future}}] > c$  then
10          $\text{nodeCost}[i + 1, t_{\text{future}}] \leftarrow c$ ;
11     end
12     for each  $t_{\text{past}}$  such that  $(x_i, x_{i+1}) \in E_{t_{\text{past}}}$  do
13         let  $c_{\text{past}} = c + f(t - t_{\text{past}})$ 
14         if  $c_{\text{past}} \leq \mathcal{C}$  and  $\text{nodeCost}[i + 1, t_{\text{past}}] > c_{\text{past}}$  then
15              $\text{nodeCost}[i + 1, t_{\text{past}}] \leftarrow c_{\text{past}}$ ;
16         end
17     end
18 let  $t_{\min}$  be the minimum time instant such that  $\exists t, \text{nodeCost}[n, t] + f(t - t_{\min}) \leq \mathcal{C}$ ;
19 if  $t_{\min}$  exists then
20     return ExtractTimeTravel( $n, t_{\min}, \text{nodeCost}$ );
21 end
22 return  $\perp$ 

```

Algorithm ExtractTimeTravel: Extract a \mathcal{C} -constrained ODOC travel to the given destination

```

input:  $i \in [0, n], \quad t \in \mathbb{N}, \quad \text{nodeCost}$ 
1 if  $i = 0$  then
2     return  $((x_0, 0), (x_0, t))$ 
3 end
4 if  $\text{nodeCost}[i, t] = \infty$  then
5     Let  $t' = \text{argmin}_{t' \in \mathbb{N}} (\text{nodeCost}[i, t'] + f(t' - t))$ ;
6     return ExtractTimeTravel( $i, t', \text{nodeCost}$ )  $\oplus ((x_i, t'), (x_i, t))$ 
7 end
8 Let  $t'$  such that  $(x_{i-1}, x_i) \in E_{t'}$  and  $\text{nodeCost}[i - 1, t'] + f(t' - t) = \text{nodeCost}[i, t]$ 
   /*  $t'$  exists by construction of nodeCost. nodeCost[i, t] either equals to nodeCost[i - 1, t'] + f(t' - t)
   with some  $t' > t$  (Line 9) or equals nodeCost[i - 1, t'] with some  $t' < t$  (Line 14), hence  $f(t' - t) = 0$ .
   In both cases, edge  $((x_{i-1}, x_i), t')$  exists. */
9 return ExtractTimeTravel( $i - 1, t', \text{nodeCost}$ )  $\oplus ((x_{i-1}, t'), (x_i, t'), (x_i, t))$ 

```

We first prove the main property satisfied by our algorithm that we then use to construct a solution. Let $\delta_{\mathcal{C}}$ be the function that returns, for each pair (x, t) where x is a node and t a time, the best backward-cost smaller or equal to \mathcal{C} , from x_0 to x , for travels arriving at time t .

Lemma 2. When a pair (i, t) is extracted from **nodeCost** at line 4, then $\delta_{\mathcal{C}}(x_i, t) = \text{nodeCost}[i, t]$.

Proof. Assume for the sake of contradiction that this is not true, and let (i, t) be the first tuple extracted

such that the property is false. Let $c_{i,t} = \text{nodeCost}[i, t]$. Let T be a \mathcal{C} -constrained-backward-cost-optimal travel to x_i arriving at time t (hence $\text{cost}(T) < c_{i,t}$ by assumption).

Let T' be the longest prefix of T , to (x_j, t') (i.e., such that $T = T' \oplus (x_j, t') \oplus T''$, for some T''), such that (j, t') was extracted from **nodeCost** and satisfies $\delta_{\mathcal{C}}(x_j, t') = \text{nodeCost}[j, t']$. Now, T' is well defined because the first element in T is $(x_0, 0)$ and, by Line 2, $(0, 0)$ is the first extracted pair, and satisfies $\text{nodeCost}[0, 0] = 0 = \delta_{\mathcal{C}}(x_0, 0)$. Hence, prefix $((x_0, 0))$ satisfies the property, so the longest of such prefixes exists. Observe that T' , resp. T'' , ends, resp. starts, with (x_j, t') , by the definition of travel concatenation.

When (j, t') is extracted from **nodeCost**, it is extended to the next future edge (Lines 7 to 9), and all past edges (Lines 11 to 14). T'' starts either (a) with $((x_j, t'), (x_j, t_a), (x_{j+1}, t_a))$, with $t_a < t'$, (b) with $((x_j, t'), (x_j, t_a), (x_{j+1}, t_a))$ with $t_a > t'$, or (c) with $(x_j, t'), (x_{j+1}, t')$.

In case (a), this means that the edge $((x_j, x_{j+1}, t_a))$ exists, hence, by Line 14, we know that $\text{nodeCost}[j+1, t_a] \leq \text{nodeCost}[j, t'] + f(t' - t_a)$. However, by induction hypothesis $\text{cost}(T') = \text{nodeCost}[j, t']$, hence

$$\text{nodeCost}[j+1, t_a] \leq \text{cost}(T' \oplus ((x_j, t'), (x_j, t_a), (x_{j+1}, t_a)))$$

and $\text{nodeCost}[j+1, t_a] \leq \mathcal{C}$, which contradicts the definition of T' . Indeed, $T' \oplus ((x_j, t'), (x_j, t_a), (x_{j+1}, t_a))$ is a longer prefix of T with the same property as T' .

