

Parameterized Complexity of Directed Traveling Salesman Problem

Václav Blažej   



Institute of Informatics, University of Warsaw, Warsaw, Poland

Andreas Emil Feldmann   

University of Sheffield, Sheffield, United Kingdom

Foivos Fioravantes  

Faculty of Information Technology, Czech Technical University in Prague, Prague, Czech Republic

Paweł Rzażewski  

Warsaw University of Technology, Warsaw, Poland

Institute of Informatics, University of Warsaw, Warsaw, Poland

Ondřej Suchý  

Faculty of Information Technology, Czech Technical University in Prague, Prague, Czech Republic

Abstract

The DIRECTED TRAVELING SALESMAN PROBLEM (DTSP) is a variant of the classical TRAVELING SALESMAN PROBLEM in which the edges in the graph are directed and a vertex and edge can be visited multiple times. The goal is to find a directed closed walk of minimum length (or total weight) that visits every vertex of the given graph at least once. In a yet more general version, DIRECTED WAYPOINT ROUTING PROBLEM (DWRP), some vertices are marked as terminals and we are only required to visit all terminals. Furthermore, each edge has its capacity bounding the number of times this edge can be used by a solution.

While both problems (and many other variants of TSP) were extensively investigated, mostly from the approximation point of view, there are surprisingly few results concerning the parameterized complexity. Our starting point is the result of Marx et al. [APPROX/RANDOM 2016] who proved that DTSP is $W[1]$ -hard parameterized by distance to pathwidth 3. In this paper we aim to initiate the systematic complexity study of variants of DIRECTED TRAVELING SALESMAN PROBLEM with respect to various, mostly structural, parameters.

We show that DWRP is FPT parameterized by the solution size, the *feedback edge number* and the *vertex integrity* of the underlying undirected graph. Furthermore, the problem is XP parameterized by treewidth.

On the complexity side, we show that the problem is $W[1]$ -hard parameterized by the distance to constant treedepth.

2012 ACM Subject Classification Theory of computation \rightarrow Parameterized complexity and exact algorithms; Theory of computation \rightarrow Graph algorithms analysis

Keywords and phrases Directed TSP, parameterized complexity, vertex integrity, treedepth

Related Version *Full Version:* <https://arxiv.org/abs/2506.22127> [15]

Funding *Václav Blažej and Paweł Rzażewski:* The work on this article is a part of project BOBR that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 948057).

Ondřej Suchý: Co-funded by the European Union under the project Robotics and advanced industrial production (reg. no. CZ.02.01.01/00/22_008/0004590).

Acknowledgements The project was initiated at the workshop Homonolo 2023. We are grateful to the organizers and other participants for a nice and productive atmosphere.

1 Introduction

The TRAVELING SALESMAN PROBLEM (TSP) is possibly the most famous combinatorial optimization problem. It is one of the few problems having entire monographs devoted solely to its aspects [6, 21, 39, 58]. In this paper we focus on the directed variant, which can be formulated as follows. The input instance is a directed graph G with edges assigned positive integers (*weights*) and an integer b , and we are to decide whether there is a closed walk in G visiting every vertex *at least* once and having total weight at most b . We refer to this problem as DIRECTED TRAVELING SALESMAN PROBLEM (DTSP). We also consider two natural generalizations of the problem. In the first of them, DIRECTED SUBSET TSP (DSTSP), we are given alongside the graph also a subset of its vertices W called *waypoints* or *terminals* and the walk is only required to visit these vertices. Obviously, an instance of DTSP can be interpreted as the instance of DSTSP by letting $W = V(G)$. In some applications it is also preferred not to traverse any (or some) of the edges of the graph too often. To this end, in DIRECTED WAYPOINT ROUTING PROBLEM (DWRP) (introduced in [3]), we augment the input with a positive integer for each edge specifying its *capacity*. The walk is then required to traverse each edge at most as many times as its capacity. A simple argument (provided later) shows that each edge is used at most $n := |V(G)|$ times in any solution. Hence, setting the capacities to n is equivalent to dropping the capacity constraints completely.

