Theoretical Computer Science 941 (2023) 221–240

Contents lists available at ScienceDirect

Theoretical Computer Science

journal homepage: www.elsevier.com/locate/tcs

The influence of maximum (s, t)-cuts on the competitiveness of deterministic strategies for the Canadian Traveller Problem ☆

Pierre Bergé [a,*], Lou Salaün [b]

[a] IRIF, CNRS, Université de Paris, 75013 Paris, France
[b] Bell Labs, Nokia Paris-Saclay, 91620 Nozay, France

A R T I C L E   I N F O

Article history:
Received 22 April 2021
Received in revised form 8 November 2022
Accepted 11 November 2022
Available online 15 November 2022
Communicated by V. Bonifaci

Keywords:
Canadian Traveller Problem
Online algorithms
Maximum cuts

A B S T R A C T

We study the k-Canadian Traveller Problem, where the objective is to lead a traveller in an undirected weighted graph G from a source s to a target t, knowing that at most k edges of the graph are blocked and cannot be traversed. The locations of blockages are unknown at the beginning of the walk but a blocked edge is revealed when the traveller visits one of its endpoints. There exist graphs for which the competitive ratio of any deterministic strategies cannot be smaller than $2k+1$. Conversely, there exists a very simple strategy, REPOSITION, which achieves this ratio $2k+1$. It consists in successively traversing shortest (s, t)-paths and coming back to s when the traveller is blocked.

We refine this analysis by detecting families of graphs for which a smaller competitive ratio can be obtained. This paper produces a global analysis to understand the impact of the size of the maximum (s, t)-cuts of G on the competitiveness of deterministic strategies. We design deterministic strategies achieving a ratio $\rho k + O(\lambda)$, with $\rho < 2$, for two different cut parameters λ. In particular, we propose a strategy called DETOUR with a competitive ratio $\sqrt{2}k + O(\mu^E_{\max})$, where $\mu^E_{\max}$ is the size of the maximum edge (s, t)-cut. Another contribution is a strategy called BYPASS with a competitive ratio $2^{\frac{3}{4}}k + O(\lambda^V_{\max})$, where $\lambda^V_{\max}$ is the size of the maximum vertex (s, t)-cut of all subgraphs of G. This produces an efficient algorithm for outerplanar graphs, which verify $\lambda^V_{\max} \leq 2$.

1. Introduction

Before presenting our contributions, we begin with a summary of the literature related to our study.

1.1. Related work

The *k-Canadian Traveller Problem* (k-CTP) was defined by Papadimitriou and Yannakakis [16]. It models the travel through a graph where some obstacles may appear suddenly. Given an undirected weighted graph $G = (V, E, \omega)$ and two of its vertices $s, t \in V$, a traveller has to walk from s to t on graph G in the shortest way despite the existence of blocked edges

☆ This article belongs to Section A: Algorithms, automata, complexity and games, Edited by Paul Spirakis.

* Corresponding author.
E-mail addresses: berge@irif.fr (P. Bergé), lou.salaun@nokia-bell-labs.com (L. Salaün).

https://doi.org/10.1016/j.tcs.2022.11.017

$E_* \subsetneq E$. The traveller does not know which edges are blocked when he begins his walk. He discovers that an edge $e = (u, v)$ is blocked, *i.e.* belongs to E_*, when he visits one of its endpoints u or v. Parameter k is an upper bound on the number of blocked edges: $|E_*| \le k$. The k-CTP is PSPACE-complete [2,16]. Variants of the k-CTP have been studied, where either edges are blocked with a certain probability [1,2,7] or there are multiple travellers [4,17] or we seek the shortest tour [15].

Graph $G = (V, E, \omega)$ is made up of $n = |V|$ vertices and $m = |E|$ edges. Edge weights are given by the function $\omega : E \to \mathbb{Q}^+$. Our objective is to make the traveller reach target t with a minimum cost (also called distance), which is the sum of the weights of the traversed edges. A pair (G, E_*) is called a *road map*. All the road maps considered are feasible: there exists an (s,t)-path in $G \backslash E_*$, the graph G deprived of the obstructed edges E_*. In other words, there always is a way to reach target t from source s despite the blockages.

A solution to the k-CTP is an online algorithm, called a *strategy*, which guides the traveller through his walk on the graph. Its quality can be assessed with competitive analysis [8]. Roughly speaking, the competitive ratio is the quotient between the distance actually traversed by the traveller and the distance he would have traversed, knowing which edges are blocked before beginning his walk. Westphal [19] proved that no deterministic strategy achieves a competitive ratio better than $2k + 1$. Said differently, for any deterministic strategy A, there is at least one k-CTP road map for which the competitive ratio of A is at least $2k + 1$. Two strategies proposed in the literature reach this optimal ratio: REPOSITION [19] and COMPARISON [20]. Here is how strategy REPOSITION works in simple terms. It consists in trying to traverse the shortest (s,t)-path (exploration phase) of G deprived of the blockages revealed: if a blocked edge is found on this path and, hence, prevents us from reaching t, we update the set of blocked edges discovered, go back to s (backtracking phase) and restart the process until we reach t.

Randomised strategies, *i.e.* strategies in which choices of direction depend on a random draw, were also studied. Westphal [19] proved that there is no randomised strategy achieving a ratio lower than $k + 1$. Bender et al. [3] studied graphs composed only of vertex-disjoint (s,t)-paths and proposed a polynomial-time strategy of ratio $k + 1$. A slight revision of this strategy is reported in [18]. Demaine et al. proposed a polynomial-time randomised strategy improving the optimal deterministic ratio on general graphs by a $o(1)$ factor [10]. To the best of our knowledge, there is no polynomial-time randomised strategy achieving a competitive ratio $\rho k + O(1)$, with $\rho < 2$, on general graphs. Such a strategy would not be memoryless [5]. Finally, Karger and Nikolova [14] studied the case of trees for the stochastic version of the CTP.

1.2. Motivation

A natural question is whether the ratio $2k + 1$ can be improved on certain instances. Our objective is to determine families of graphs for which the competitive ratio of deterministic strategies could attain $\rho k + O(1)$ with $\rho < 2$. We quickly observed that without fixing a condition on the weight function, most well-known families of graphs (bipartite, planar, chordal,...) admit worst-case instances for which the competitive ratio $2k + 1$ is optimal. Nevertheless, we pursued looking for a class of graphs without weight restriction on which we could design a deterministic strategy with ratio $\rho k + O(1)$, $\rho < 2$.

In this article, we study the impact of the minimal (s,t)-cuts size on the competitiveness of deterministic strategies. An (s,t)-cut is a set of edges or vertices which separates s from t when it is withdrawn from G, and it is *minimal* whenever none of its proper subset is an (s,t)-cut. We design strategies outperforming REPOSITION on graphs where k is larger than the size of the largest minimal (s,t)-cuts. We put in evidence a family of graphs which is, in our opinion, of great interest for the k-CTP as a competitive ratio of $\rho k + O(1)$, $\rho < 2$, can be achieved on it. Below, we list our results.

1.3. Results

To analyse the influence of minimal (s,t)-cuts on the competitive ratio of deterministic strategies, we define three parameters:

- the size of the largest minimal edge (s,t)-cut of G: $\mu^E_{\max}$,
- the size of the largest minimal vertex (s,t)-cut of G: $\mu^V_{\max}$,
- the maximum $\mu^V_{\max}$ over all subgraphs of G: $\lambda^V_{\max}$.

For a given graph G, these parameters satisfy $\mu^V_{\max} \le \lambda^V_{\max} \le \mu^E_{\max}$. For each of them, we study their impact on the competitive ratio of deterministic strategies.

First, we propose a strategy called DETOUR. Its competitive ratio when $k > \mu^E_{\max}$ is $\sqrt{2}(k - \mu^E_{\max}) + 2\mu^E_{\max} + 1$. As a consequence, the ratio of DETOUR can be written $\sqrt{2}k + O(\mu^E_{\max})$. When $k \le \mu^E_{\max}$, its competitive ratio is $2k + 1$, similar to REPOSITION. This contribution was already introduced in our conference paper [6] whereas all other results below are new.

Theorem 1. *For any graph satisfying $k > \mu^E_{max}$, there is a strategy, called* DETOUR*, which guarantees a competitive ratio $\sqrt{2}(k - \mu^E_{max}) + 2\mu^E_{max} + 1$. Moreover, for $k \le \mu^E_{max}$, the competitive ratio of* DETOUR *is at most $2k + 1$.*

	$\mu^V_{\max}$	$\lambda^V_{\max}$	$\mu^E_{\max}$
$k \leq$	$2k+1$	$2k+1$	$2k+1$
$k >$	$2k+1$	$2^{\frac{3}{4}}k + O(\lambda^V_{\max})$	$\sqrt{2}k + O(\mu^E_{\max})$

(a) Best competitive ratios obtained for different intervals of k.

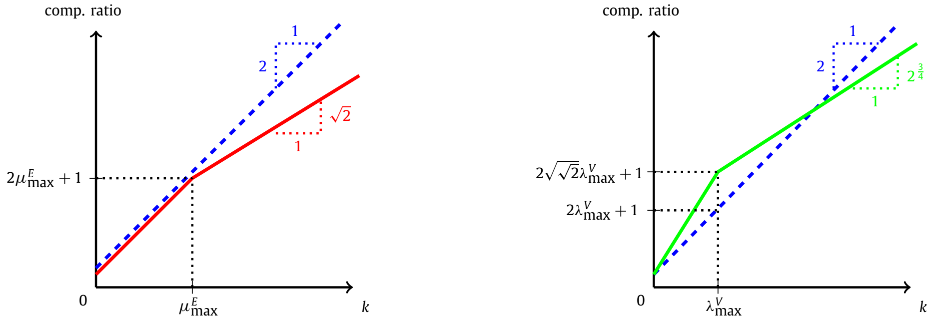

(b) Competitiveness of REPOSITION (blue) and DETOUR (red) versus k

(c) Competitiveness of REPOSITION (blue) and BYPASS (green) versus k

Fig. 1. A summary of our contributions. (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)

Second, we propose another strategy called BYPASS. Its competitive ratio when $k > \lambda^V_{\max}$ is $\frac{2}{\sqrt{\sqrt{2}}}(k - \lambda^V_{\max}) + 2\sqrt{\sqrt{2}}\lambda^V_{\max} + 1$. It can be written $2^{\frac{3}{4}}k + O(\lambda^V_{\max})$. When $k \leq \lambda^V_{\max}$, its competitive ratio is $2\sqrt{\sqrt{2}}k + 1$.

Theorem 2. *For any graph satisfying $k \geq \lambda^V_{max}$, there is a strategy, called* BYPASS, *which guarantees a competitive ratio $2^{\frac{3}{4}}(k - \lambda^V_{max}) + 2\sqrt{\sqrt{2}}\lambda^V_{max} + 1$.*

Both DETOUR and BYPASS are based on the same principles as REPOSITION. However, they contain extra instructions imposing the traveller to traverse suitable "detours" and "bypasses" (if they exist), instead of the original paths. Detours and bypasses are paths with certain properties and whose cost will guarantee the expected ratio.

In parallel with these results, we show that graphs with $\lambda^V_{\max} \geq 3$ have $K_{2,3}$ as a minor. In other words, any graph that does not have $K_{2,3}$ as a minor satisfies $\lambda^V_{\max} \leq 2$. This is true for outerplanar graphs [9] which are the graphs having a planar embedding where all vertices are on the outer face. Therefore, the competitive ratio of BYPASS strategy is $2^{\frac{3}{4}}k + O(1)$ on outerplanar graphs.

In summary, we identify two deterministic strategies which achieve ratios $\rho k + O(\lambda)$, $\rho < 2$, where λ is a cut parameter. The competitive slope of DETOUR, $\sqrt{2}$, is less than the one of BYPASS, $2^{\frac{3}{4}}$, however it applies to a smaller set of instances.

In addition to these two strategies, we prove that the parameter $\mu^V_{\max}$ has no impact on the competitiveness of deterministic strategies. Indeed, there are instances with $\mu^V_{\max} = 1$ for which the competitive ratio $2k+1$ is optimal. It is quite unexpected to see that the set of minimal vertex (s,t)-cuts of subgraphs of G allow us to outperform ratio $2k+1$, whereas the largest minimal vertex (s,t)-cut of G itself has no impact on it.

Main results. *Here is a summary of the main results of the paper.*

1. *No strategy achieves a ratio smaller than $2k+1$ on graphs satisfying $\mu^V_{max} = 1$. (Theorem 4, Section 3)*
2. *There is a strategy achieving a ratio $\sqrt{2}k + O(\mu^E_{max})$. (Theorem 1)*
3. *There is a strategy achieving a ratio $2^{\frac{3}{4}}k + O(\lambda^V_{max})$. (Theorem 2)*

Fig. 1 represents graphically the evolution of the competitive ratio of DETOUR and BYPASS as a function of k.

A natural question arises: is there a cut parameter λ, $\mu^V_{\max} \leq \lambda \leq \lambda^V_{\max}$ and a strategy such that its competitive ratio would be $\rho k + O(\lambda)$, $\rho < 2$? A positive answer would widen the set of instances for which we can design improved strategies. However, we prove in this paper that: considering any unweighted graph satisfying $k \leq \lambda^V_{\max}$, we can assign it a weight function such that the best competitive ratio achievable on it is at least $2k+1$. Therefore, such parameter λ cannot exist. This result highlights the fact that $\lambda^V_{\max}$ is the lowest cut parameter which has an impact on the competitive ratio when considering graphs with bounded cut sizes.