In case (b), this means that the edge $((x_j, x_{j+1}, t_a))$ exists, hence, by Line 9, we know that $\text{nodeCost}[j+1, t_a] \leq \text{nodeCost}[j, t']$. By assumption, we have $\text{cost}(T') = \text{nodeCost}[j, t']$, hence

$$\text{nodeCost}[j+1, t_a] \leq \text{cost}(T' \oplus ((x_j, t'), (x_j, t_a), (x_{j+1}, t_a)))$$

and $\text{nodeCost}[j+1, t_a] \leq \mathcal{C}$, which contradicts the definition of T' .

In case (c), this means that the edge $((x_j, x_{j+1}, t'))$ exists, which implies, using a similar argument, a contradiction. \square

The previous lemma says that **nodeCost** contains correct information about the cost to reach a node, but actually, it does not contain all the information. Indeed, a node x_i can be reachable by a travel at a given time t and still $\text{nodeCost}[i, t] = \infty$. This fact helps our algorithm to be efficient. We now prove that we can still find all existing travel using **nodeCost**.

Lemma 3. *For all $i \in [0, n]$, $t \in \mathbb{N}$, there exists a \mathcal{C} -constrained travel T to x_i arriving at time t , if and only if there exists $t' \in \mathbb{N}$ such that $\text{nodeCost}[i, t'] + f(t' - t) \leq \mathcal{C}$.*

Proof. We just need to prove the implication since the converse follows from the previous Lemma. Indeed, if $\text{nodeCost}[i, t'] + f(t' - t) \leq \mathcal{C}$, then $\text{nodeCost}[i, t']$ is finite and is the optimal cost of \mathcal{C} -constrained travels to i arriving at time t' , so a \mathcal{C} -constrained travel to x_i arriving at time t exists.

The proof of the implication is done by induction on the pair (i, t) (in the lexicographic order). The result is clearly true when $i = 0$ for any t . Now let $i \in [1, n]$ and $t \in \mathbb{N}$ and assume the result true for any pair $(j, k) < (i, t)$. Assume for the sake of contradiction that there exists a \mathcal{C} -constrained travel T to x_i arriving at time t such that

$$\forall t' \in \mathbb{N}, \text{nodeCost}[i, t'] + f(t' - t) > \mathcal{C}. \quad (2)$$

We can assume that T is simple and cost optimal. T does not arrive before time t otherwise the inductive hypothesis implies directly a contradiction.

Hence T goes through, x_{i-1} at time $t_{i-1} \geq t$:

$$T = T_1 \oplus ((x_{i-1}, t_{i-1}), (x_i, t_{i-1}), (x_i, t))$$

By inductive hypothesis, there exists $t' \in \mathbb{N}$ such that

$$\text{nodeCost}[i-1, t'] + f(t' - t_{i-1}) \leq \mathcal{C} \quad (3)$$

If $t' < t_{i-1}$, then, when the pair $(i-1, t')$ was extracted from **nodeCost** (in Line 4), since an edge exists between x_{i-1} and x_i at time t_{i-1} , then the variable t_{future} is at most t_{i-1} .

$$\mathbf{nodeCost}[i-1, t']$$

By Lemma 2, the optimal cost of travels to node x_i arriving at time t_{future} is $\mathbf{nodeCost}[i, t_{future}]$. Hence $\mathbf{nodeCost}[i, t_{future}] \leq \text{cost}(T') + f(t_{i-1} - t_{future})$. So we have, using the sub-additivity f ,

$$\begin{aligned} \mathbf{nodeCost}[i, t_{future}] + f(t_{future} - t) &\leq \text{cost}(T') + f(t_{i-1} - t_{future}) + f(t_{future} - t) \\ &\leq \text{cost}(T') + f(t_{i-1} - t) = \text{cost}(T) \leq C \end{aligned}$$

which contradicts Property (2).

If $t' \geq t_{i-1}$, then, when the pair $(i-1, t')$ was extracted from **nodeCost** (in Line 4), since an edge exists between x_{i-1} and x_i at time t_{i-1} , there is an iteration of the *for* loop where the variable $t_{past} = t_{i-1}$ and we have

$$\mathbf{nodeCost}[i, t_{i-1}] = \mathbf{nodeCost}[i-1, t'] + f(t' - t_{i-1}) \leq C$$

So $\mathbf{nodeCost}[i, t_{i-1}]$ is finite so, using Lemma 2, $\mathbf{nodeCost}[i, t_{i-1}] = \text{cost}(T')$. We can use this to have:

$$\text{cost}(T) = \text{cost}(T') + f(t_{i-1} - t) = \mathbf{nodeCost}[i, t_{i-1}] + f(t_{i-1} - t) \leq C$$

which contradicts Property (2). □

Now we show that **ExtractTimeTravel** is correct and return a solution to the problem, if it exists.

Lemma 4. *Assuming **nodeCost** is constructed by Algorithm 1, then **ExtractTimeTravel**($i, t, \mathbf{nodeCost}$) returns, if it exists, a travel with optimal cost to node x_i arriving at time t .*

Proof. First recall that, by Lemma 2, for any pair (i, t) , if $\mathbf{nodeCost}[i, t] < \infty$, then $\mathbf{nodeCost}[i, t]$ is the optimal cost of C -constrained travels arriving at node x_i and at time t .