Our focus is on exact parameterized algorithms for the problem. Let us first summarize the known results for the undirected variant of the problem. Since HAMILTONIAN CYCLE (HAMC) is NP-hard even on planar graphs of maximum degree three [35], TSP is para-NP-hard with respect to (w.r.t.) the maximum degree of the input graph as well as w.r.t. the maximum number of visits to any vertex. As the solution size is equal to the number of vertices, the dynamic programming of Bellman [10] and Held and Karp [40] which runs in $O(2^n \cdot n^2)$ time can be considered as an FPT algorithm w.r.t. solution size. It can be also used to solve SUBSET TSP in $O(2^{|W|} \cdot |W|^2)$ after an initial computation of distances between all pairs of waypoints. This can be improved to $O(2^{\sqrt{|W|} \log |W|} n^{O(1)})$ if the graph is planar and the weights are bounded by a polynomial [44].

On the one hand, a single-visit variant of TSP is known to be FPT w.r.t. treewidth (see also [25]) and this result also extends to our multi-visit variant and to WAYPOINT ROUTING (WPR) [56]. On the other hand, TSP is hard on cliques with edge weights 1 and 2, shown by a simple reduction from HAMC. Therefore, it is para-NP-hard w.r.t. any parameter which is constant on cliques. Moreover, already HAMC is W[1]-hard w.r.t. cliquewidth [33] (while FPT w.r.t. the neighborhood diversity [47]).

Recently, the (undirected) TSP was studied from the perspective of kernelization [13, 14]. It was shown to admit a polynomial kernel w.r.t. the feedback edge number and w.r.t. distance to constant size components. Here, for a graph class \mathcal{C} , *distance to \mathcal{C}* refers to the minimum number of vertices that must be deleted from the input graph to obtain a graph that belongs to \mathcal{C} . In the above case, the class contains all graphs with every connected component of size at most c , for any fixed c . TSP was shown not to admit polynomial kernel w.r.t. vertex integrity, unless $\text{NP} \subseteq \text{coNP/poly}$, whereas SUBSET TSP does not admit polynomial kernel even w.r.t. distance to disjoint cycles, under the same assumption [14]. Here *vertex integrity* of a graph is the minimum k such that we can delete at most k vertices from the graph in such a way that each connected component of the resulting graph has at most k vertices.

Much less is known about the parameterized complexity of the directed case of the problem. The mentioned hardness results translate to the directed case by symmetrically orienting the input graph. While the algorithmic results do not transfer automatically to the

directed setting, it is often possible to adapt them. For example, the dynamic programming w.r.t. number of vertices [10, 40] applies also to this case. This can be again improved to $O(2^{\sqrt{|W|}} \log |W| n^{O(1)})$ for DsTSP in planar graphs [51]; also lifting the weight-restriction from the undirected case. In contrast, DWRP is NP-hard already for two waypoints [3].

Similarly, the single visit variant is FPT w.r.t. the treewidth. However, Marx et al. [52] showed, that DTSP is W[1]-hard w.r.t. the pathwidth pw of the input graph and there is no algorithm for DTSP running in $f(\text{pw}) \cdot n^{o(\text{pw})}$ time for any computable function f , unless the Exponential Time Hypothesis (ETH) [41, 42] fails. More precisely, the proof shows that this is the case even w.r.t. distance to graphs of pathwidth 3. Furthermore, a more careful analysis of the proof (see Theorem 26 in the appendix) shows that the W[1]-hardness holds true even w.r.t. distance to pathwidth 2. Thus, the fpt-algorithm for WRP w.r.t. treewidth cannot be transferred to the directed case.

Apparently, these are the only results concerning the parameterized complexity of DTSP, and, in particular, no fpt-algorithms w.r.t. structural parameters are known for the problem.