1.4. Organisation

In Section 2, we present the definitions and notations used in this article. A large part of this section relates to cuts, as we remind the proofs of many folklore results dealing with the relationship between edge and vertex (s,t)-cuts. Section 3 provides a detailed analysis of our results. It contains the proof that parameter $\mu_{\max}^{V}$ has no influence on the competitiveness of deterministic strategies. The following sections are dedicated to the two strategies we propose. DETOUR is described in Section 4 and we derive its competitive ratio. BYPASS is presented in Section 5 and its competitive analysis is inspired by the one of DETOUR. In brief, Section 4 (resp. Section 5) contains the proof of Theorem 1 (resp. Theorem 2).

2. Preliminaries

We begin with a reminder of some graph theory concepts to be used throughout this paper. Then, we present the definition of competitive ratio and some notions associated with cuts.

2.1. Graphs and paths

We work on undirected weighted graphs $G=(V,E,\omega)$, where $n=|V|$, $m=|E|$, and $\omega : E \to \mathbb{Q}^+$. The notation $V(G)$ and $E(G)$ may refer to V and E respectively if there is any risk of confusion with another graph. An *equal-weight* graph is such that all $\omega(e)$, $e \in E$, are equal.

A subgraph G' of G is a graph $G'=(V',E',\omega')$, where $V' \subseteq V$, $E' \subseteq E \cap (V' \times V')$, and $\omega' = \omega_{|E'}$. In this case, we use the notation $G' \subseteq G$. For any $U \subseteq V$, we denote by $E[U]$ the set of edges of G with two endpoints in U. Let $G[U]$ be the subgraph of G induced by U, $G[U]=(U,E[U])$. We denote by $G\backslash U$ the graph deprived of vertices in U: $G\backslash U = G[V\backslash U]$. Similarly, for any set of edges $E' \subseteq E$, the graph G deprived of E' is denoted by $G\backslash E' = \big(V, E\backslash E'\big)$.

We denote by $K_{p,q}$ the complete bipartite graph with an independent set of size p and another of size q. Graph H is a *minor* of graph G if it can be obtained after removing vertices and edges and/or contracting edges in G.

A *simple path* P is a sequence of pairwise different vertices $v_1 \cdot v_2 \cdots v_i \cdot v_{i+1} \cdots v_\ell$, with departure v_1 and arrival v_ℓ, where two successive vertices (v_i, v_{i+1}) are adjacent in G. All paths mentioned in this article are simple: they do not form cycles. To improve readability, we abuse notations: $v_1 \in P$ and $(v_1,v_2) \in P$ mean that vertex v_1 and edge (v_1,v_2) are on path P, respectively.

If vertices u and v belong to path P, then $P^{(u,v)}$ denotes the section of path P between vertices u and v. Any path is naturally associated with a direction, from the departure to the arrival. We define the *successor* of edge e in P as the edge positioned just after e in P. The *descendants* of e are all edges positioned after e in P. The *predecessor* of e is the edge positioned just before e and its *ancestors* are all edges located between the departure and e in P. These notions can also be defined naturally for vertices.

Graphs may contain several shortest (s,t)-paths. Our algorithms require to compute only one of these shortest (s,t)-paths on any graph in a deterministic way. To achieve this, a solution is to associate each vertex with a different identifier in $\{1,\dots,n\}$. If two paths have the same distance, we consider their sequence of vertices and compare the lexicographic order of their identifiers. Dijkstra's algorithm [12] is adapted to this extra criterion: for any vertex v, it stores the shortest path from the start point to v with the smallest lexicographic order. Whenever we refer to "the shortest (u,v)-path", for any vertices u and v, this procedure is executed.

2.2. Road maps

Let $G=(V,E,\omega)$ be a graph and E_* represents a set of blocked edges. We define below the concept of *road maps* which are the instances of the k-CTP problem.

Definition 1 *(Road maps).* A pair (G,E_*) is a *road map* if s and t remain connected in $G\backslash E_*$.

In other words, there must be an (s,t)-path in graph G deprived of the blocked edges E_*. As road maps have been defined, we can introduce formally the k-Canadian Traveller Problem.

Definition 2 *(k-CTP).*
Input: Graph $G=(V,E,\omega)$, vertices $s,t \in V$, and a set E_* of blocked edges which are unknown and such that (G,E_*) is a road map.
Objective: Traverse graph G from s to t with minimum distance.

The set of blocked edges E_* is a hidden input at the beginning of the walk. The traveller discovers whether an edge is blocked when visiting one of its endpoints.

With the k-CTP, we say a path is *blocked* if we know that it contains a blocked edge, *i.e.* it was already discovered by the traveller. We say a path is *open* if we are sure that it does not contain any blocked edge. That is, an open path is only

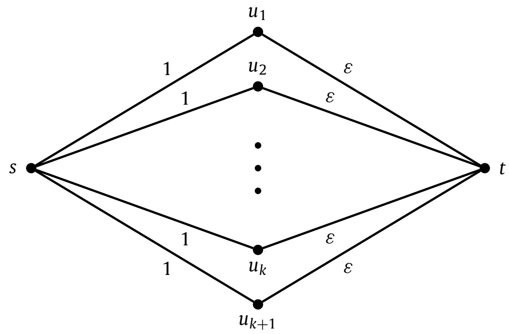

Fig. 2. Graph W_k, as defined in [19].

made up of edges which have been "revealed" (the traveller has visited one of their endpoints) and are not blocked. Finally, we say a path is *apparently open* if no blocked edge was discovered on it for now. However, it may contain a blocked edge which has not been discovered yet.

2.3. Competitive ratio

For any subset of blocked edges $E'_* \subseteq E_*$, $\omega_{\min}\left(G, E'_*\right)$ is the cost of the shortest (s,t)-path in graph $G\backslash E'_*$. Value $\omega_{\text{opt}} = \omega_{\min}(G, E_*)$ is the *optimal offline cost* for the road map (G, E_*). Concretely, this is the distance the traveller would have traversed if he had known the blockages in advance.

The competitive ratio is defined in [8]. We denote by $\omega_A(G, E_*)$ the distance traversed by the traveller guided by strategy A on graph G from source s to target t with blocked edges E_*. The competitive ratio $c_A(G, E_*)$ of A over a road map (G, E_*) is defined as $c_A(G, E_*) = \frac{\omega_A(G,E_*)}{\omega_{\text{opt}}}$. For $k \in \mathbb{N}$, the competitive ratio c_A of A for the k-CTP is:

$$c_A = \max_{\substack{\text{road map } (G,E_*) \\ |E_*| \le k}} c_A(G, E_*) . \tag{1}$$

Similarly, we say strategy A is $c_{A,\mathcal{F}}$-competitive for a family $\mathcal{F}$ of road maps if it is the maximum value $c_A(G, E_*)$ over all road maps (G, E_*) such that $(G, E_*) \in \mathcal{F}$.

Given a graph G, its family of road maps, denoted by $\mathcal{F}(G)$, is the set of road maps made up of graph G and any configuration E_*, $|E_*| \le k$, of blocked edges such that s and t remain connected in $G\backslash E_*$. Put formally,

$$\mathcal{F}(G) = \{(G, E_*) : |E_*| \le k \text{ and there is an open } (s,t)\text{-path in } G\backslash E_*\} .$$

Computing the competitive ratio of a strategy A over family $\mathcal{F}(G)$ allows us to determine the worst-case performance of A on graph G. All the strategies proposed in this paper are executed in polynomial time.

We remind the state of the art on the competitive ratio of deterministic strategies for k-CTP. We present here some known worst-case road maps, *i.e.* road maps on which the competitive ratio $2k + O(1)$ cannot be beaten. Westphal [19] identified, for any integer k, a family of road maps for which any deterministic strategy achieves at least ratio $2k+1$. These road maps are all based upon the graph W_k made up of $k+1$ disjoint (s,t)-paths, *i.e.* they do not share common edge or vertex except s and t. Each path has two edges: one with weight 1, another one with weight $\varepsilon \ll 1$. Family $\mathcal{F}(W_k)$ contains in particular the road maps composed of graph W_k and a set of k blocked edges among the $k+1$ edges of weight ε. Fig. 2 illustrates graphs W_k.

As the $k+1$ disjoint (s,t)-paths are indistinguishable, any deterministic strategy has no choice but selecting arbitrarily the first path traversed. In this situation, there exists a configuration of blockages such that the only open path is the last one visited. In this case, the total distance traversed is $2k+1+\varepsilon$. The optimal offline cost is $1+\varepsilon$. Making ε tend to 0 produces the bound $2k+1$.

Conversely, there are two strategies in the literature achieving the competitive ratio $2k+1$: REPOSITION [19] and COMPARISON [20]. The first one is very simple and we describe it briefly as a two-phase algorithm. The traveller traverses the shortest (s,t)-path of G (exploration phase). If he is blocked, then he goes back to s via the same path (backtracking phase). Then, he restarts this two-phase process on the updated graph $G\backslash E'_*$, which is graph G after removing the blocked edges discovered E'_*. Until he reaches t, the traveller traverses the shortest (s,t)-path in graph G deprived of the blocked edges discovered E'_* (exploration) and comes back to s if he is blocked (backtracking).

2.4. Cuts

We begin with the definition of edge cuts. A set $X \subseteq E$ is an edge (s,t)-cut if source s and target t are separated in graph G deprived of edges X, *i.e.* in $G\backslash X$. In other words, there is no (s,t)-path in $G\backslash X$. This definition can be extended to

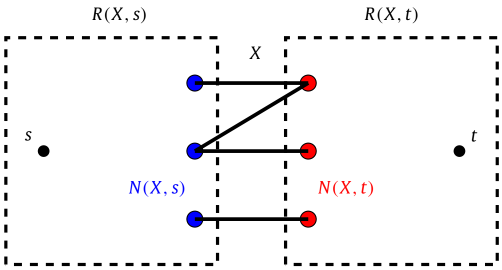

Fig. 3. A minimal edge (s,t)-cut X with its source/target sides.

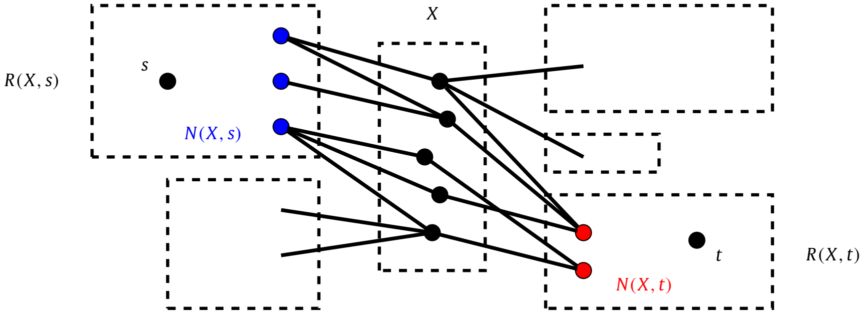

Fig. 4. Schematic view of some graph G regarding the minimal vertex (s,t)-cut X.

two sets of vertices instead of a single source and target: an edge (A, B)-cut is a set of edges separating the sets of vertices $A, B \subseteq V$, $A \cap B = \emptyset$. Then, we say that cut X is *minimal* if none of its proper subsets $X' \subsetneq X$ is an (s,t)-cut. Let $\mu^E_{\max}$ be the maximum cardinality of a minimal edge (s,t)-cut:

$$\mu^E_{\max} = \max_{\substack{X \subseteq E \text{ minimal} \\ (s,t)-\text{cut}}} |X| . \tag{2}$$

We say a *maximum* (s,t)*-cut* of a graph G is one of its largest minimal (s,t)-cuts. By definition, any (s,t)-cut X such that $|X| > \mu^E_{\max}$ is not minimal. We know that if X is a minimal (s,t)-cut, then graph $G \backslash X$ contains exactly two connected components [11,13] : $R(X,s)$[1] containing s and $R(X,t)$ containing t.

Let $N(X,s)$ (resp. $N(X,t)$) be the vertices of $R(X,s)$ (resp. $R(X,t)$) which are adjacent to cut X. Fig. 3 shows a schematic view of the notations $R(X,\cdot)$ and $N(X,\cdot)$ with a cut X made up of four edges. The vertices in $N(X,s)$ are in blue while the vertices in $N(X,t)$ are in red.

Given a vertex set $A \subseteq V$, we denote by $\delta(A)$ the set of edges with one endpoint in A and the other one outside. If A induces a connected subgraph containing s and not t, then $\delta(A)$ is a minimal edge (s,t)-cut. Given a minimal (s,t)-cut X of G, we have $X = \delta(R(X,s))$.

Now we define vertex cuts. A set $X \subseteq V$ is a vertex (s,t)-cut if s and t are separated in graph G deprived of vertices X, *i.e.* in $G \backslash X$. As for edges, a *minimal* vertex (s,t)-cut X is such that no subset $X' \subsetneq X$ is an (s,t)-cut. We denote by $\mu^V_{\max}$ the maximum cardinality of a minimal vertex (s,t)-cut:

$$\mu^V_{\max} = \max_{\substack{X \subseteq V \text{ minimal} \\ (s,t)-\text{cut}}} |X| .$$

A minimal vertex (s,t)-cut X may admit more than two connected components, in contrary to the edge case. Among them, $R(X,s)$ contains s while $R(X,t)$ contains t. Set $N(X,s)$ (resp. $N(X,t)$) contains the neighbours of vertices in X belonging to $R(X,s)$ (resp. $R(X,t)$).

Fig. 4 represents these notions in an example. Cutset X and the connected components of graph $G \backslash X$ are drawn. The vertices of $N(X,s)$ are represented in blue, the vertices of $N(X,t)$ in red. We draw all edges having exactly one endpoint in X. Observe through this example that, for minimal vertex (s,t)-cuts X, we may have $X \cup R(X,s) \cup R(X,t) \neq V$.