We first prove the Lemma assuming $\mathbf{nodeCost}[i, t] < \infty$. By definition of t' in Line 8 we know that recursive call to **ExtractTimeTravel** also satisfy this property.

We want to show that the travel returned by **ExtractTimeTravel**($i, t, \mathbf{nodeCost}$) has cost $\mathbf{nodeCost}[i, t]$. We prove this result by induction. If $i = 0$ the result is clear. Otherwise, we have by construction of **nodeCost**:

$$\mathbf{nodeCost}[i-1, t'] + f(t' - t) = \mathbf{nodeCost}[i, t] \quad (4)$$

for some t' . Indeed, $\mathbf{nodeCost}[i, t]$ is obtained from $\mathbf{nodeCost}[i-1, t']$ either by extension on Line 9 or on Line 14 of Algorithm 1. Observe that in the former case $t' < t$, but in this case $f(t' - t) = 0$, so the equation remains true.

Since $\mathbf{nodeCost}[i-1, t']$ is finite, then, by induction hypothesis, **ExtractTimeTravel**($i-1, t', \mathbf{nodeCost}$) is a travel with optimal cost to x_{i-1} arriving at time t' .

Hence we have that the cost of the returned travel is in fact

$$\begin{aligned} \text{cost}(\mathbf{ExtractTimeTravel}(i-1, t', \mathbf{nodeCost}) \oplus ((x_{i-1}, t'), (x_i, t'), (x_i, t))) \\ = \mathbf{nodeCost}[i-1, t'] + f(t' - t) \end{aligned}$$

This is equal to $\mathbf{nodeCost}[i, t]$, using Equation 4.

Now, assume $\mathbf{nodeCost}[i, t] = \infty$. If a cost optimal C -constrained travel T exists to node x_i arriving at time t , then T must goes through a temporal edge $((x_{i-1}, x_i), t_1)$ for some $t_1 \in \mathbb{N}$ i.e.,

$$T = T_1 \oplus ((x_i, t_1), (x_i, t))$$

and by construction $\text{nodeCost}[i, t_1] = \text{cost}(T_1)$. So, when computing t' in Line 5 we have

$$\text{nodeCost}[i, t'] + f(t' - t) \leq \text{cost}(T)$$

Then the recursive call of **ExtractTimeTravel** is made with (i, t') so that the previous property holds and the returned travel has the optimal cost. \square

Theorem 3. *If the cost function f is in \mathcal{UF} and Algorithm 1 outputs a travel T , then T is a solution of the \mathcal{C} -constrained ODOC problem.*

Proof. By Lemma 3, t_{\min} is the smallest time such that a travel exists to node x_n . Then, using **nodeCost**, we can easily construct in a backward manner, a solution using a **ExtractTimeTravel** procedure, as proved by Lemma 4 \square

We now consider the complexity of Algorithm 1. We assume that retrieving the next edge (resp. the previous edge) that occurs after (resp. before) a given time is in $O(1)$. For instance, the graph can be stored as a dictionary that maps each node to an array mapping each time to the current, the previous, and the next temporal edges (the array can be made sparser easily with low complexity overhead to gain space if few edges occur per time-instant).

Since each temporal edge is extracted from the **nodeCost** at most once, and the inner *for* loop iterates over a subset of the edges, the time complexity is clearly polynomial in the number of temporal edges (we also have to consider the time to extract the min from **nodeCost**, which is polynomial). Moreover, if there is an infinite number of temporal edges², we show that our algorithm always terminates (even if no solution exists), and that its complexity is polynomial in the size of a finite subset of edges that we define next.

We start by defining the *slow-travel* to node x_i , denoted $\text{slowT}(x_i)$ inductively as follows. $\text{slowT}(x_i)$ is a simple travel from x_0 arriving at node x_i and passing through a temporal edge between x_{j-1} and x_j at time t_j^s (for $0 < j \leq i$) with:

$$t_0^s = 0 \quad \wedge \quad t_j^s = \begin{cases} \min(\{t \mid t \geq t_{j-1}^s \wedge (x_{j-1}, x_j) \in E_t\}) & \text{if the set is not empty} \\ \max(\{t \mid (x_{j-1}, x_j) \in E_t\}) & \text{otherwise} \end{cases} \quad (5)$$

t_i^s is well defined because, if $\{t \mid t \geq t_{j-1}^s \wedge (x_{j-1}, x_j) \in E_t\}$ is empty, then the set $\{t \mid (x_{j-1}, x_j) \in E_t\}$ is finite and the maximum exists (it is not empty because we assumed that the footprint is a line, so at least one edge exists between any two consecutive nodes). Hence,

$$\text{slowT}(x_0) = ((x_0, 0)) \quad \wedge \quad \text{slowT}(x_i) = \text{slowT}(x_{i-1}) \oplus ((x_{i-1}, t_{i-1}^s), (x_{i-1}, t_i^s), (x_i, t_i^s))$$

with t_j^s defined by Equation 5. Figures 5 and 6 illustrate the definition of $\text{slowT}(x_i)$ for two particular infinite evolving graphs.