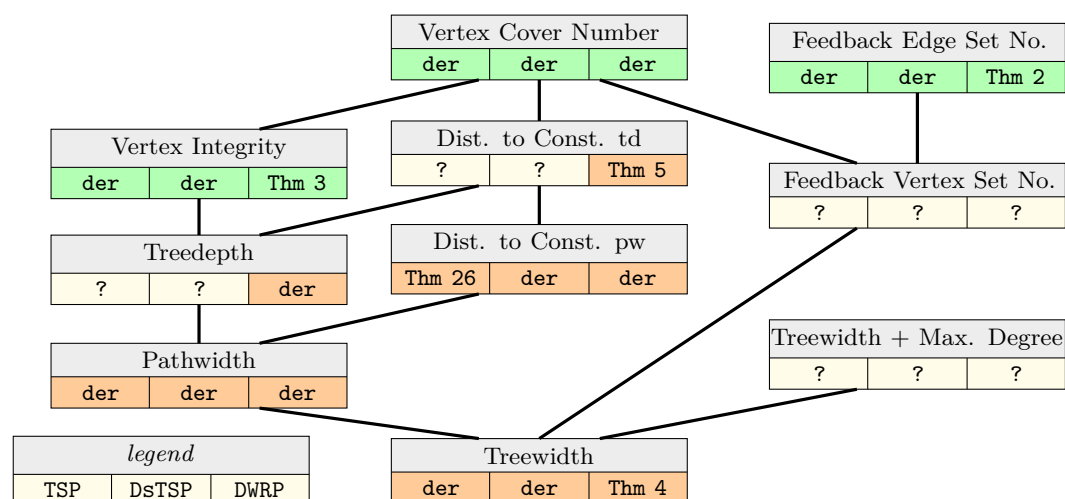


Figure 1 Complexity picture for the studied problems. Box fill shows tractability: green stands for FPT, orange for W[1]-hard and in XP, and yellow for in XP, but open whether FPT or W[1]-hard. The boxes either contain references to the corresponding theorem, or *der* if the result can be derived from the other results. All the hardness results apply even for unit weights, whereas the algorithms do not assume any restriction on the weights.

Our Contribution. We initiate the systematic study of the parameterized complexity of DTSP and its generalizations focusing on structural parameters. In particular, we concentrate on parameters measuring the structure of the underlying undirected graph of the input. We aim to fill the apparent gap in knowledge and address the lack of fpt-algorithms.

First, we revisit the standard parameterization, i.e., the parameterization by the solution size. The budget bounds the number of edges in the solution, hence also the number of vertices visited, and in particular the size of W (for any yes-instance of DTSP and DsTSP) Thus for these problems, the known single exponential algorithms [10, 40] and subexponential algorithm for planar graphs [51] show that the problem is FPT w.r.t. the budget b . For DWRP the situation is not as straightforward. Nevertheless, we show the following result.¹

¹ Proofs of statements marked with (★) can be found in the appendix.

► **Theorem 1** (★). *In $e^k 4^k k^{O(\log k)} \cdot m \log m$ time we can decide whether the given m -edge instance of DWRP has a solution consisting of at most k edge occurrences.*

Next we turn to the structural parameters (see Figure 1 for an overview). The result of Marx et al. [52] shows that DTSP is already W[1]-hard w.r.t. distance to graphs of pathwidth 3, i.e., for graphs with rather restricted structure. In fact, as we show in Theorem 26 (appendix), this holds also for unweighted graphs. Therefore, in search for fpt-algorithms for the problems, the structure of the graph after removal of parameter many vertices or edges should be more restricted than having pathwidth 3. Our first positive result is w.r.t. the feedback edge number, i.e., the minimum number of edges that needs to be removed to obtain a forest.