We define another cut parameter. Let $Y_{\max}$ be the largest minimal vertex (s,t)-cut we can find in all subgraphs of G. We denote by $\lambda^V_{\max}$ its size.

[1] If necessary, the graph considered is added to the notation: $R(X,s) = R_G(X,s)$.

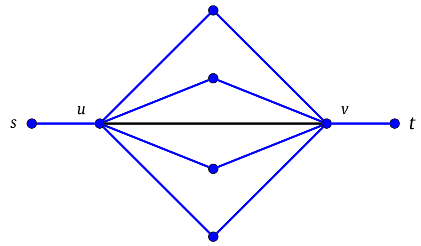

(a) Graph F_4

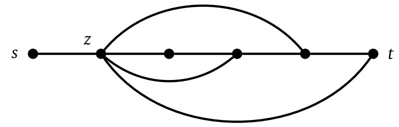

(b) Graph J_4

Fig. 5. Gaps between parameters $\mu^V_{\max}$, $\lambda^V_{\max}$, and $\mu^E_{\max}$: graphs F_4 and J_4.

$$\lambda^V_{\max} = \max_{\substack{Y\subseteq V \text{ minimal} \\ (s,t)-\text{cut of } G'\subseteq G}} |Y|.$$

If a subgraph G' does not contain either s or t, it does not admit a minimal (s,t)-cut and the parameters $\mu^E_{\max}$ and $\mu^V_{\max}$ of this subgraph are zero. The same conclusion holds if s and t are separated in G'.

Inequalities exist between the three cut parameters we introduced. We could define similarly $\lambda^E_{\max}$ as the largest minimal edge (s,t)-cut among all subgraphs $G' \subseteq G$. However, one can observe that, on any graph G, $\lambda^E_{\max} = \mu^E_{\max}$. This assertion stands as a folklore result on cuts. Indeed, any minimal edge (s,t)-cut of a subgraph of G which does not separate s and t in G can be completed with extra edges to form a minimal edge (s,t)-cut of G. The detailed proof of this statement is left to the reader.

We present now an inequality involving the three parameters $\mu^V_{\max}$, $\lambda^V_{\max}$ and $\mu^E_{\max}$. Observe that the equality established above does not hold for vertex cuts.

Lemma 1. *$\mu^V_{max} \le \lambda^V_{max} \le \mu^E_{max}$.*

Proof. As G is its own subgraph, we have $\mu^V_{\max} \le \lambda^V_{\max}$ by definition. Let $Y_{\max}$ be the largest minimal vertex (s,t)-cut in a subgraph G' of G. We denote by R'_s (resp. R'_t) the source (resp. target) side of $Y_{\max}$ in G'.

Let X be the set containing all edges which have one endpoint in $Y_{\max}$ and one in R'_t. Set X is an edge (s,t)-cut in G': as any (s,t)-path must pass through cut $Y_{\max}$, it necessarily traverses one edge of X to reach t. It is also minimal because if we re-open one edge of X connecting $Y_{\max}$ and R'_t, then there exists an open (s,t)-path in G' passing through this edge.

For any vertex $v \in Y_{\max}$, there is at least one edge incident to v with another endpoint in R'_t. Otherwise, set $Y_{\max}\setminus\{v\}$ would be an (s,t)-cut: this is a contradiction with the minimality of $Y_{\max}$. As a consequence, $|Y_{\max}| \le |X|$ and $\lambda^V_{\max} \le \lambda^E_{\max} = \mu^E_{\max}$. □

The inequalities provided in Lemma 1 can be strict. We can identify two families of graphs, say F_n and J_n, $n \in \mathbb{N}$ such that:

(i) $\mu^V_{\max}(F_n) = 1$, $\lambda^V_{\max}(F_n) = n$ and $\mu^E_{\max}(F_n) = n+1$,
(ii) $\mu^V_{\max}(J_n) = \lambda^V_{\max}(J_n) = 1$ and $\mu^E_{\max}(J_n) = n$.

Points (i) and (ii) show that these three parameters can be different. Furthermore, the gap between them can be as large as we want. Fig. 5 represents graphs F_4 and J_4. For F_n, value n is the number of (u,v)-paths of length 2. If we withdraw the edge (u,v), the size of the largest minimal vertex (s,t)-cut of the subgraph obtained (drawn in blue) is n, while it is 1 for the entire graph. For J_n, value n is the degree of vertex z minus one. The edges incident to z form a minimal (s,t)-cut of size n, while the size of the largest minimal vertex (s,t)-cut of any subgraph of J_n is 1.

In the remainder of this paper, we analyse the impact of these three parameters on the competitiveness of deterministic strategies for the k-CTP.

3. A cut-based competitive analysis: state of the art and first observations

In this section, we explain why the concept of minimal (s,t)-cut is critical to the competitiveness of deterministic strategies for the k-CTP. In particular, we show the impact of parameters $\mu^E_{\max}$ and $\lambda^V_{\max}$ defined earlier.

We prove that parameter $\mu^V_{\max}$ has no influence on the competitive ratio: put formally, for any $k \ge 1$, there are graphs with $\mu^V_{\max} = 1$ and for which the best competitive ratio obtained by any deterministic strategy is no less than $2k+1$. In other words, the lower bound $2k+1$ cannot be reduced even if the input graph has small maximum vertex (s,t)-cuts. Fortunately, we also design two strategies, called DETOUR and BYPASS which admit respectively a competitive ratio $\sqrt{2}k + O(\mu^E_{\max})$ and $2^{\frac{3}{4}}k + O(\lambda^V_{\max})$. So, the lower bound $2k+1$ can be reduced when parameters $\mu^E_{\max}$ and $\lambda^V_{\max}$ are small compared to k.

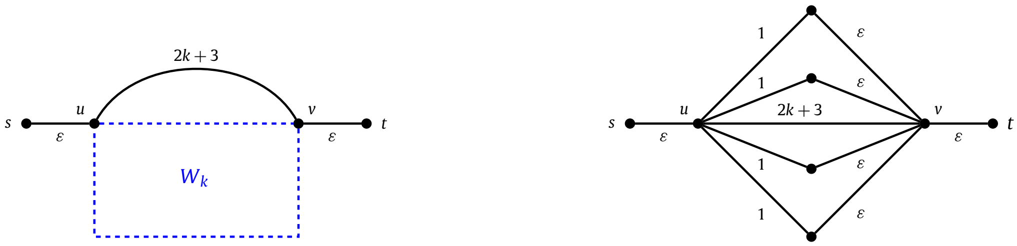

(a) Graph W_k^* (b) Graph W_3^*

Fig. 6. The structure of W_k^* and graph W_3^* as an example.

Graphs W_k satisfy the following inequality: $k < \mu_{\max}^V = \lambda_{\max}^V = \mu_{\max}^E = k+1$. The result of [19] can be presented in those terms.

Theorem 3 *(Competitive ratio of deterministic strategies [19]). For any integer k, there is a graph G such that the best competitive ratio achievable by any deterministic strategy on it is* $2k+1$*. Moreover, it satisfies* $k < \lambda$*, for any* $\lambda \in \{\mu_{max}^V, \lambda_{max}^V, \mu_{max}^E\}$*.*

A natural question is whether a better performance can be obtained for graphs fulfilling $\lambda \leq k$, for some $\lambda \in \{\mu_{\max}^V, \lambda_{\max}^V, \mu_{\max}^E\}$. In particular, we wonder if a smaller competitive bound could be reached for graphs with constant λ. This is all the more interesting in that such criterion is fully "topological", it does not involve the weights of graph G. Hence, we focus on the competitiveness of deterministic strategies as a function of the cut parameters we introduced.

Let us begin with parameter $\mu_{\max}^V$: we show a family of graphs with $\mu_{\max}^V = 1$ and for which the ratio $2k+1$ cannot be improved.

Theorem 4. *There is a family of graphs* W_k^* *with* $\mu_{max}^V = 1$ *on which any deterministic strategy is at least* $(2k+1)$*-competitive.*

Proof. First, we describe graph W_k^* which is an extension of the graph W_k. Graph W_k^* without weights is the graph F_k already introduced in Section 2.4 (see Fig. 5a). We remind that $\mu_{\max}^V(F_k) = 1$. We fix the weights of edges (s,u) and (v,t) equal to ε. The weight of edge (u,v) is $2k+3$. The weights of the two-edge (u,v)-paths are defined as in graph W_k: the first edge has weight 1, the second one has weight ε (Figs. 6a and 6b).

Second, we analyse the competitive ratio of strategies on the weighted graph W_k^*. Considering $\varepsilon \ll 1$, the cost of the (s,t)-path passing through (u,v) is larger than the distance needed for the return crossing of all the other (s,t)-paths which each have a cost $1+3\varepsilon$. So, the most competitive way to traverse graph W_k^* is to handle first the paths passing through the subgraph W_k. Edges (s,u) and (v,t) are necessarily open, otherwise the road map would not be feasible. The best achievable competitive ratio for graph W_k is $2k+1+O(\varepsilon)$ if one of its path is open. Assume one (u,v)-path in W_k is open (then it is the optimal offline path if concatenated with edges (s,u) and (v,t)). Any deterministic strategy may traverse the open path last, as in the original graph W_k. As the optimal offline cost is $1+O(\varepsilon)$, making ε tend to zero produces the lower bound $2k+1$. $\square$

We see in this proof that adding edge (u,v) does not improve the competitiveness of deterministic strategies on graph W_k because its weight is too large. However, it decreases the size of the maximum vertex (s,t)-cut to 1. This shows that parameter $\mu_{\max}^V$ is not relevant to identify families of graphs for which a better ratio can be obtained.

For parameter $\mu_{\max}^E$, there is a strategy with competitive ratio $\rho k + O(\mu_{\max}^E)$, $\rho < 2$, called DETOUR (Theorem 1). Its competitive slope is $\rho = \sqrt{2}$.

DETOUR refines REPOSITION: it can also be seen as a two-phase algorithm. As with REPOSITION, the traveller starts by traversing the shortest (s,t)-path of the current discovered graph $G \backslash E'_*$ (exploration phase). However, during the backtracking phase (renamed *detour-backtracking*) instead of going back to s directly, the traveller verifies, on his way back, whether certain "detours" exist to reach t. Detours are short paths (to be defined formally) from his current position to t.

DETOUR not only provides a competitive ratio $\sqrt{2}k + O(\mu_{\max}^E)$ for graphs with $k \geq \mu_{\max}^E$ but also ensures the optimal competitive ratio $2k+1$ for general graphs otherwise. In particular, its ratio grows in $\sqrt{2}k$ when k is larger than parameter $\mu_{\max}^E$. As a consequence, DETOUR offers, for now, better guarantees than the existing strategies, REPOSITION and COMPARISON, for which the only known worst-case ratio is $2k+1$ in all cases. Section 4 is dedicated to the description and the competitive analysis of DETOUR.

We also propose a strategy, called BYPASS, benefiting from small values of $\lambda_{\max}^V$. Its competitive slope is $\rho = \frac{2}{\sqrt[4]{2}} = 2^{\frac{3}{4}}$ (Theorem 2).

BYPASS is in fact DETOUR with an extra instruction; instead of systematically traversing the shortest (s,t)-path during the exploration phase, it may select "bypasses". Bypasses are (s,t)-paths which are not much longer than the shortest one and satisfying a certain distance property. Section 5 is dedicated to the description and the competitive analysis of BYPASS.

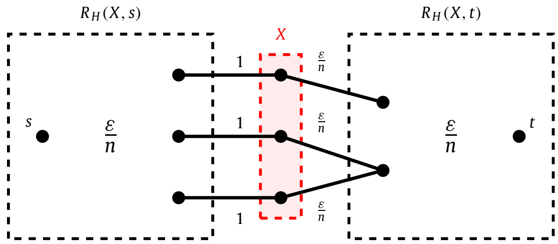

Fig. 7. Structure of the graph obtained with weight function ω_ε.

In contrast with DETOUR, strategy BYPASS does not guarantee a ratio of at most $2k+1$ for graphs satisfying $k < \lambda^V_{\max}$. Its competitive ratio can potentially reach $2\sqrt{\sqrt{2}k}+1$ on certain graphs.

On one hand, this result is stronger than Theorem 1. Indeed, strategy BYPASS expands the family of graphs on which we can apply a strategy achieving a competitive ratio $\rho k + O(\lambda)$, $\rho < 2$ and cut parameter λ. On the other hand, the competitive slope of BYPASS is larger than the one of DETOUR.

In the following, we focus on the graphs fulfilling $\lambda^V_{\max} \leq 2$. BYPASS guarantees a competitive ratio $2^{\frac{3}{4}}k + O(1)$ on them. We prove that the graphs which do not admit $K_{2,3}$ as a minor satisfy $\lambda^V_{\max} \leq 2$ and, therefore, benefits from the competitiveness of BYPASS.

Lemma 2. *If $\lambda^V_{max} \geq 3$, then $K_{2,3}$ is a minor of G.*

Proof. Let $\lambda = \lambda^V_{\max} \geq 3$: we show that graph G admits $K_{2,\lambda}$ as a minor. As $K_{2,3}$ is itself a subgraph of $K_{2,\lambda}$, then the lemma holds.