²an evolving graph with an infinite number of edges can exist in practice even with bounded memory, e.g. when the graph is periodic.




Figure 5: Example of an infinite evolving graph where all edges occur periodically (with different periodicity patterns). The red travel corresponds to $slowT(x_7)$.




Figure 6: Example of an infinite evolving graph where edges (x_0, x_1) , (x_1, x_2) , (x_2, x_3) , (x_3, x_4) , and (x_5, x_6) occur periodically, but where edges (x_4, x_5) and (x_6, x_7) stop occurring after a while. The red travel corresponds to $slowT(x_7)$.

We now show that if a solution exists to the ODOC problem, then one solution do not uses temporal edges occurring after the $slow$ -travel. Formally, let \mathcal{E} be the set of temporal edges used by the $slow$ -travel or occurring before.

$$\mathcal{E} = \bigcup_{i \in [1, n]} \{((x_{i-1}, x_i), t) \mid (x_{i-1}, x_i) \in E_t \wedge t \leq t_i^s\}$$

Lemma 5. *Assuming the cost function is in \mathcal{UF} , let T be a solution of the ODOC problem, then there exists a travel T' that has a cost equal or smaller than the cost of T and such that its temporal edges are in \mathcal{E} . Moreover, if T is \mathcal{H} -history constrained, then T' is too.*

Proof. We show the result by induction. Assume the result is true for temporal edges up to node x_{i-1} . Let T be a solution of the ODOC problem with the desired property up to node x_{i-1} , i.e., such that

$$T = T_1 \oplus ((x_{i-1}, t_{i-1}), (x_{i-1}, t_i), (x_i, t_i), (x_i, t_{i+1})) \oplus T_2$$

where all temporal edges of T_1 are in \mathcal{E} and $((x_{i-1}, x_i), t_i) \notin \mathcal{E}$. We have that $t_{i-1} \leq t_{i-1}^s$ and $t_i > t_i^s$.

If $t_i \leq t_i^s$ (which occurs only when $t_i^s = \max(\{t \mid (x_{j-1}, x_j) \in E_t\})$), then T has the desired property up to node x_i (its temporal edges up to x_i are in \mathcal{E} and, obviously, its cost is not greater than the cost of T).

Otherwise, $t_i^s = \min(\{t \mid t \geq t_{i-1}^s \wedge (x_{i-1}, x_i) \in E_t\})$, which implies $t_{i-1} \leq t_{i-1}^s \leq t_i^s < t_i$ so the temporal edge of T between x_{i-1} and x_i could occur at time t_i^s instead of t_i without increasing the cost. Formally, the following travel has its temporal edges in \mathcal{E} up to node x_i :

$$T' = T_1 \oplus ((x_{i-1}, t_{i-1}), (x_{i-1}, t_i^s), (x_i, t_i^s), (x_i, t_{i+1})) \oplus T_2$$

Moreover

$$\begin{aligned} cost(T) - cost(T') &= f(t_{i-1} - t_i) + f(t_i - t_{i+1}) - f(t_{i-1} - t_i^s) - f(t_i^s - t_{i+1}) \\ &= f(t_i - t_{i+1}) - f(t_i^s - t_{i+1}) \geq 0 \end{aligned}$$

because $t_{i-1} \leq t_i$ implies $f(t_{i-1} - t_i) = f(t_{i-1} - t_i^s) = 0$ and $t_i^s \leq t_i$ implies $f(t_i - t_{i+1}) \leq f(t_i^s - t_{i+1})$. So T' has the desired property up to node x_i . Moreover, if T is \mathcal{H} -history-constrained, since $t_{i-1} \leq t_i^s \leq t_i$, then T' is too.

The base case being trivial (travels do not have temporal edges up to node x_0 because x_0 is the departure node), then the results is true up to the last node x_n . \square

Corollary 1. *Assuming the cost function is in \mathcal{UF} , if there is no solution to the ODOC problem having all its temporal edges in \mathcal{E} , then there is no solution to the ODOC problem.*

Theorem 4. *If the cost function \mathfrak{f} is in \mathcal{UF} , then Algorithm 1 terminates in $O(|\mathcal{E}|^2)$. If no solution is found, then there is no solution to the ODOC problem.*

Proof. When a pair (i, t) is extracted, we saw that a travel exists to x_i arriving at time t with cost $\text{nodeCost}[i, t]$ and the aim in the following lines is to extend this travel with temporal edges to create possibly optimal paths towards x_{i+1} . One can see that at most one such temporal edge occurs at time t' greater or equal to t . If $t \leq t_i^s$ (using the definition of t_i^s given in Equation 5), then clearly $t' \leq t_{i+1}^s$. By induction, all the temporal edges that we consider and extract are in \mathcal{E} . Moreover, when a pair (i, t) is extracted, the inner *for* loop iterates over the temporal edges that occur before t , which is less than $|\mathcal{E}|$ times. Retrieving the minimum from nodeCost is also in $O(|\mathcal{E}|)$. In the end, the complexity is $O(|\mathcal{E}|^2)$.