► **Theorem 2.** *DWRP can be solved in $k^{O(k)} + n^{O(1)}$ time, where k is the feedback edge number of the input graph.*

Then, we present our main algorithmic result, showing that DWRP is FPT w.r.t. vertex integrity. Recall that vertex integrity allows, roughly speaking, to delete parameter many vertices so as to obtain a graph with connected components of parameter size.

► **Theorem 3.** *DWRP can be solved in $k^{O(k^6)} \cdot n \cdot \log^2 n (\log n + \log U)$ time, where k is the vertex integrity of the input graph and U is the maximum weight of an edge in G .*

The algorithm first guesses a part of the solution ensuring its connectivity and then employs a well structured ILP (*N-fold IP* [45]) to find the cheapest way to complete the initial guess into a solution.

We also show that, while the problems are probably not FPT w.r.t. treewidth alone, they are FPT w.r.t. treewidth combined with the bound on the number of visits to every city.

► **Theorem 4** (★). *DWRP on n -vertex graphs of treewidth t can be solved in $n^{O(t)}$ time.*

This means that all considered problems are in XP w.r.t. treewidth and, hence, w.r.t. all structural parameters considered in this paper.

We complement the positive results by showing that DWRP is W[1]-hard w.r.t. distance to constant treedepth, i.e., if the input graph becomes of constant treedepth after deleting parameter many vertices.

► **Theorem 5.** *DWRP is W[1]-hard with respect to modulator to constant treedepth, even if all weights are 1.*

Note that this result is incomparable to that of Marx et al. [52] in that it holds for a more restricted graph class, while it only applies to a more general problem.

Further Related Work. Another popular formulation of (DIRECTED) TSP is that the input graph is complete and, thus, we are given a matrix of weights for all ordered pairs of vertices. Since this matrix is asymmetric for the case of directed graphs, the directed variant of TSP is often called ASYMMETRIC TSP (ATSP). It is easy to switch from graph G to a complete graph on the same set of vertices by performing the so called *metric closure*, i.e., assigning the edge (u, v) of the complete graph the weight given by distance from vertex u to vertex v in the graph G . The weights obtained this way satisfy *triangle inequality*. Whenever this is the case, there is no reason to revisit any vertex, as it cannot make the tour shorter. However, if the matrix is not required to satisfy the triangle inequality, it makes a difference whether visiting a vertex more than once is allowed or not (single visit version).

For the single visit version with general matrices a folklore result shows that it cannot be approximated within any polynomial factor, unless $P=NP$ [22]. For undirected case with triangle inequality it was known that there is $\frac{3}{2}$ -approximation since 1976 [20], which was later improved [36, 53, 54] up to the currently best known factor 1.4 [57] for the graphic case, where the weights originate from the distances in an unweighted graph. However, for the directed case, it was long open, whether a constant factor approximation exists for matrices satisfying the triangle inequality [5, 7]. Constant-factor approximation algorithms appeared only recently, first for unweighted directed graphs [59], later for general distances [60], with the currently best approximation factor being $22 + \varepsilon$ for any $\varepsilon > 0$ [61]. The best known lower bound on the approximability is $\frac{75}{74}$ [43].

It is surprising, given the lack of parameterized analysis for the problem, that there are several studies concerning parameterized approximation algorithm for DTSP. First, Böckenhauer [16] considered (undirected) TSP with deadlines on the latest time to visit some of the vertices and showed a 2.5-approximation algorithm in $k! \cdot n^{O(1)}$ time, where k is the number of deadlines, and complemented that by several inapproximability results. The main result of the already mentioned paper of Marx et al. [52] is a polynomial-time constant-factor approximation for DTSP in nearly embeddable graphs, where both the factor and the time depend on the actual parameters of the class of graphs considered. Bonnet et al. [19] provided a $(\log r)$ -approximation for ATSP in $2^{O(\frac{n}{r})}$ time for any r . Behrendt et al. [9] considered the case that there are few edges with (significantly) asymmetric costs and presented a 2.5-approximation parameterized by the vertex cover number of the graph formed by such edges and a 3-approximation parameterized by the minimum number of asymmetric edges in a minimum arborescence of the graph.