We describe the steps of deletion and contraction on G in order to obtain $K_{2,\lambda}$. Let G' be the subgraph of G which admits a minimal vertex (s,t)-cut X of size λ. We remove all the edges of G which are not in G'. For any vertex $x \in X$, since $X \backslash \{x\}$ is not an (s,t)-cut in G', there exists at least one (s,t)-path, denoted by P_x, which passes through x but not through any other element of X. We denote by x_s (resp. x_t) the vertex of P_x adjacent to x in $R_{G'}(X,s)$ (resp. $R_{G'}(X,t)$). We delete the edges which are not on any path P_x, $x \in X$. Moreover, for each $x \in X$, we contract all edges in sections $P_x^{(s,x_s)}$ and $P_x^{(x_t,t)}$. The edges remaining either connect an element of X with a contracted node in $R(X,s)$ "representing" s or an element of X with a contracted node in $R(X,t)$ "representing" t. The graph obtained is a complete bipartite graph $K_{2,\lambda}$. □

As a consequence, the competitive ratio of BYPASS on $K_{2,3}$-minor-free graphs is $2^{\frac{3}{4}}k + O(1)$. Among them, we find outerplanar graphs since $K_{2,3}$ is one of their forbidden minor [9]. In brief, Theorem 2 and Lemma 2 give us a well-known family of graphs, without weight restrictions, for which the bound $2k+1$ is not optimal.

Corollary 1. *There is a strategy, called* BYPASS, *which achieves a competitive ratio $2^{\frac{3}{4}}k + C$ on outerplanar graphs, where C is a constant.*

We conclude this overview with a last result, stating that any unweighted graph satisfying $k < \lambda^V_{\max}$ is a worst-case graph. Formally, for any graph G such that the number of blocages is less than $\lambda^V_{\max}$, there is an edge weighting of G, with some weights $\varepsilon \ll 1$, such that the best ratio we can obtain on G is $2k+1$ when ε tends to zero.

Theorem 5. *For any graph G satisfying $k < \lambda^V_{max}$ and any $\varepsilon > 0$, there is a weighting $\omega_\varepsilon : E \to \mathbb{Q}^+$ such that the competitive ratio of any deterministic strategy on the family $\mathcal{F}(G)$ is at least $\frac{2k+1}{1+\varepsilon}$.*

Proof. Let H be a subgraph of G and X a maximum vertex (s,t)-cut of H such that $|X| = \lambda^V_{\max}$. Set $R_H(X,s)$ is the source side of X in H, $R_H(X,t)$ is its target side.

We define weights ω_ε. For all edges with two endpoints in $R_H(X,s) \cup R_H(X,t)$, *i.e.* $e \in E[R_H(X,s)] \cup E[R_H(X,t)]$, we set $\omega_\varepsilon(e) = \frac{\varepsilon}{n}$. As cut X is minimal, any vertex $x \in X$ is adjacent to both $R_H(X,s)$ and $R_H(X,t)$, otherwise x would be useless for the separation of s and t. For any $x \in X$, we select an arbitrary edge from $R(X,s) \times \{x\}$ and fix its weight to value 1. Furthermore, we select an arbitrary edge e of $\{x\} \times R(X,t)$ and fix $\omega_\varepsilon(e) = \frac{\varepsilon}{n}$. Finally, the edges of graph H that have not been treated yet and the edges in $E(G) \backslash E(H)$ are set to infinite weight: more rigorously, assigning weight $(2k+1)n+1$ is equivalent to setting a weight $+\infty$ as traversing such edges becomes necessarily inefficient. Consequently, we obtain a graph (Fig. 7) almost equivalent to W_k, where $k+1 = \lambda^V_{\max}$.

Suppose that set E_* contains k $(\frac{\varepsilon}{n})$-weighted edges from $X \times R(X,t)$. The traveller must pass through X at some point to arrive at t. At worst, each attempt to go through vertex $x \in X$ ends with a blockage (x,y) with $y \in R(X,t)$. As $k < \lambda^V_{\max}$, this can occur k times. The distance traversed between each blockagee discovery is longer than 2. After discovering the k^{th} blocked edge and returning to set $R(X,s)$, the traveller goes to t with a distance of at least 1. Hence, the total worst-case distance traversed is at least $2k+1$ and no deterministic strategy has a competitive ratio lower than $\frac{2k+1}{\omega_{\mathrm{opt}}}$. With the weight function ω_ε, we have $\omega_{\mathrm{opt}} \le 1+\varepsilon$. Therefore, the obtained lower bound is $\frac{2k+1}{1+\varepsilon}$. □

This theorem shows that $\lambda^V_{\max}$ is of significant importance in the competitive analysis of the k-CTP. Indeed, all graphs satisfying $k < \lambda^V_{\max}$, associated with a certain edge weighting, belong to worst-case road maps. Therefore, there is no hope for extending these families of bounded-cut graphs and obtaining a ratio $\rho k + O(\lambda)$, $\rho < 2$, as in Theorems 1 and 2. Future open challenges for bounded-cut graphs are narrowed down to finding strategies with smaller slopes than DETOUR for $k \ge \mu^E_{\max}$, or BYPASS, for $k \ge \lambda^V_{\max}$.

Another consequence of Theorem 5 is that, contrary to outerplanar graphs, many well-known families of graphs do not have efficient deterministic strategies if we do not impose weight restrictions. For instance, for any integer k, there exist chordal, bipartite, and planar graphs satisfying $k < \lambda^V_{\max}$. For these "topological" families of graphs where the weights can take any value, a ratio $\rho k + O(1)$, $\rho < 2$, cannot be achieved by any deterministic strategy. Nevertheless, if we impose equal weights for example, *i.e.* $\omega(e) = 1$ for each $e \in E$, then it might be possible to find efficient strategies for these families.

4. A competitive strategy when $\mu^E_{\max} \le k$

We remind that the objective of this section is to prove Theorem 1. We propose a strategy, called DETOUR, which admits a competitive ratio $\sqrt{2}k + O(\mu^E_{\max})$. First, we introduce in Subsection 4.1 a parameterised strategy called α-DETOUR. It takes as input a graph G, source s, target t, and a parameter $0 \le \alpha \le 1$. In Subsection 4.2, we provide an upper bound on its competitive ratio. This bound is minimised for $\alpha = \frac{\sqrt{2}}{2}$ and is $2\mu^E_{\max} + \sqrt{2}(k - \mu^E_{\max}) + 1$ in this case. Strategy DETOUR mentioned earlier corresponds to $\frac{\sqrt{2}}{2}$-DETOUR.

4.1. Description of strategy α-DETOUR

We present the α-DETOUR strategy in Algorithm 1. Variable pos keeps track of the traveller's current position. The idea is to perform successively two phases: an *exploration* followed by a *detour-backtracking* (replacing the backtracking phase of REPOSITION). The exploration starts when the traveller is on source s (line 8). He traverses the shortest (s,t)-path $P^{(s,t)}_{\min}$ called the *exploration path*. Its cost $\omega^{(s,t)}_{\min} = \omega_{\min}(G, E'_*)$ is stored in $\omega_{\exp}$ (line 7). At this point, there are two possibilities: (i) the traveller reaches t and the execution terminates (line 13), (ii) the traveller arrives at $\mathrm{pos} = u$ and discovers a blocked edge $(u,v) \in P^{(s,t)}_{\min}$. All the vertices of the exploration path which have been visited before arriving on u are stored in a stack, denoted by the variable stack. Set V_{stack} refers to the set of vertices in the stack (without ordering). Then, the detour-backtracking phase begins.

When the traveller is blocked on $P^{(s,t)}_{\min}$, we ask whether an α-*detour* exists, *i.e.* an apparently open (pos,t)-path with cost at most $\alpha\omega_{\exp}$ in graph $G \backslash V_{\mathrm{stack}}$. If an α-detour exists, the traveller traverses the shortest path $P^{(\mathrm{pos},t)}_{\min}$ from the current position pos to target t in graph $G' \backslash V_{\mathrm{stack}}$, where $G' = G \backslash E'_*$ is graph G deprived of the blocked edges already discovered E'_* (line 9). Obviously, its cost satisfies $\omega^{(\mathrm{pos},t)}_{\min} \le \alpha\omega_{\exp}$. Otherwise, the traveller backtracks to the vertex before $\mathrm{pos} = u$ on the exploration path (lines 14-16) and we withdraw this vertex from stack.

We do not allow an α-detour $P^{(\mathrm{pos},t)}_{\min}$ to pass through any vertex $v \in V_{\mathrm{stack}}$, since the section $P^{(v,t)}_{\min}$ will be considered later on when $\mathrm{pos} = v$. The vertices of an exploration path traversed by the traveller are naturally put in stack. Moreover, when the traveller is blocked on an α-detour $P^{(\mathrm{pos},t)}_{\min}$, the vertices of $P^{(\mathrm{pos},t)}_{\min}$ from pos to the endpoint of the blocked edge visited are put in stack. Finally, if the traveller backtracks to s, the algorithm goes back to the exploration phase. At this moment, the stack is empty.

Let E'_* denote the set of discovered blocked edges. Variable G' contains the graph G deprived of the discovered blockages E'_* at any moment of the execution. At each iteration of the while loop, the variables are updated as follows: if the path $P^{(u_0,t)}_{\min}$ currently traversed (lines 8-9) - either an exploration path ($u_0 = s$) or a detour - does not contain any blockage, then the traveller reaches t, *i.e.* $\mathrm{pos} \leftarrow t$. In this case, the algorithm terminates since the destination is reached. Otherwise, let $P^{(u_0,t)}_{\min} = u_0 \cdots u_i \cdot u_{i+1} \cdots u_r \cdot t$, where (u_i, u_{i+1}) is its first blocked edge. The traveller's position is updated from u_0 to u_i (line 10). Then, we update E'_* with the newly discovered blockages including (u_i, u_{i+1}), and $G' \leftarrow G \backslash E'_*$ (line 12). In addition, we push the traversed vertices $u_0, \ldots, u_{i-1}$ on the stack (except u_i) and update accordingly $V_{\mathrm{stack}} \leftarrow V_{\mathrm{stack}} \cup \{u_0, \ldots, u_{i-1}\}$. In case there is no α-detour $P^{(u_i,t)}_{\min}$ in $G' \backslash V_{\mathrm{stack}}$, the algorithm backtracks by popping u_{i-1} from the stack and setting $\mathrm{pos} \leftarrow u_{i-1}$ (lines 14-16), etc.

If $\alpha = 0$, the algorithm does not take any detour. As a consequence, 0-DETOUR is equivalent to REPOSITION, as both procedures perform an exploration phase followed by backtracking without taking any detour. In the following, we provide an upper bound of α-DETOUR's competitive ratio.

Algorithm 1: The α-DETOUR strategy.

Input: graph G, source s, target t, parameter $\alpha \in (0, 1)$
$E'_* \leftarrow \emptyset$; $G' \leftarrow G \backslash E'_*$; pos $\leftarrow s$; $u_0 \leftarrow s$; $\omega_{\text{exp}} \leftarrow \omega_{\min}(G, \emptyset)$;
stack $\leftarrow$ Empty Stack; $V_{\text{stack}} \leftarrow \emptyset$;
while true **do**
  $u_0 \leftarrow$ pos;
  if $u_0 = s$ **then**
    $\omega_{\text{exp}} \leftarrow \omega_{\min}(G, E'_*)$;
    traverse the shortest (s,t)-path $P_{\min}^{(u_0,t)}$ in G'; # exploration
  else
    traverse the shortest (u_0,t)-path $P_{\min}^{(u_0,t)}$ in $G' \backslash V_{\text{stack}}$; # detour
  endif
  update pos;
  push the vertices visited in $P_{\min}^{(u_0,t)}$ except pos on stack and V_{stack};
  update E'_* and $G' \leftarrow G \backslash E'_*$;
  if $pos = t$ **then break**;
  while $pos \neq s$ **and** *there is no* $P_{min}^{(pos,t)}$ *in* $G' \backslash V_{stack}$ *such that* $\omega_{min}^{(pos,t)} \leq \alpha\omega_{exp}$ **do**
    pos $\leftarrow$ pop(stack);
    $V_{\text{stack}} \leftarrow V_{\text{stack}} \backslash \{\text{pos}\}$;
  end
end

4.2. Competitive analysis of detour

We denote by $P_1, \dots, P_\ell$ the exploration paths such that the distance from s to the blocked edge discovered on it is greater than α multiplied by their own cost, *i.e.* $\alpha\omega_i$. In other words, the distance d_i traversed by the traveller on the exploration path P_i, $1 \leq i \leq \ell$, satisfies $d_i \geq \alpha\omega_i$. Paths P_i are sorted in order to fulfil $\omega_1 \leq \cdots \leq \omega_\ell$. This corresponds to the order in which these paths are traversed. The exploration paths $P_1,\dots,P_{\ell-1}$ are necessarily blocked, while path P_ℓ may be open. If P_ℓ does not contain any blockage, then the algorithm terminates after the traveller traverses it. Otherwise, it means the traveller is blocked on P_ℓ and reaches target t via an α-detour.

From now on, we impose $\alpha \geq \frac{1}{2}$. Let us partition $P_1, \dots, P_\ell$ into two sequences $S_1 = P_1, \dots, P_{h-1}$ and $S_2 = P_h, \dots, P_\ell$ such that $2\alpha\omega_{h-1} < \omega_\ell \leq 2\alpha\omega_h$. In the particular case where $\omega_\ell \leq 2\alpha\omega_1$, then $h = 1$ and the two sets are $S_1 = \emptyset$ and $S_2 = P_1, \dots, P_\ell$. We denote by $G[P_h, \dots, P_\ell]$ the subgraph of G induced by paths $P_h, \dots, P_\ell$, *i.e.* containing only the vertices and edges of paths P_i, $h \leq i \leq \ell$.