Using Corollary 1, we know that if Algorithm 1 does not find a solution, then no solution exists. \square

5 Offline \mathcal{H} -history-constrained ODOC Algorithm

Section 4 made the assumption that a given agent was able to go back to *any* previous snapshot of the network. However, this hypothesis cannot hold with entry-level BTT devices (and even some mid-tier devices), that are only able to send the agent back a given number of snapshots from a particular node, due to cutting costs in a context of memory chip shortage. Hence, we consider in this section that \mathcal{H} denotes the maximum number of time instants one agent can travel back to. In more detail, once an agent reaches time instant t , it cannot go back to time instants $t' < t - \mathcal{H}$, even after multiple jumps. **In this section, it is important to notice that the capability of BTT devices does not depend on the time when the agent uses it but rather on the largest time reached by the agent.** Again, we consider that graph G has $n + 1$ nodes denoted x_0, x_1, \dots, x_n .

We present Algorithm 2 that solve the \mathcal{H} -history-constrained ODOC problem. The algorithm uses dynamic programming to store intermediary results. At each iteration, we update the optimal cost based on the best cost of previous nodes. For each node x_i and time t we need to store the best cost depending on the maximum time reached by the agent.

Lemma 6. *If the cost function \mathfrak{f} is in \mathcal{UF} , then, in Algorithm 2, for any tuple (i, t, t_M) , if $c[i, t - h, t] < \infty$, then $c[i, t - h, t]$ is the cost of an \mathcal{H} -history-constrained cost optimal travel towards node x_i , that arrives before or at time instant $t - h$, and never reaches a time instant greater than t .*

Proof. We show by induction that, at each iteration, $c[i, t - h, t]$ is set to the cost of an \mathcal{H} -history-constrained cost optimal travel towards node x_i , that arrives before or at time instant $t - h$, and never reaches a time instant greater than t .

It is clear that if $i = 0$, regardless of t and h , the property is true, due to the initialization Line 1.

Now assume that the property is true for any tuple smaller than $(i, t - h, t)$ (using the lexicographical order). Assume for the sake of contradiction that there exists a backward-cost optimal travel T towards node i , arriving before or at time $t - h$ that is \mathcal{H} -absolute-constrained, and never reaches a round greater than t , such that $\text{cost}(T) = c' < c[i, t - h, t]$. We now consider three cases, corresponding to the three members appearing in the min function in Line 5, and show that each case implies a contradiction.

If T never reaches a time instant greater than $t - 1$ (which implies in particular that $t \geq 1$), then, using the induction hypothesis on T , we have that $c[i, t - h, t - 1] \leq c'$. But by definition (Line 5), we also have $c[i, t - h, t] \leq c[i, t - h, t - 1]$, which is a contradiction (that would mean $c[i, t - h, t] \leq c' < c[i, t - h, t]$).

If $((x_{i-1}, t_{i-1}, (x_i, t_{i-1}), (x_i, t - h))$ is the end of T and $t_{i-1} < t - h$. Then T arrives at x_i before time $t - h$, hence, by induction hypothesis, $c[i, t - h - 1, t] \leq c'$. But by definition $c[i, t - h, t] \leq c[i, t - h - 1, t]$, which is a contradiction.

If $((x_{i-1}, t_{i-1}, (x_i, t_{i-1}), (x_i, t - h))$ is the end of T , and $t_{i-1} \in [t - h, t]$. Let T' be such that $T = T' \oplus ((x_{i-1}, t_{i-1}, (x_i, t_{i-1}), (x_i, t - h))$. So we have that T' is an \mathcal{H} -absolute-constrained travel to x_{i-1}

arriving before or at time t_{i-1} and that never reaches a time instant greater than t . Hence, by induction hypothesis on T' , we get:

$$c[i-1, t_{i-1}, t] \leq \text{cost}(T') = \text{cost}(T) - f(t_{i-1} - (t - h)).$$

But by definition (third case of Line 5), we have

$$c[i, t - h, t] \leq c[i-1, t_{i-1}, t] + f(t_{i-1} - (t - h)) = c',$$

which is a contradiction.

Otherwise $((x_{i-1}, t-h), (x_i, t-h))$ **is the end of T** . This case is treated similarly as before, because we can append $(x_i, t-h)$ without changing the travel, and now it matches the previous case. \square

Lemma 7. *Assuming c is constructed by Algorithm 2, then*

$\text{ExtractHistoryConstrainedTravel}(i, t, t + \mathcal{H}, c)$ returns, if it exists, an \mathcal{H} -history-constrained travel with optimal cost to node x_i arriving at time t .

Proof. Recall that, for any tuple (i, t, t_M) , if $c[i, t, t_M] < \infty$, then $c[i, t, t_M]$ is the cost of an \mathcal{H} -history-constrained cost optimal travel towards node i , that arrives before or at time instant t , and never reaches a time instant greater than t_M .

Then, observe that an \mathcal{H} -history-constrained travel with optimal cost to node x_i arriving at time t never reaches a time greater than $t + \mathcal{H}$, so it is enough to prove that the travel returned by $\text{ExtractHistoryConstrainedTravel}(i, t, t_M, c)$ has cost $c[i, t, t_M]$, never reaches a time greater than t_M and is \mathcal{H} -history-constrained. We prove this result by induction. If $i = 0$ the result is clear. Otherwise, we have by construction of c :

$$c[i-1, t', t'_M] + f(t' - t) = c[i, t, t_M] \quad (6)$$

for some $t' \leq t'_M \leq t_M$.