Further results are known for a variant of DWRP where the waypoints have a prescribed order [2] and for the undirected case [4].

Another studied variant of DTSP is the MANY-VISIT TSP, where each vertex has a prescribed number of visits and the task is to find a closed walk which visits each vertex *exactly* the given number of times. After a significant effort, an algorithm was found for the problem that is exponential only in the number of vertices, i.e., independent of the demands which can be huge [12, 46]. The problem was recently studied from parameterized perspective [49] under the name CONNECTED FLOW. Here only a subset of vertices has a prescribed number of visits, the other vertices can be visited arbitrary number of times (or not visited at all) and edges have capacities (as in DWRP). They showed that the problem is para-NP-hard w.r.t. number of vertices with demand even in unweighted graph, while FPT w.r.t. this parameter, if there are no capacities. Further, they gave a polynomial kernel w.r.t. vertex cover number, an $n^{O(\text{tw})}$ -time algorithm and a matching $n^{o(\text{tw})}$ -time ETH lower bound.

Several parameterized and fine-grained complexity studies considered the local search for the undirected single-visit TSP parameterized the size of the neighborhood to be searched [11, 18, 24, 37, 48, 50]. None of these studies seems to extend to directed graphs.

2 Notations, Preliminaries, and Basic Observations

We follow standard graph-theoretic notation [26]. Also, we refer the reader to [23] for a textbook on parameterized complexity.

(Generalized) N -fold integer programming (N -fold IP) is the problem of minimizing a separable convex objective (for us it suffices to minimize a linear objective) over a set

of constraints of a specific form. Let $r \in \mathbb{N}$ and $s_i, t_i \in \mathbb{N}$ for every $i \in [N]$. The IP has $d = \sum_{i \in [N]} t_i$ variables partitioned into N so-called *bricks*. The i -th brick is denoted $x^{(i)}$ and contains t_i variables. The constraints have the following form:

$$\begin{aligned} D_1 x^{(1)} + D_2 x^{(2)} + \dots + D_N x^{(N)} &= \mathbf{b}_0 && \text{(linking (global) constraints)} \\ A_i x^{(i)} &= \mathbf{b}_i && \forall i \in [N] \text{ (local constraints)} \\ \mathbf{l}_i \leq x^{(i)} \leq \mathbf{u}_i &&& \forall i \in [N] \text{ (box constraints)} \end{aligned}$$

where we have $D_i \in \mathbb{Z}^{r \times t_i}$, $A_i \in \mathbb{Z}^{s_i \times t_i}$, $\mathbf{b}_0 \in \mathbb{Z}^r$, $\mathbf{b}_i \in \mathbb{Z}^{s_i}$ and $\mathbf{l}_i, \mathbf{u}_i \in \mathbb{Z}^{t_i}$ for every $i \in [N]$. Let us denote $s = \max_{i \in [N]} s_i$ and recall that the dimension is $d = \sum_{i \in [N]} t_i$. There are r global constraints and, apart from these, each variable is only involved in at most s local constraints and one box constraint. We call variables not appearing in linking constraints *local*.

The current best algorithm solving the N -fold IP in $(rs\Delta)^{O(r^2s+rs^2)} d \log d \log D \log \text{val}_{\max}$ time is by Eisenbrand et al. [30, Cor. 97]; see also [29, 45], where D is the maximum range of a variable, $\Delta = \max_{i \in [N]} \left(\max(\|D_i\|_\infty, \|A_i\|_\infty) \right)$, and val_{\max} is the maximum feasible value of the objective.

Variants of the Problem and Basic Observations. Let us formally introduce the considered problems and establish notation used in the next sections.

DIRECTED TRAVELING SALESMAN PROBLEM (DTSP)

Input: A directed graph $G = (V, E)$, edge weights $\omega: E \rightarrow \mathbb{N}$, and a budget $b \in \mathbb{N}$.