Theorem 6. *The size of maximum edge (s,t)-cuts on graph $G[P_h, \dots, P_\ell]$ is at least $\ell - h + 1$.*

Proof. We denote by b_i the blocked edge discovered on P_i, for $i \in \{h, \dots, \ell\}$. We construct inductively a set $\{e_h, \dots, e_\ell\}$ of edges satisfying the following induction hypotheses, for all $i \in \{h, \dots, \ell\}$:

H_1(i): $\{e_h, \dots, e_i\}$ is a minimal (s,t)-cut of $G[P_h, \dots, P_i]$,
H_2(i): Either $e_i = b_i$ or e_i is an ancestor of b_i in P_i,
H_3(i): For $j \in \{i+1, \dots, \ell\}$, P_j cannot pass through e_i.

Basis: For $i = h$, $G[P_h, \dots, P_i]$ contains only one path P_h. We choose $e_h = b_h$, which fulfils $H_2(h)$. Since any edge of P_h is a max-(s,t)-cut of $G[P_h]$, it satisfies $H_1(h)$. Statement $H_3(h)$ is also true, as e_h is blocked.

Inductive step: Assume that $H_1(i)$ to $H_3(i)$ are true for a certain integer i in $\{h, \dots, \ell-1\}$. We will construct e_{i+1} and prove the induction hypotheses $H_1(i+1)$ to $H_3(i+1)$. For simplicity, we denote sets $R(\{e_h, \dots, e_i\}, s)$ and $R(\{e_h, \dots, e_i\}, t)$ in graph $G[P_h, \dots, P_i]$ by $R_i(s)$ and $R_i(t)$, respectively.

Let $P_{i+1}^{(v_0,v_p)} = v_0 \cdot v_1 \cdots v_p$ be the longest section in P_{i+1}, starting from $v_0 = s$, such that $v_0, \dots, v_p \in R_i(s)$. Section $P_{i+1}^{(v_0,v_p)}$ contains at least vertex $v_0 = s$. For $j \in \{h, \dots, i\}$, all ancestors of e_j in P_j belong to $R_i(s)$, and all descendants belong to $R_i(t)$. Therefore, according to $H_2(i)$, all exploration paths' sections of the form $P_j^{(s,u)}$ are open and equal to the shortest path from s to u, for $u \in R_i(s) \cap P_j$ and $j \in \{h, \dots, i\}$. In particular, since $P_{i+1}^{(v_0,v_p)}$ is the shortest (v_0, v_p)-path, we deduce that it is open as v_p belongs to some P_j by definition of $R_i(s)$.

According to $H_3(i)$, $P_{i+1}^{(v_p,t)}$ is a new path connecting $R_i(s)$ to $R_i(t)$, which does not traverse any edge of the cut $\{e_h, \dots, e_i\}$. Furthermore, we show that no vertex in $P_{i+1}^{(v_{p+1},t)}$ belongs to $R_i(s)$. Indeed, suppose for the sake of contradiction that $u \in P_{i+1}^{(v_{p+1},t)}$ and $u \in R_i(s)$. There would exist $j \in \{h, \dots, i\}$, such that $P_j^{(s,u)}$ is the shortest (s,u)-path, and all its vertices belong to $R_i(s)$. This contradicts with the fact that $P_{i+1}^{(s,u)}$ is also the shortest (s,u)-path and $v_{p+1} \notin R_i(s)$, by definition. Let $v_{p'}$ be the first vertex of P_{i+1} belonging to $R_i(t)$, *i.e.* $v_{p'} \in R_i(t)$ and $p < p'$. Such a vertex exists as t is

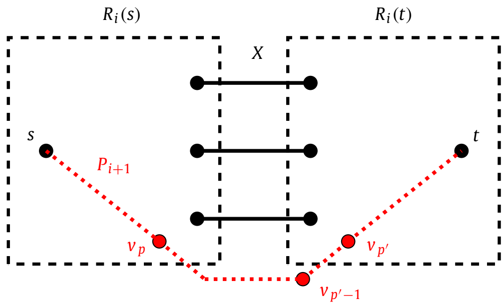

Fig. 8. Cut $X = \{e_h, \ldots, e_i\}$, path P_{i+1}, and vertices v_p, $v_{p'-1}$, $v_{p'}$.

a candidate. We derive that $P_{i+1}^{(v_p,v_{p'})}$ is a path both connecting $R_i(s)$ to $R_i(t)$ and avoiding cut X. Fig. 8 represents cut $\{e_h, \ldots, e_i\}$, path P_{i+1} and its vertices v_p and $v_{p'}$.

We fix e_{i+1} differently depending on the position of b_{i+1}. We already proved that $b_{i+1} \notin P_{i+1}^{(v_0,v_p)}$, the remaining cases are:

- If $b_{i+1} \in P_{i+1}^{(v_p,v_{p'})}$, then we set $e_{i+1} = b_{i+1}$. As $e_{i+1} \in E_*$, $H_3(i+1)$ is true.
- Otherwise, if $b_{i+1} \in P_{i+1}^{(v_{p'},t)}$, we select $e_{i+1} = (v_{p'-1}, v_{p'})$. We prove that the cost of the current shortest $(s, v_{p'})$-path, $P_{i+1}^{(s,v_{p'})}$, is at least $\alpha\omega_h$. Indeed, as vertex $v_{p'}$ belongs to a certain path $P_{j'}$, $j' \in \{h, \ldots, i\}$, the cost of $P_{i+1}^{(s,v_{p'})}$ is at least the cost of $P_{j'}^{(s,v_{p'})}$. If we have $\omega_{i+1}^{(s,v_{p'})} \le \alpha\omega_h$, the distance traversed by the traveller on $P_{j'}$ is less than $\alpha\omega_h \le \alpha\omega_{j'}$, as $v_{p'} \in R_i(t)$. This contradicts with the fact that $P_{j'} \in \{P_h, \ldots, P_\ell\}$. Moreover, after the $(i+1)$-th detour-backtracking phase, all remaining open $(v_{p'}, t)$-paths are longer than $\alpha\omega_{i+1} \ge \alpha\omega_h$, otherwise they would be α-detours and would have been traversed. Therefore, the cost of any exploration (s,t)-path passing through $v_{p'}$ is greater than $\alpha\omega_h + \alpha\omega_{i+1} \ge 2\alpha\omega_h$. This is impossible since the last exploration path P_ℓ satisfies $\omega_\ell \le 2\alpha\omega_h$. As a consequence, no exploration path passes through $v_{p'}$ and $H_3(i+1)$ is true.

Both cases fulfil naturally $H_2(i+1)$. It only remains to prove statement $H_1(i+1)$. We showed that $P_{i+1}^{(v_p,v_{p'})}$ is the only path connecting $R_i(s)$ to $R_i(t)$, and $e_{i+1} \in P_{i+1}^{(v_p,v_{p'})}$. Thus, $\{e_h, \ldots, e_{i+1}\}$ is an (s,t)-cut of $G[P_h, \ldots, P_{i+1}]$.

If we re-open edge e_{i+1}, path $P_{i+1}^{(v_p,v_{p'})}$ connects $R_i(s)$ to $R_i(t)$. If we re-open e_j, $j < i+1$, there is a path in $G[P_h, \ldots, P_i]$ which connects $R_i(s)$ to $R_i(t)$ independently of $P_{i+1}^{(v_p,v_{p'})}$, according to the minimality of $\{e_h, \ldots, e_i\}$ in $H_1(i)$. As a consequence, no proper subset of $\{e_h, \ldots, e_{i+1}\}$ is an (s,t)-cut. Cut $\{e_h, \ldots, e_{i+1}\}$ is minimal.

In summary, we derive by induction that $\{e_h, \ldots, e_\ell\}$ is a minimal (s,t)-cut of $G[P_h, \ldots, P_\ell]$. Thus, the size of the maximum edge (s,t)-cut is at least $\ell - h + 1$. □

From $\lambda_{\max}^E = \mu_{\max}^E$, we know that the maximum edge (s,t)-cut size $\mu_{\max}^E$ of G is greater or equal to the size of any minimal edge (s,t)-cut of a subgraph of G. As a consequence of Theorem 6, a relationship exists between values ℓ, h, and $\mu_{\max}^E$, which is $\ell - h + 1 \le \mu_{\max}^E$.

After traversing an exploration path P_i, the traveller performs a detour-backtracking phase. The number of blockages discovered during this i-th detour-backtracking phase is denoted by q_i. We analyse the cost of traversing P_i and performing the i-th detour-backtracking phase in Lemma 3.

Lemma 3. *The total cost of both the i-th exploration phase and the i-th detour-backtracking phase is not greater than* $(2 + 2\alpha q_i)\omega_i$.

Proof. The stack in Algorithm 1 ensures that each edge is only traversed twice: first time when moving towards t on an exploration path or a detour, and a second time when backtracking. The exploration path costs ω_i and each detour costs no more than $\alpha\omega_i$. Besides, the number of detours is at most q_i. Hence, the total cost is at most $2\omega_i + q_i 2\alpha\omega_i$, which concludes the proof. □

We denote by k_1 (resp. k_2) the number of blocked edges discovered during the exploration and detour-backtracking phases associated with paths $P_1, \ldots, P_{h-1}$ (resp. $P_h, \ldots, P_\ell$). Let k_3 be the number of blockages discovered during the other phases, so that $k_1 + k_2 + k_3 = k$. We derive in Theorem 7 an upper-bound on the competitive ratio as a function of k_1, k_2, k_3, and α.

Theorem 7. *The competitive ratio of α-DETOUR is upper-bounded by:*

$$\frac{k_1}{\alpha} + 2\mu^E_{max} + 2\alpha(k_2 + k_3 - \mu^E_{max}) + 1. \tag{3}$$

Proof. Since path P_ℓ is the shortest (s,t)-path of a certain graph $G\backslash E'_*$ where $E'_* \subseteq E_*$, the offline optimal cost satisfies

$$\omega_{\text{opt}} \geq \omega_\ell. \tag{4}$$

According to Lemma 3, the distance traversed during the exploration and detour-backtracking phases of $P_1, \ldots, P_{h-1}$ is not greater than

$$\sum_{j=1}^{h-1} (2 + 2\alpha q_j)\omega_j \leq 2\omega_{h-1} \sum_{j=1}^{h-1} (1 + q_j) = 2k_1\omega_{h-1}. \tag{5}$$

Inequality (5) comes from the fact that $\omega_1 \leq \cdots \leq \omega_{h-1}$, $\alpha \leq 1$ and $\sum_{j=1}^{h-1}(1+q_j) = k_1$.

We evaluate the cost of the phases associated with $P_h, \ldots, P_\ell$. Path P_ℓ is either open and traversed in one direction only (Case 1) or it is blocked and the traveller reaches t via a detour (Case 2).

Case 1: If P_ℓ does not contain any blockage, then the algorithm terminates after traversing it. This final exploration phase costs ω_ℓ. We have $q_\ell = 0$ and $k_2 = \sum_{j=h}^{\ell-1}(1+q_j)$. Given Lemma 3, the cost of the h-th to ℓ-th phases is less than:

$$\begin{aligned}
\sum_{j=h}^{\ell-1} (2 + 2\alpha q_j)\omega_j + \omega_\ell &= \sum_{j=h}^{\ell-1} (2\alpha + 2\alpha q_j)\omega_j + \sum_{j=h}^{\ell-1} (2 - 2\alpha)\omega_j + \omega_\ell, \\
&\leq 2\alpha k_2\omega_\ell + (2-2\alpha)(\ell - h)\omega_\ell + \omega_\ell, &(6)\\
&< 2\alpha k_2\omega_\ell + (2-2\alpha)\mu^E_{\max}\omega_\ell + \omega_\ell, &(7)\\
&= 2\alpha(k_2 - \mu^E_{\max})\omega_\ell + 2\mu^E_{\max}\omega_\ell + \omega_\ell.
\end{aligned}$$

We deduce Inequality (6) from $\omega_h \leq \cdots \leq \omega_\ell$. By applying Theorem 6 on $S_2 = P_h, \ldots, P_\ell$, we derive that $\ell - h \leq \mu^E_{\max} - 1 < \mu^E_{\max}$ in Inequality (7).

Case 2: Suppose that P_ℓ is blocked. The ℓ-th exploration and detour-backtracking phases cost at most $(2+2\alpha q_\ell)\omega_\ell + \alpha\omega_\ell$. Moreover, we have $k_2 = \sum_{j=h}^{\ell}(1+q_j)$. The distance traversed from the h-th to the ℓ-th phases is not greater than:

$$\begin{aligned}
\sum_{j=h}^{\ell-1} (2 + 2\alpha q_j)\omega_j + (2 + 2\alpha q_\ell + \alpha)\omega_\ell &= \sum_{j=h}^{\ell} (2 + 2\alpha q_j)\omega_j + \alpha\omega_\ell, \\
&\leq 2\alpha k_2\omega_\ell + (2-2\alpha)(\ell - h + 1)\omega_\ell + \alpha\omega_\ell, &(8)\\
&\leq 2\alpha k_2\omega_\ell + (2-2\alpha)\mu^E_{\max}\omega_\ell + \alpha\omega_\ell, &(9)\\
&\leq 2\alpha(k_2 - \mu^E_{\max})\omega_\ell + 2\mu^E_{\max}\omega_\ell + \omega_\ell. &(10)
\end{aligned}$$

Inequality (8) follows from $\omega_h \leq \cdots \leq \omega_\ell$. We obtain (9) from $\ell - h + 1 \leq \mu^E_{\max}$. Finally, $\alpha \leq 1$ implies Eq. (10).