Since $c[i-1, t', t'_M]$ is finite, then, by induction hypothesis, $\text{ExtractHistoryConstrainedTravel}(i-1, t', t'_M)$ is an \mathcal{H} -history-constrained travel with optimal cost $c[i-1, t', t'_M]$ and never reaches time greater than t'_M . Hence, the returned path never reaches time t_M (because $t'_M \leq t_M$, and its cost is

$$\begin{aligned} \text{cost}(\text{ExtractHistoryConstrainedTravel}(i-1, t', t'_M) \oplus ((x_{i-1}, t'), (x_i, t'), (x_i, t))) \\ = c(i-1, t', t'_M) + f(t' - t) \end{aligned}$$

This is equal to $c(i, t, t_M)$, using Equation 6. \square

Theorem 5. *If the cost function f is in \mathcal{UF} , then Algorithm 2 solves the \mathcal{H} -history-constrained ODOC problem and has $O(n\mathcal{H}(t_{\min} + \mathcal{H}))$ complexity, with t_{\min} the delay of a solution.*

Proof. Using Lemma 6 and Lemma 7 we know that Algorithm 2 returns a solution to the \mathcal{H} -history-constrained ODOC problem.

Regarding the complexity, we exit the main loop after reaching $t = t_{\min} + \mathcal{H}$, and the two inner loops have complexity $O(n\mathcal{H})$, so the complexity of the algorithm is in $O(n\mathcal{H}(t_{\min} + \mathcal{H}))$ \square

Algorithm 2: Offline \mathcal{H} -history-constrained ODOC Algorithm

```
/*  $c[i, t-h, t]$  stores the cost of a cost optimal travel to node  $x_i$ , arriving before or at time  $t-h$ ,
   that is  $\mathcal{H}$ -history-constrained, and never reaches a time instant greater than  $t$ . */
1  $c[*] \leftarrow \infty$ ;  $c[0, *] \leftarrow 0$ ;
2 for  $t = 0, 1, 2, \dots$  do
3   for  $i = 1, 2, \dots, n$  do
4     for  $h = \mathcal{H}, \mathcal{H}-1, \dots, 0$  do
5        $c[i, t-h, t] \leftarrow \min \left( \begin{array}{l} c[i, t-h, t-1] \\ c[i, t-h-1, t] \\ \min_{\substack{t' \in [t-h, t] \\ e_{i-1} \in E_{t'}}} (c[i-1, t', t] + f(t' - (t-h))) \end{array} \right)$ ;
       /* for simplicity, we assume that  $c[i, t-h, t]$  remains equal to  $\infty$  if  $t-h < 0$  */
6     end
7   end
8   if the minimum time instant  $t_{\min}$  such that  $c[n, t_{\min}, t_{\min} + \mathcal{H}] < \infty$  exists then
9     return ExtractHistoryConstrainedTravel( $n, t_{\min}, t_{\min} + \mathcal{H}, c$ );
10  end
11 end
```

Algorithm ExtractHistoryConstrainedTravel: Extract an \mathcal{H} -history-constrained ODOC travel to the given destination

```
input:  $i \in [0, n]$ ,  $t \in \mathbb{N}$ ,  $t_M \in \mathbb{N}$ ,  $c$ 
1 if  $i = 0$  then
2   return  $((x_0, 0), (x_0, t))$ 
3 end
4 Let  $t'$  and  $t'_M$  such that  $(x_{i-1}, x_i) \in E_{t'}$  and  $c[i-1, t', t'_M] + f(t' - t) = c[i-1, t, t_M]$ 
   /*  $t'$  and  $t'_M$  exist by construction of  $c$ .  $c[i-1, t, t_M]$  either (a) equals  $c[i-1, t', t'_M]$  with some  $t' < t$ 
   if the minimum at Line 5 is obtained by one of the first two arguments (in this case
    $f(t' - t) = 0$ ), or (b) equals  $c[i-1, t', t_M] + f(t' - t)$  with some  $t' \geq t$  (in this case  $t'_M = t_M$ ). In both
   cases, edge  $((x_{i-1}, x_i), t')$  exists. */
5 return ExtractHistoryConstrainedTravel( $i-1, t', t'_M, c$ )  $\oplus ((x_{i-1}, t'), (x_i, t'), (x_i, t))$ 
```

6 Online ODOC Algorithm with Optimal Competitive Ratio

In this section, we consider the online version of ODOC problem, *i.e.*, with an unknown schedule. In other words, the future of the evolving graph is unknown to the algorithm: at a time t , only the snapshots at time instants $t' \leq t$ are known.

Because it does not know the entirety of the input data, an online algorithm does not always return an optimal solution. Here in particular, an agent must wait in order to discover potentially optimal (or better) solutions. But if none appear, the latter must go back in time to the best known solutions, which in turn increases the cost of the travel. Thus, online algorithms are usually analyzed through their *competitive ratio*, *i.e.*, the ratio between the cost of the travel found by the online algorithm and the (optimal) cost Opt of the travel found by its offline counterpart. An algorithm has competitive ratio ρ if in any graph G , the solution T returned by the algorithm satisfies $cost(T)/Opt \leq \rho$.

In this section, we present the online Algorithm 3 for the online ODOC problem. We show that Algorithm 3 has an optimal competitive ratio.