By definition, any path $\widehat{P}$ not in $P_1, \ldots, P_\ell$ is such that the distance traversed on it is at most α multiplied by its own cost $\widehat{\omega}$. The distance traversed during the phases which are not associated with $P_1, \ldots, P_\ell$ is the cost of these exploration paths $\widehat{P}$ and their α-detours. As $\widehat{\omega} \leq \omega_{\text{opt}}$, this distance is at most $2\alpha k_3\omega_{\text{opt}}$. Applying Eq. (4) on the above inequalities, the competitive ratio of α-DETOUR admits the following upper-bound:

$$\begin{aligned}
\frac{\omega_{\alpha-\text{DETOUR}}}{\omega_{\text{opt}}} &\leq \frac{2k_1\omega_{h-1} + 2\alpha(k_2 + k_3 - \mu^E_{\max})\omega_{\text{opt}} + 2\mu^E_{\max}\omega_{\text{opt}} + \omega_{\text{opt}}}{\omega_{\text{opt}}}, \\
&\leq \frac{k_1\omega_\ell}{\alpha\omega_{\text{opt}}} + 2\mu^E_{\max} + 2\alpha(k_2 + k_3 - \mu^E_{\max}) + 1, &(11)\\
&\leq \frac{k_1}{\alpha} + 2\mu^E_{\max} + 2\alpha(k_2 + k_3 - \mu^E_{\max}) + 1. &(12)
\end{aligned}$$

Inequality (11) follows from the partition $\{S_1, S_2\}$ which imposes $2\alpha\omega_{h-1} < \omega_\ell$. □

Let $c_{\text{det}}(k_1, k_2, k_3, \alpha)$ denote the value in (12). Parameters k_1, k_2, and k_3 depend on the road map (G, E_*), so only $\alpha \in (0,1)$ can be tuned. Value $\alpha = \frac{\sqrt{2}}{2}$ minimises $c_{\text{det}}(k_1, k_2, k_3, \alpha)$ under the condition $k_1 + k_2 + k_3 = k$ for any $k > \mu^E_{\max}$. Formally,

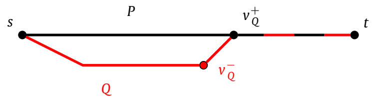

Fig. 9. A bypass Q of the current shortest (s,t)-path P.

$$\frac{\sqrt{2}}{2} = \underset{0\le\alpha\le 1}{\operatorname{argmin}} \max_{\substack{k_1,k_2,k_3\in\mathbb{N}\\ k_1+k_2+k_3=k}} c_{\det}(k_1,k_2,k_3,\alpha).$$

Corollary 2. *The competitive ratio of* DETOUR *is at most* $2\mu^E_{max} + \sqrt{2}(k - \mu^E_{max}) + 1$.

Proof. We obtain this ratio by setting $\alpha = \frac{\sqrt{2}}{2}$ and $k_1 + k_2 + k_3 = k$ in Eq. (3). □

In summary, strategy DETOUR offers the same performance guarantee as REPOSITION for the range $\mu^E_{\max} \ge k$ but is more competitive for the range $\mu^E_{\max} < k$. The slope of the competitive ratio of DETOUR when k varies is only $\sqrt{2}$ for $\mu^E_{\max} < k$.

DETOUR strategy needs to identify the shortest (s,t)-path and (pos,t)-path at any moment of its execution. To achieve it, Dijkstra's algorithm [12] is computed once between two discoveries of blocked edges with t as the start point. Hence, the running time of DETOUR is $O(k(m + n\log n))$.

Similarly to REPOSITION and COMPARISON, the execution of DETOUR strategy is independent of the value of k. Thus, it can be used when no upper bound on the number of blockages is known and its competitive ratio is $2\mu^E_{\max} + \sqrt{2}(|E_*| - \mu^E_{\max}) + 1$.

DETOUR strategy can be executed without knowing the value $\mu^E_{\max}$. Indeed, the competitive ratio of DETOUR depends on $\mu^E_{\max}$ but no decision is made based on $\mu^E_{\max}$ in Algorithm 1. Consequently, determining $\mu^E_{\max}$ before the execution of DETOUR only offers a guarantee on the distance to be traversed: it is not necessary to launch this strategy.

5. A competitive strategy when $\lambda^V_{\max} \le k$

Strategy BYPASS, presented in this section, is an extension of the previous strategy DETOUR. Its competitive ratio is $2^{\frac{3}{4}}k + O(\lambda^V_{\max})$, where $\lambda^V_{\max}$ is the size of the largest minimal vertex (s,t)-cut we can find in a subgraph of G.

5.1. Description of strategy bypass

The main difference between BYPASS and DETOUR is that BYPASS guides the traveller through a new kind of paths called *bypasses*. Under certain conditions, these paths are traversed instead of the current shortest (s,t)-path. The idea behind this additional rule is to avoid multiple passings through the vertices of a minimal vertex (s,t)-cut. We introduce a parameter $\beta \ge 1$ which limits the cost of bypasses. If their cost is too large, the competitive ratio of BYPASS may shoot up. Below, we give the definition of bypasses and Fig. 9 represents the introduced notations.

Definition 3 *(Bypass Q of the shortest (s,t)-path P_{min}).* We consider graph $G\backslash E'_*$, where E'_* is the set of blocked edges already discovered. Let P_{min} denote the shortest (s,t)-path in $G\backslash E'_*$. A bypass Q is the concatenation of three sections: a shortest (s, v_Q^-)-path, an edge (v_Q^-, v_Q^+), and a shortest (v_Q^+, t)-path such that:

- vertex v_Q^+ belongs to P_{min}: $v_Q^+ \in P_{\min}$,
- the cost of Q is at most $\beta\omega_{min}$: $\omega_Q \le \beta\omega_{min}$,
- the cost of the section $Q^{(s,v_Q^-)}$ is at most $\frac{1}{\beta}\omega_{min}$.

The edge (v_Q^-, v_Q^+) of bypass Q is called the *transition edge* and v_Q^+ the transition vertex. The shortest (s,t)-path P_{min} is called the *fictive path* of Q. As $v_Q^+ \in P_{min}$, sections $P_{min}^{(v_Q^+,t)}$ and $Q^{(v_Q^+,t)}$ are identical.

Section $Q^{(s,v_Q^+)}$ is longer than $P^{(s,v_Q^+)}$. We observe with Definition 3 that $P_{\min}$ is its own bypass. Considering any edge (v^-, v^+) of $P_{\min}$, where $\omega(P_{\min}^{(s,v^-)}) \le \frac{1}{\beta}\omega_{\min}$, guarantees the conditions listed. The idea of strategy BYPASS is to traverse a bypass Q of $P_{\min}$ such that the transition vertex v_Q^+ is as close as possible from target t. As a consequence, the path traversed is either $P_{\min}$ itself or a bypass $Q \ne P_{\min}$ such that the transition vertex $v_Q^+ \in P_{\min}$ fulfils $\omega(P_{\min}^{(s,v^+)}) > \frac{1}{\beta}\omega_{\min}$.

A short description of BYPASS follows. As DETOUR, it can be divided into two phases: exploration and detour-backtracking. The exploration phase is modified as it now allows to traverse bypasses. Strategy BYPASS makes the traveller traverse, as a priority, the bypass Q such that v_Q^+ is as close as possible from t. The exploration path is either the current shortest

Algorithm 2: The BYPASS strategy.

1: **Input:** graph G, source s, target t
2: $E'_* \leftarrow \emptyset$; $G' \leftarrow G \backslash E'_*$; pos $\leftarrow s$; $u_0 \leftarrow s$; $\omega_{\text{exp}} \leftarrow \omega_{\min}(G, \emptyset)$;
3: stack $\leftarrow$ Empty Stack; $V_{\text{stack}} \leftarrow \emptyset$;
4: **while** true **do**
5: $u_0 \leftarrow$ pos;
6: **if** $u_0 = s$ **then**
7: $\omega_{\text{exp}} \leftarrow \omega_{\min}(G, E'_*)$; $P_{\min}^{(u_0,t)} \leftarrow$ shortest (u_0,t)-path in G';
8: $v^+ \leftarrow t$; bypassfound $\leftarrow$ false;
9: **while** $v^+ \neq s$ **and** *bypassfound* = false **do**
10: **if** *there is a bypass Q passing through v^+* **then**
11: bypassfound $\leftarrow$ true;
else
12: $v^+ \leftarrow$ predecessor of v^+ on $P_{\min}^{(u_0,t)}$;
endif
end
13: **if** *bypassfound* **then** traverse Q **else** $P_{\min}^{(u_0,t)}$;
else
14: traverse the shortest (u_0,t)-path $P_{\min}^{(u_0,t)}$ in $G' \backslash V_{\text{stack}}$; # detour
endif
15: update pos; # either t or the endpoint of a blocked edge
16: push the vertices visited in $P_{\min}^{(u_0,t)}$ except pos on stack and V_{stack};
17: update E'_* and $G' \leftarrow G \backslash E'_*$;
18: **if** *pos* $= t$ **then break**;
19: **while** *pos* $\neq s$ **and** *there is no* $P_{\min}^{(pos,t)}$ *in* $G' \backslash V_{stack}$ *such that* $\omega_{\min}^{(pos,t)} \leq \frac{1}{\beta}\omega_{\text{exp}}$ **do**
20: pos $\leftarrow$ pop(stack); $V_{\text{stack}} \leftarrow V_{\text{stack}} \backslash \{\text{pos}\}$;
end
end

(s,t)-path $P_{\min}$ or one of its bypasses different from $P_{\min}$. The detour-backtracking phase follows the same principle as with DETOUR: each time the traveller meets a detour, he traverses it. However, the definition of a detour changes. From now on, we say a detour is a path from the traveller's position to target t which is apparently open and with a cost at most $\frac{1}{\beta}\omega_{\min}$.

Algorithm 2 gives the pseudocode of BYPASS. Lines 1-5 and 14-20 stay unchanged compared to Algorithm 1 which presents DETOUR, except for line 19 where the maximum cost of detours is modified. Lines 6-13 describe the exploration phase of BYPASS. They replace lines 6-8 of Algorithm 1. Indeed, instead of traversing directly the current shortest (s,t)-path, BYPASS checks whether some bypass joins $P_{\min}$ closer to target t.

Once the current shortest (s,t)-path $P_{\min}^{(s,t)} = P_{\min}$ is computed, the existence of a bypass is verified. For each vertex v^+ of $P_{\min}$, starting from t, we check whether there is a bypass Q such that v^+ is the target-side endpoint v_Q^+ of the transition edge. In other words, we check whether there is a bypass Q with some transition edge (v, v^+), $v \in V$. This verification is given in line 10. We deliberately omitted the details of this procedure in the pseudocode to keep it concise and understandable, so we explain it now. For each v^+, we look at all its neighbours v and, for each of them, we check whether the concatenation of the shortest (s,v)-path with edge (v,v^+) and $P_{\min}^{(v^+,t)}$ produces a path fulfilling the conditions given by Definition 3. Once a bypass Q is identified, we stop the listing of neighbours and update the boolean bypassfound $\leftarrow$ true.

5.2. Competitive analysis of bypass

The starting point of this analysis is similar to the one of DETOUR. The exploration paths $P_1, \ldots, P_\ell$ are either shortest (s,t)-paths or bypasses. We divide the exploration paths $P_1, \ldots, P_\ell$ in two disjoint sequences: $S_1 = P_1, \ldots, P_{h-1}$ and $S_2 = P_h, \ldots, P_\ell$. The distance traversed on each path of S_1 is at most the threshold value $\frac{1}{\beta}\omega_{\text{opt}}$, while S_2 contains all paths on which the distance traversed is greater than $\frac{1}{\beta}\omega_{\text{opt}}$. As in Section 4.2, paths in S_1 and all detours will be seen as "short paths".

To preserve the advantages of strategy DETOUR, the cost of the current shortest (s,t)-path must be at least $\frac{\sqrt{2}}{2}\omega_{\text{opt}}$ when exploring paths in S_2. Considering an exploration path of S_2, we distinguish two cases. If the exploration path is the shortest (s,t)-path in $G \backslash E'_*$, then we want its cost to be at least $\frac{\sqrt{2}}{2}\omega_{\text{opt}}$. To satisfy this condition, we have necessarily $\frac{1}{\beta}\omega_{\text{opt}} \geq \frac{\sqrt{2}}{2}\omega_{\text{opt}}$, thus $\beta \leq \sqrt{2}$. Otherwise, if the exploration path is a bypass, its cost is at least $\frac{1}{\beta}\omega_{\text{opt}}$, therefore the cost of its fictive path can potentially reach $(\frac{1}{\beta})^2\omega_{\text{opt}}$ at least. In order to maintain the cost of all fictive paths above $\frac{\sqrt{2}}{2}\omega_{\text{opt}}$, we must have $\beta \leq \sqrt{\sqrt{2}} = 2^{\frac{1}{4}}$. In brief, $1 \leq \beta \leq 2^{\frac{1}{4}}$. We will see in the remainder that value $\beta = 2^{\frac{1}{4}}$ minimises the competitive ratio of BYPASS.

Paths $P_h, \dots, P_\ell$ are sorted in the order in which they are considered. However, their weights $\omega_h, \dots, \omega_\ell$ are not necessarily increasing. It may happen that a bypass P_i has a larger cost than the next exploration path P_{i+1}. In summary, any path P_i, $h \le i \le \ell$ belongs to one of these categories:

- **Type A**: path P_i is the current shortest (s,t)-path when the traveller begins his walk on it. We have $\frac{1}{\beta}\omega_{\text{opt}} \le \omega_i \le \omega_{\text{opt}}$.
- **Type B**: path P_i is a bypass. It admits a fictive path F_i, the current shortest (s,t)-path, and $F_i \neq P_i$. The cost of P_i may be larger than ω_{opt}. It satisfies $\frac{1}{\beta}\omega_{\text{opt}} \le \omega_i \le \beta\omega_{\text{opt}}$. The cost of F_i is such that $\frac{\sqrt{2}}{2}\omega_{\text{opt}} \le \omega(F_i) \le \omega_{\text{opt}}$.

When the exploration path $P_i \in S_2$ is a bypass, then the blocked edge discovered on it is placed after the transition vertex.

Lemma 4. *Assume $P_i \in S_2$ is a B-type exploration path,* i.e. *a bypass with a fictive path $F_i \neq P_i$. Let (v^-, v^+) be the transition edge of P_i. If P_i is not open, the blocked edge discovered on it belongs to section $P_i^{(v^+,t)} = F_i^{(v^+,t)}$.*

Proof. For the sake of contradiction, suppose that the traveller is blocked on $P_i^{(s,v^+)}$. The distance traversed by the traveller on P_i is thus at most $\omega(P_i^{(s,v^-)}) \le \frac{1}{\beta}\omega(F_i)$, according to Definition 3. As $\omega(F_i) \le \omega_{\text{opt}}$, the distance traversed on P_i is at most $\frac{1}{\beta}\omega_{\text{opt}}$. This is a contradiction with the definition of collection S_2 which contains the explorations for which the distance traversed on them is greater than $\frac{1}{\beta}\omega_{\text{opt}}$. $\square$

As with DETOUR, the competitive analysis of BYPASS consists in considering a sequence $G_h, \dots, G_\ell$ of subgraphs of G. Each subgraph G_i of G, $h \le i \le \ell$, contains the vertices and edges of paths $P_h, \dots, P_i$. For any $h \le i \le \ell$, we can identify a minimal vertex (s,t)-cut X_i on G_i of size $|X_i| = i - h + 1$. We build this cut X_i incrementally. Each time a new exploration path P_{i+1} is considered, we add a vertex x_{i+1} to the cut X_i such that $X_{i+1} = \{x_{i+1}\} \cup X_i$ is a minimal vertex (s,t)-cut in G_{i+1}. This way, the source and target sides of X_i, $R(X_i, s)$ and $R(X_i, t)$, are growing when i is increasing. Formally, $R(X_i, t) \subseteq R(X_{i+1}, t)$ and $R(X_i, s) \subseteq R(X_{i+1}, s)$. We use an inductive process to extend the cutset from any index i to $i+1$. Below are listed the induction hypotheses that will be satisfied by cut $X_i = \{x_h, \dots, x_i\}$, for $h \le i \le \ell$:

$H'_1(i)$: For any $h \le j \le i$, the (s,t)-path P_j traversed by the traveller contains exactly one vertex of X_i and it is x_j.
$H'_2(i)$: For any $h \le j \le i$, vertex x_j is an ancestor of the blocked edge discovered on path P_j. That is, the blocked edge discovered on P_j is on $P_j^{(x_j,t)}$, not on $P_j^{(s,x_j)}$.
$H'_3(i)$: All vertices v of graph G_i belonging to $X_i \cup R(X_i, t)$ are such that the cost of the current shortest (s,v)-path (after being blocked on P_i) is larger than $\frac{\sqrt{2}}{2\beta}\omega_{\text{opt}}$.

Induction hypotheses $H'_1(i)$ and $H'_2(i)$ are similar to $H_1(i)$ and $H_2(i)$, proposed for DETOUR in Section 4. Hypothesis $H'_3(i)$ is a stronger statement than $H_3(i)$ and allows us to handle the minimality of vertex (s,t)-cuts.

The base case follows, for which $i = h$. We prove there is a cut $X = \{x_h\}$ such that the statements $H'_1(h)$, $H'_2(h)$ and $H'_3(h)$ are true on the subgraph G_h containing only path P_h.

Theorem 8. *There is a vertex $x_h \in P_h$ which forms a minimal (s,t)-cut $X_h = \{x_h\}$ in G_h and which satisfies the hypotheses $H'_1(h)$, $H'_2(h)$ and $H'_3(h)$.*

Proof. Let (u_h, v_h) denote the blocked edge on P_h: u_h is the predecessor of v_h. We take $x_h = u_h$. This way, $H'_1(h)$ and $H'_2(h)$ are clearly satisfied. The set $X_h \cup R(X_h, t)$ contains u_h and its successors. The proof that $H'_3(h)$ is satisfied depends on the nature of P_h.

Type A. Path P_h is the current shortest (s,t)-path. So, $\omega_h = \omega_{\min} \le \omega_{\text{opt}}$. By definition, the cost of the traversed section $P_h^{(s,u_h)}$ is larger than $\frac{1}{\beta}\omega_{\text{opt}}$. The current distance between s and any vertex v of $X_h \cup R(X_h, t)$, *i.e.* any vertex v of $P_h^{(u_h,t)}$, is more than $\frac{1}{\beta}\omega_{\text{opt}} > \frac{\sqrt{2}}{2\beta}\omega_{\text{opt}}$. Throughout the execution of BYPASS, the distance between s and such v will either stay unchanged or increase due to the discovery of certain blockages.

Type B. Path P_h is a bypass of the shortest (s,t)-path $F_h = P_{\min}$ and $\omega_{\min} \ge \frac{\sqrt{2}}{2}\omega_{\text{opt}}$. The blocked edge (u_h, v_h) discovered on P_h also belongs to $P_{\min}$ as it arrives after the transition vertex $v^+_{P_h}$ of P_h, according to Lemma 4. Let v^- and v^+ be two successive vertices of F_h. If $\omega(F_h^{(s,v^-)}) \le \frac{1}{\beta}\omega(F_h)$, then (v^-, v^+) is a transition edge of F_h as a bypass of itself. We know by construction that $v^+_{P_h}$ is the transition vertex closest to the target t, therefore it is a descendant of any v^- satisfying the above inequality. We deduce that the distance from s to $v^+_{P_h}$ on F_h is greater than $\frac{1}{\beta}\omega(F_h)$. By combining this with

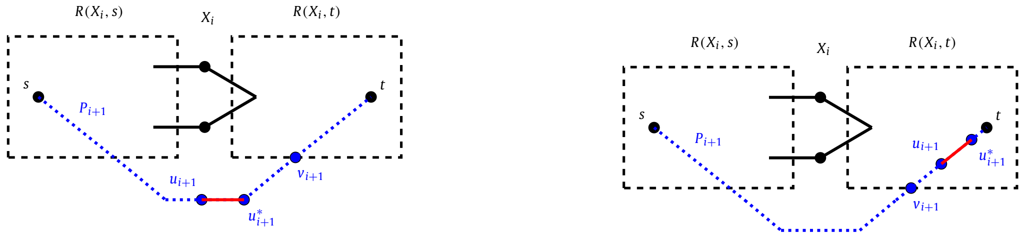

(a) v_{i+1} is a either u^*_{i+1} or one of its descendants (b) v_{i+1} is either u_{i+1} or one of its predecessors

Fig. 10. Position of v_{i+1} with respect to the blocked edge b_{i+1} (in red).

$\omega(F_h) \geq \frac{\sqrt{2}}{2}\omega_{\text{opt}}$ and the fact that u_h is a descendant of $v^+_{P_h}$, we get that the distance between s and u_h on F_h is at least $\frac{\sqrt{2}}{2\beta}\omega_{\text{opt}}$. Hence, condition $H'_3(h)$ holds.

It follows that in both cases the distance between s and $v \in X_h \cup R(X_h, t)$ will be greater than $\frac{\sqrt{2}}{2\beta}\omega_{\text{opt}}$ throughout the execution of BYPASS. $\square$

Let us now put in evidence an intermediate property that will be useful to prove the induction step in Theorem 9. It states that the vertices in $X_i \cup R(X_i, t)$, once they are visited by the traveller, cannot be visited anymore.

Lemma 5. *Suppose the hypothesis $H'_3(i)$ holds true for collection $P_h, \ldots, P_i$. Let $v \in X_i \cup R(X_i, t)$ be a vertex visited for the first time by the traveller when he traversed a path P_j, $h \leq j \leq i$. Then, vertex v will not be visited by the traveller anymore.*

Proof. We denote by $P_{\min}[j]$ the shortest (s, t)-path after P_j is traversed. We know that $\frac{\sqrt{2}}{2}\omega_{\text{opt}} \leq \omega(P_{\min}[j]) \leq \omega_{\text{opt}}$. We suppose by way of contradiction that vertex v is traversed again by the traveller on path P_q, $j \leq q \leq \ell$.

According to $H'_3(i)$, the cost of $P_q^{(s,v)}$ is greater than $\frac{\sqrt{2}}{2\beta}\omega_{\text{opt}}$. Section $P_q^{(v,t)}$ is not identical to section $P_j^{(v,t)}$ as the latter contains the blocked edge discovered on P_j. Therefore, the cost of $P_q^{(v,t)}$ is greater than $\frac{1}{\beta}\omega(P_{\min}[j]) \geq \frac{\sqrt{2}}{2\beta}\omega_{\text{opt}}$, otherwise it would have been traversed as a detour during the backtracking phase on P_j. Hence, we have:

$$\omega(P_q) = \omega(P_q^{(s,v)}) + \omega(P_q^{(v,t)}) > \frac{\sqrt{2}}{2\beta}\omega_{\text{opt}} + \frac{\sqrt{2}}{2\beta}\omega_{\text{opt}} \geq \beta\omega_{\text{opt}}.$$

The contradiction appears as the cost of path P_q is greater than $\beta\omega_{\text{opt}}$. $\square$

Lemma 5 guarantees that each vertex of the cut is visited exactly once. Indeed, vertex x_i is visited when the traveller traverses P_i as it is an ancestor of the blocked edge on P_i, according to $H'_1(i)$ and $H'_2(i)$. As a consequence, it will not be traversed anymore on paths $P_{i+1}, \ldots, P_\ell$. This ensures the minimality of cuts $X_h, \ldots, X_\ell$ because, for any cut X_q, $q \geq i$, the (s, t)-path P_i does not pass through any vertex of $X_q \setminus \{x_i\}$.

Lemmas 4 and 5 are used to prove the inductive step in the next theorem which is the keystone of our competitive analysis for BYPASS.

Theorem 9. *Assume hypotheses $H'_1(i)$, $H'_2(i)$, and $H'_3(i)$ are satisfied for X_i, $h \leq i \leq \ell$. There exists a vertex x_{i+1} such that $X_{i+1} = X_i \cup \{x_{i+1}\}$ is an (s, t)-cut satisfying hypotheses $H'_1(i+1)$, $H'_2(i+1)$ and $H'_3(i+1)$.*

Proof. We denote by $b_{i+1} = (u_{i+1}, u^*_{i+1})$ the blocked edge to be discovered on P_{i+1} by the traveller. The traveller stands on vertex u_{i+1} when this blockage is revealed. Let v_{i+1} be the closest-to-s vertex of P_{i+1} which is in $R(X_i, t)$. We distinguish two cases, represented in Fig. 10, depending on the locations of b_{i+1} and v_{i+1}. In Case 1, the blocked edge of P_{i+1} arrives "before" vertex v_{i+1}. We show that adding vertex v_{i+1} into the cut X_i preserves the induction hypotheses. In Case 2, the blocked edge P_{i+1} arrives "after" vertex v_{i+1}. Depending on the nature of the predecessor v^-_{i+1} of v_{i+1} on P_{i+1}, either the induction hypotheses stay true when v^-_{i+1} is added to the cut or a contradiction appears.

Case 1: First, assume that v_{i+1} is either u^*_{i+1} or one of its descendants (Fig. 10a). In this case, the traveller does not reach set $R(X_i, t)$ when he traverses P_{i+1}. We fix $x_{i+1} = u_{i+1}$. With this choice, X_{i+1} is an (s, t)-cut as x_{i+1} belongs to P_{i+1}, condition $H'_2(i+1)$ is therefore true. In the following, we will show that $X_{i+1} = X_i \cup \{x_{i+1}\}$ satisfies $H'_1(i+1)$ and $H'_3(i+1)$.

Set $R(X_i, s)$ only contains vertices which have already been visited by the traveller since they are "before" X_i, and consequently before the blockages discovered on their respective exploration paths. As a consequence, vertex u_{i+1} does not belong to $R(X_i, s)$, otherwise the blocked edge b_{i+1} would have been revealed earlier and P_{i+1} would not be the current exploration path, which is apparently open. In brief, vertex $x_{i+1} = u_{i+1}$ is neither in $R(X_i, s)$ nor in $R(X_i, t)$. Therefore, no path P_j contains x_{i+1}, for $h \leq j \leq i$. This, in addition to $H'_1(i)$, implies that P_j contains exactly one vertex of X_{i+1} and it is

x_j. From $H_3'(i)$ and Lemma 5, we know that P_{i+1} does not pass through a vertex of X_i. So, P_{i+1} contains exactly one vertex of X_{i+1} and it is x_{i+1}. In summary, $H_1'(i+1)$ holds.