We assume in this section that the cost function is the identity: $f : x \mapsto x$. Note that f is in \mathcal{UF} , so an optimal solution by an offline algorithm exists (see previous sections). We discuss later why relaxing this assumption is challenging (but BTT devices with such cost function remain in majority until 2052).

Lemma 8. *Assuming $f : x \mapsto x$, if the future is unknown, there exist no online algorithm with competitive*




Figure 7: Definition of the evolving graphs G^i , with 9 nodes ($n = 8$). The blue travel T_8 has a backward-cost of 8. The red travel T_i has a backward cost of i .

ratio $2 - \varepsilon$, with $\varepsilon > 0$, for the ODOC problem.

Proof. Assume for the sake of contradiction that algorithm A solves the ODOC problem and has a competitive ratio of $2 - \varepsilon$, with $\varepsilon > 0$. Let n be an even integer greater than $\frac{5}{\varepsilon}$. For any $i > 3$, let G^i be an evolving graph with $n + 1$ nodes defined in the following way:

- $G^i(0)$ is the graph where half of the edges are present:

$$E^i(0) = \{\{x_k, x_{k+1}\} \mid k \in [0, n] \wedge k \equiv 1 \pmod{2}\}.$$

- $G^i(2)$ is the graph where the other half of the edges are present:

$$E^i(1) = \{\{x_k, x_{k+1}\} \mid k \in [0, n] \wedge k \equiv 0 \pmod{2}\}.$$

- $G^i(i)$ is a line graph : $E^i(i) = \{\{x_k, x_{k+1}\} \mid k \in [0, n - 1]\}$.
- for all $j \notin \{0, 2, i\}$, $G^i(j)$ is a graph with no edge : $E^i(j) = \emptyset$.

It is clear that, in all such graphs G^i , there exists a travel, denoted by T_n , with backward-cost n , using the edges present at time 0 and 2 (the blue travel in Figure 7). In addition, there exists a travel, denoted T_i , of backward-cost i in the evolving graph G^i (the red travel in Figure 7).

If $i > n$, the optimal travel is T_n , and if $i < n$ the optimal travel is T_i .

Let us now run Algorithm A on the evolving graph G^{2n} . Clearly, the algorithm cannot wait until time instant $2n$ otherwise the backward cost would be at least $2n$, which is two times more than the backward cost of the optimal path T_n . This implies that Algorithm A cannot distinguish between (hence runs exactly in the same way in) graphs G^i , with $i \geq 2n$. Let t_{\max} be the maximum time instant reached by Algorithm A in G^{2n} . Then we can even say that A cannot distinguish between graphs G^i with $i > t_{\max}$.

Claim 1: In G^i , with $i > t_{\max}$, Algorithm A outputs a travel with a backward-cost of at least $n + t_{\max} - 2$

Proof of the Claim: The travel $T = A(G^i)$ that A outputs must contains the same temporal edges as T_n because those are the only edges that exists before time i (recall that $t_{\max} < i$). Let t_j be the time instant reached by Algorithm A at node x_j , for all $j = 0, \dots, n$. To move from node x_j to node x_{j+1} the travel T includes a backward trip of cost $\max(2, t_j)$, if $j \equiv 1 \pmod{2}$, and of cost $\max(0, t_j - 2)$, otherwise. Let $t_{j_{\max}} = t_{\max} = \max(t_j)$, we have that

$$\begin{aligned} \text{cost}(T) &= \sum_{j \in \text{Odd}(n)} \max(2, t_j) + \sum_{j \in \text{Even}(n)} \max(0, t_j - 2) \\ &\geq \begin{cases} 2|\text{Odd}(n)| + t_{\max} - 2 & \text{if } j_{\max} \text{ is odd} \\ 2(|\text{Odd}(n)| - 1) + t_{\max} & \text{if } j_{\max} \text{ is even} \end{cases} \end{aligned}$$

Where $\text{Even}(n)$, resp. $\text{Odd}(n)$, denotes the set of even, resp. odd, numbers smaller or equal to n . Since $|\text{Odd}(n)| = n/2$, we obtain in both case $\text{cost}(T) \geq n + t_{\max} - 2$

Claim 2: $t_{\max} \leq n - 4$

Proof of the Claim: Since algorithm A has a competitive ratio of $2 - \varepsilon$, then, if it runs in the evolving graph G^{2n} , it must return a path of backward-cost at most

$$(2 - \varepsilon) \text{Opt}(G^{2n}) = (2 - \varepsilon)n < 2n - 5$$

(recall that $n\varepsilon > 5$), so it cannot reach time instant $n - 3$. Indeed, if the algorithm waits until time instant $t_{\max} \geq n - 3$, then, using the previous claim, the backward-cost of the travel would be at least $n + n - 5$.

Now we run Algorithm A on graph $G^{t_{\max}+1}$. Using Claim 1, we know that A returns a travel of cost at least $n + t_{\max} - 2$. However, in $G^{t_{\max}+1}$, since $t_{\max} + 4 \leq n$ (Claim 2), the optimal travel is $T_{t_{\max}+1}$ having a cost of $t_{\max} + 1$. We obtain the following inequality:

$$\text{cost}(A(G^{t_{\max}+1})) \geq n + t_{\max} - 2 \geq t_{\max} + 4 + t_{\max} - 2 \geq 2(t_{\max} + 1) \geq 2\text{Opt}(G^{t_{\max}+1})$$

This contradicts the fact that A has competitive ratio of $2 - \varepsilon$. □

Algorithm 3: Online ODOC Algorithm

```

1  $c_{\max} \leftarrow \infty$ ;
2 while  $\text{time} < c_{\max}$  do
3   wait 1 time instant;
4   if there exists a travel  $T$ , starting at time 0, from node  $x_0$  to node  $x_n$  with cost  $c < c_{\max}$  then
5      $c_{\max} \leftarrow c$ ;
6      $T_{\max} \leftarrow T$ ;
7   end
8 end
9 go to the past at time 0, then send the agent using travel  $T_{\max}$ 

```

Theorem 6. Assuming $\mathfrak{f} : x \mapsto x$, Algorithm 3 is an online algorithm with a competitive ratio of 2 for the ODOC problem.

Proof. First, we prove that T_{\max} is the optimal offline travel. Indeed, the algorithm reached time c_{\max} so all the other travels that are not discovered by our algorithm require temporal edges appearing after time c_{\max} , so their backward-costs are greater than c_{\max} . Hence c_{\max} is the optimal backward-cost. When the algorithm terminates, the travel T that is returned is $((x_0, 0), (x_0, c_{\max}), (x_0, 0)) \oplus T_{\max}$. It has a backward-cost of $2c_{\max}$ (c_{\max} to go back to time 0 plus c_{\max} to use T_{\max}), which proves the Lemma. □

7 Conclusion

We presented the first solutions to the optimal delay optimal cost space-time constrained travel planning problem in dynamic line networks, and demonstrated that the problem can be solved in polynomial time, even in the case when backward time jumps can be made up to a constant, for any sensible pricing policy. By contrast, we showed that assuming an identity cost function, no online algorithm can exhibit a competitive ratio of less than two, and we present a very simple online algorithm with a competitive ratio of two. We conjecture our algorithms can be extended to arbitrary dynamic graphs, which is left for future work.

One obvious open question is to further study the impact of cost functions on the competitive ratio. Not all functions in \mathcal{UC} allow a finite competitive ratio:

Lemma 9. Assuming $f : x \mapsto c$, where c is a constant, if the future is unknown, there exist no algorithm with finite competitive ratio, for the ODOC problem.

Proof. The existence of a single backward-free travel (that is, a journey) to the destination implies a travel of optimal cost $Opt = c$ (travel to the destination and perform a single backward time travel to reach time 0), and this travel can be discovered arbitrarily far in the future. \square

Considering other functions in \mathcal{UC} with respect to online competitive ratio is left for future work.

References

- [1] Arnaud Casteigts, Paola Flocchini, Bernard Mans, and Nicola Santoro. Shortest, fastest, and foremost broadcast in dynamic networks. *Int. J. Found. Comput. Sci.*, 26(4):499–522, 2015.
- [2] Arnaud Casteigts, Paola Flocchini, Walter Quattrociocchi, and Nicola Santoro. Time-varying graphs and dynamic networks. *Int. J. Parallel Emergent Distributed Syst.*, 27(5):387–408, 2012.
- [3] Arnaud Casteigts, Anne-Sophie Himmel, Hendrik Molter, and Philipp Zschoche. Finding temporal paths under waiting time constraints. *Algorithmica*, 83(9):2754–2802, 2021.
- [4] Arnaud Casteigts, Joseph G. Peters, and Jason Schoeters. Temporal cliques admit sparse spanners. *J. Comput. Syst. Sci.*, 121:1–17, 2021.
- [5] Shigang Chen and Klara Nahrstedt. An overview of qos routing for the next generation high-speed networks: Problems and solutions. *Network, IEEE*, 12:64 – 79, 12 1998.
- [6] Afonso Ferreira. On models and algorithms for dynamic communication networks: The case for evolving graphs. In *Quatrièmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications (ALGOTEL 2002)*, pages 155–161, Mèze, France, May 2002. INRIA Press.
- [7] Rosario G. Garroppo, Stefano Giordano, and Luca Tavanti. A survey on multi-constrained optimal path computation: Exact and approximate algorithms. *Computer Networks*, 54(17):3081–3107, Dec 2010.
- [8] Jochen W. Guck, Amaury Van Bemten, Martin Reisslein, and Wolfgang Kellerer. Unicast qos routing algorithms for sdn: A comprehensive survey and performance evaluation. *IEEE Communications Surveys & Tutorials*, 20(1):388–415, 2018.
- [9] Giuseppe Antonio Di Luna, Paola Flocchini, Giuseppe Prencipe, and Nicola Santoro. Black hole search in dynamic rings. In *41st IEEE International Conference on Distributed Computing Systems, ICDCS 2021, Washington DC, USA, July 7-10, 2021*, pages 987–997. IEEE, 2021.
- [10] George Pal. The time machine, 1960.
- [11] Anthony Russo and Joe Russo. Avengers: Endgame, 2019.
- [12] Krishnaiyan Thulasiraman, Subramanian Arumugam, Andreas Brandstädt, and Takao Nishizeki. *Handbook of graph theory, combinatorial optimization, and algorithms*. 2016.
- [13] Robert Zemeckis. Back to the future, 1985.