We prove hypothesis $H_3'(i+1)$ in this paragraph. The vertices that are in $X_{i+1} \cup R(X_{i+1}, t)$ but not in $X_i \cup R(X_i, t)$ belong to the section $P_{i+1}^{(u_{i+1}^*, t)}$. Let us consider two scenarios depending on whether P_{i+1} is the current shortest (s,t)-path or a bypass:

- If P_{i+1} is the current shortest (s,t)-path, then the distance traversed on it must be larger than $\frac{1}{\beta}\omega_{\text{opt}}$. So, the cost of the shortest path between s and any vertex in $P_{i+1}^{(u_{i+1}, t)}$ is greater than $\frac{1}{\beta}\omega_{\text{opt}} > \frac{\sqrt{2}}{2\beta}\omega_{\text{opt}}$.
- If P_{i+1} is a bypass, the blocked edge b_{i+1} is also on its fictive path F_{i+1}, according to Lemma 4. Moreover, the cost of the section between s and u_{i+1} on F_{i+1} is greater than $\frac{1}{\beta}\omega(F_{i+1})$ because the transition vertex of P_{i+1} is closer to t than any transition vertex of F_{i+1} as a bypass of itself. As $\omega(F_{i+1}) \geq \frac{\sqrt{2}}{2}\omega_{\text{opt}}$, we know that the cost of F_{i+1} between s and u_{i+1} is larger than $\frac{\sqrt{2}}{2\beta}\omega_{\text{opt}}$.

It follows from $H_3'(i)$ that the shortest path between s and any vertex of $X_{i+1} \cup R(X_{i+1}, t)$ has a cost larger than $\frac{\sqrt{2}}{2\beta}\omega_{\text{opt}}$, *i.e.* $H_3'(i+1)$ is satisfied.

Case 2: We assume here that v_{i+1} is located before the blockage b_{i+1} (Fig. 10b). We denote by v_{i+1}^- the predecessor of v_{i+1} on P_{i+1}. We fix $x_{i+1} = v_{i+1}^-$.

First, we assume that $x_{i+1} = v_{i+1}^- \in R(X_i, s)$. This means that vertex x_{i+1} was already visited. Consequently, $P_{i+1}^{(x_{i+1}, t)}$ is not a detour, otherwise it would have already been traversed. That is, $\omega(P_{i+1}^{(x_{i+1}, t)}) > \frac{1}{\beta}\omega_{\min} \geq \frac{\sqrt{2}}{2\beta}\omega_{\text{opt}}$. Using this, we derive the following inequalities:

$$\omega(P_{i+1}^{(s,x_{i+1})}) = \omega_{i+1} - \omega(P_{i+1}^{(x_{i+1},t)}) < \beta\omega_{\text{opt}} - \frac{\sqrt{2}}{2\beta}\omega_{\text{opt}}, \tag{13}$$

$$\leq \frac{\sqrt{2}}{\beta}\omega_{\text{opt}} - \frac{\sqrt{2}}{2\beta}\omega_{\text{opt}} \leq \frac{\sqrt{2}}{2\beta}\omega_{\text{opt}} \leq \frac{1}{\beta}\omega_{\min}. \tag{14}$$

The cost of P_{i+1} is $\omega_{i+1} \leq \beta\omega_{\text{opt}}$, whether P_{i+1} is a shortest path or a bypass, which implies (13). Inequality (14) comes from the fact that $\beta \leq \frac{\sqrt{2}}{\beta}$.

Moreover, vertex v_{i+1}, which is in $R(X_i, t)$, the successor of x_{i+1} on P_{i+1}, belongs to a former exploration path P_j and it is located after the blockage on P_j, otherwise, it would not be on P_{i+1} (Lemma 5). As a conclusion, P_{i+1} is a bypass of either P_j or the fictive path F_j (if P_j is a bypass) and its transition vertex is closer to t than the transition vertex of P_j. We obtain a contradiction: P_{i+1} should have been traversed during a previous round.

Given this contradiction, we know that $v_{i+1}^- \notin R(X_i, s)$. So, it does not belong to G_i. Set $X_{i+1} = X_i \cup \{x_{i+1}\}$ is an (s,t)-cut, where each exploration path contains exactly one of its vertices and each element of the cut arrives before the blockage: $H_1'(i+1)$ and $H_2'(i+1)$ are satisfied.

Let P_j, $h \leq j \leq i$ be the last exploration path traversed containing v_{i+1}^- but v_{i+1} was not visited by the traveller on it (Lemma 5). We prove by contradiction that if the cost of P_{i+1} between s and v_{i+1}^- is at most $\frac{\sqrt{2}}{2\beta}\omega_{\text{opt}} \leq \frac{1}{\beta}\omega_{\min}$, then P_{i+1} should have been the bypass traversed instead of P_j during step j. To justify this assertion, we distinguish two scenarios as follows:

- If P_j was a shortest (s,t)-path, then P_{i+1} would have been its bypass. On the one hand, v_{i+1} is a descendant of the blocked edge (u_j, u_j^*) on P_j, since v_{i+1} was not visited during the exploration of P_j. On the other hand, the transition vertex of P_j (as a bypass on itself) appears before the blocked edge. Hence, the transition vertex of P_{i+1}, v_{i+1}, is closer to t than the transition vertex of P_j itself.
- If P_j was a bypass of F_j, then a similar argument works. Indeed, the transition vertex of P_j on F_j is located before the blocked edge (u_j, u_j^*), so it is an ancestor of v_{i+1}.

In both cases, the transition vertex of P_{i+1} is closer to t than the transition vertex of P_j itself. This contradiction implies $H_3'(i+1)$.

In summary, when v_{i+1} is located before the blocked edge b_{i+1}, we necessarily have $v_{i+1}^- \notin R(X_i, s)$ and we proved that taking $x_{i+1} = v_{i+1}^-$ in this case satisfies the hypotheses $H_1'(i+1)$, $H_2'(i+1)$, and $H_3'(i+1)$. □

This theorem ensures the existence of a minimal vertex (s,t)-cut $X = X_\ell$ of cardinality $\ell - h + 1$. As a consequence, $\ell - h + 1 \leq \lambda_{\max}^V$. The cost of any detour and any exploration path which is not in $P_h, \ldots, P_\ell$ is at most $\frac{1}{\beta}\omega_{\text{opt}}$. Let k_1 be the blockages which are on paths $P_h, \ldots, P_\ell$ and $k_2 = k - k_1$. Now, we can state the competitive ratio of BYPASS.

Theorem 10. *The competitive ratio of* BYPASS *is upper-bounded by:*

$$2\beta\lambda^V_{max} + \frac{2}{\beta}(k - \lambda^V_{max}) + 1. \tag{15}$$

As β must be less than or equal to $2^{\frac{1}{4}}$, the best competitive slope is obtained when $\beta = 2^{\frac{1}{4}} = \sqrt{\sqrt{2}}$ and the competitive ratio is at most:

$$2\sqrt{\sqrt{2}}\lambda^V_{max} + 2^{\frac{3}{4}}(k - \lambda^V_{max}) + 1.$$

Proof. The cost of the traversal of "short paths", *i.e.* paths which are not in collection $S_2 = P_h, \dots, P_\ell$, is at most $\frac{2}{\beta}k_2\omega_{\text{opt}}$. The cost of each path in S_2 is at most $\beta\omega_{\text{opt}}$ because of the existence of bypasses. Therefore, the distance traversed on paths $P_h, \dots, P_\ell$ is at most $2\beta k_1\omega_{\text{opt}}$. It follows that the total distance traversed is upper-bounded by:

$$2\beta k_1\omega_{\text{opt}} + \frac{2}{\beta}k_2\omega_{\text{opt}} + \omega_{\text{opt}} = 2\beta k_1\omega_{\text{opt}} + \frac{2}{\beta}(k - k_1)\omega_{\text{opt}} + \omega_{\text{opt}},$$

with an extra term ω_{opt} which is at most the distance traversed on the path when the traveller reaches t. Since $k_1 = \ell - h + 1 \leq \lambda^V_{\max}$ and $\beta \geq 1$, the total distance is upper-bounded by $2\beta\lambda^V_{\max}\omega_{\text{opt}} + \frac{2}{\beta}(k - \lambda^V_{\max})\omega_{\text{opt}}$, which gives us Equation (15) as an upper-bound of the competitive ratio. □

6. Perspectives

Both strategies DETOUR and BYPASSES achieve a competitive bound asymptotically smaller than the global one $2k + 1$ on some families of bounded-cut graphs. A natural improvement of our contributions would be to design strategies with smaller competitive slopes, *i.e.* $\rho < \sqrt{2}$ when $\mu^E_{\max}$ is bounded or $\rho < \sqrt{\sqrt{2}}$ when $\lambda^V_{\max}$ is bounded. A significant breakthrough would consist in identifying a family of bounded-cut graphs which admits deterministic strategies with a constant competitive ratio, *i.e.* independently of k.

In our study, we focused on outerplanar graphs as they satisfy $\lambda^V_{\max} \leq 2$. For this reason, BYPASS performs well on these graphs with a competitive slope $\sqrt{\sqrt{2}}$. We believe that a strategy with a smaller competitive ratio could be designed specifically for outerplanar graphs. Indeed, our reasoning only uses the fact that $K_{2,3}$ is a forbidden minor of outerplanar graphs. The inequality $\lambda^V_{\max} \leq 2$ could suggest the existence of strategies more efficient that the one treated in this article achieving $\rho k + O(\lambda^V_{\max})$ in general.

An interesting extension of this work would be to expand the list of graph parameters such that, when these parameters are small, efficient strategies can be designed. We know from our study that this list contains both cut parameters $\mu^E_{\max}$ and $\lambda^V_{\max}$. We wonder if parameters coming from other horizons - different from cuts - could be used to design deterministic strategies outperforming ratio $2k + 1$. We can already state that some well-known notions on graphs do not belong to this list:

- **Tree decompositions**: the treewidth and pathwidth of the pathological graph W_k are equal to 2 for any $k \in \mathbb{N}$.
- **Degrees**: the degeneracy of graph W_k is equal to 2. Moreover, there is a binary apex tree (maximum degree 2) with a structure similar to W_k for which ratio $2k + 1$ cannot be outperformed [5].
- **Planarity**: graph W_k is planar bipartite.

For all these parameters, even if they are constant, there is no hope of finding a deterministic strategy with competitive slope smaller than 2.

Eventually, the major open question in the field of k-CTP is the existence of a polynomial-time randomised strategy with competitive ratio $\rho k + O(1)$, $\rho < 2$ for general graphs. We believe that generalising the algorithm of Bender and Westphal [3] for node-disjoint paths is the best approach to answer this question. The difficulty lies in determining an implicit representation of the graph on which random draws could be handled recursively, as it was done for simple disjoint (s, t)-paths.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] V. Aksakalli, O.F. Sahin, I. Ari, An AO* based exact algorithm for the Canadian traveler problem, INFORMS J. Comput. 28 (1) (2016) 96–111.
[2] A. Bar-Noy, B. Schieber, The Canadian Traveller Problem, in: Proc. of ACM/SIAM SODA, 1991, pp. 261–270.

[3] M. Bender, S. Westphal, An optimal randomized online algorithm for the k-Canadian Traveller Problem on node-disjoint paths, J. Comb. Optim. 30 (1) (2015) 87–96.
[4] P. Bergé, J. Desmarchelier, W. Guo, A. Lefebvre, A. Rimmel, J. Tomasik, Multiple canadians on the road: minimizing the distance competitive ratio, J. Comb. Optim. 38 (4) (2019) 1086–1100.
[5] P. Bergé, J. Hemery, A. Rimmel, J. Tomasik, On the competitiveness of memoryless strategies for the k-Canadian Traveller Problem, in: Proc. of COCOA, 2018, pp. 566–576.
[6] P. Bergé, L. Salaün, Improved deterministic strategy for the Canadian traveller problem exploiting small max-(s, t)-cuts, in: Proc. of WAOA, 2019, pp. 29–42.
[7] Z. Bnaya, A. Felner, S.E. Shimony, Canadian traveler problem with remote sensing, in: Proc. of IJCAI, 2009, pp. 437–442.
[8] A. Borodin, R. El-Yaniv, Online Computation and Competitive Analysis, Cambridge Univ. Press, 1998.
[9] G. Chartrand, F. Harary, Planar permutation graphs, Ann. Inst. Henri Poincaré B, Probab. Stat. 3 (1967) 433–438.
[10] E.D. Demaine, Y. Huang, C. Liao, K. Sadakane, Approximating the Canadian traveller problem with online randomization, Algorithmica 83 (5) (2021) 1524–1543.
[11] R. Diestel, Graph Theory, 4th edition, Graduate Texts in Mathematics, vol. 173, Springer, 2012.
[12] E.W. Dijkstra, A note on two problems in connexion with graphs, Numer. Math. 1 (1) (1959) 269–271.
[13] H. Eto, T. Hanaka, Y. Kobayashi, Y. Kobayashi, Parameterized algorithms for maximum cut with connectivity constraints, in: Proc. of IPEC, vol. 148, 2019, pp. 13:1–13:15.
[14] D. Karger, E. Nikolova, Exact algorithms for the Canadian Traveller Problem on paths and trees, Technical report, MIT CSAIL, 2008.
[15] C. Liao, Y. Huang, The covering Canadian traveller problem, Theor. Comput. Sci. 530 (2014) 80–88.
[16] C. Papadimitriou, M. Yannakakis, Shortest paths without a map, Theor. Comput. Sci. 84 (1) (1991) 127–150.
[17] D. Shiri, F.S. Salman, On the online multi-agent $O-D$ k-Canadian Traveler Problem, J. Comb. Optim. 34 (2) (2017) 453–461.
[18] D. Shiri, F.S. Salman, On the randomized online strategies for the k-Canadian traveler problem, J. Comb. Optim. (2019) 1–14.
[19] S. Westphal, A note on the k-Canadian Traveller Problem, Inf. Process. Lett. 106 (3) (2008) 87–89.
[20] Y. Xu, M. Hu, B. Su, B. Zhu, Z. Zhu, The Canadian traveller problem and its competitive analysis, J. Comb. Optim. 18 (2) (2009) 195–205.