Question: Is there a closed directed walk in G of total weight at most b that traverses each vertex of G at least once?

DIRECTED WAYPOINT ROUTING PROBLEM (DWRP)

Input: A directed graph $G = (V, E)$, set of waypoints $W \subseteq V$, edge weights $\omega: E \rightarrow \mathbb{N}$, edge capacities $\kappa: E \rightarrow \mathbb{N}$, budget $b \in \mathbb{N}$.

Question: Is there a closed directed walk C in G of total weight at most b , that traverses each vertex in W (at least once) and such that for each edge e the number of times C traverses e is at most $\kappa(e)$?

To avoid degenerate cases, we assume that $|W|$ is always at least 2 and that the input graph is strongly connected. We use undirected parameters, and, when speaking about a width of a directed graph, we mean the width of the underlying undirected graph.

► **Lemma 6** (Folklore, see, e.g., Bang-Jansen and Gutin [8, Exercise 1.12]). *Any closed directed walk (in particular a solution to the problem) can be decomposed into simple cycles, such that the number of times an edge is traversed by the walk equals the number of the cycles it appears in.*

► **Observation 7** (\star). *In each optimal solution to DWRP, each vertex v is visited at most $|W \setminus \{v\}|$ times.*

Hence, an instance of DsTSP can be interpreted as an equivalent instance of DWRP, by letting $\kappa(e) = |V|$ for every arc.

3 DWRP is FPT with Respect to the Feedback Edge Number

In this section we prove Theorem 2 that we restate here.

► **Theorem 2.** *DWRP can be solved in $k^{O(k)} + n^{O(1)}$ time, where k is the feedback edge number of the input graph.*

Proof. Let $(G = (V, E), W, \omega, \kappa, b)$ be an instance of DWRP, where $n = |V|$. Let G' be the underlying undirected graph of G . Let $F \subseteq E(G')$ be a feedback edge set of G' of size at most k , i.e., $G' \setminus F$ is a forest. Note that it can be computed in polynomial time.

Our algorithm consists of three phases. First, we exhaustively apply some straightforward reduction rules in order to simplify (or already reject) the input. Then, we branch into $2^{O(k)}$ of possibilities, guessing the structure of the solution. Finally, for each such guess, we create an equivalent instance of DWRP with $O(k)$ vertices and edges. The original instance is a yes-instance if and only if at least one of the created instances is a yes-instance. Thus, our algorithm can be seen as a compression to OR of $2^{O(k)}$ instances, each of size $O(k)$. By Observation 7, each such instance can be solved in time $k^{O(k)}$ by brute force. Thus, the overall running time is $k^{O(k)} + n^{O(1)}$.

Preprocessing. The following reduction rule lets us deal with vertices of degree 1 in G' .

- **Reduction Rule 1.** *Suppose there is a vertex v in G which has only one neighbor u in G' .*
- *If $v \notin W$, remove v from the graph.*
 - *If $v \in W$ and both $(u, v) \in E$ and $(v, u) \in E$, then remove the vertex v , include u into W , and decrease the budget by $\omega((u, v)) + \omega((v, u))$.*
 - *If $v \in W$ and one of the arcs $(u, v), (v, u)$ is not present, then reject the instance.*

The correctness of the rule is straightforward. Note that, as $\{u, v\}$ is not part of any cycle in G' , we may assume that $\{u, v\} \notin F$, and thus F is still a feedback edge set of the graph resulting from the reduction rule.

We apply the reduction rule exhaustively, and, for simplicity, we keep denoting the resulting instance by $(G, W, \omega, \kappa, b)$ and the underlying undirected graph by G' .

Let R be the set of all vertices that are incident to the edges in F . Note that since the reduction rule was exhaustively applied, each vertex has degree at least two in G' . Therefore, each leaf of $G'' = G' \setminus F$ is in R and, thus, G'' has at most $2k$ leaves. Let D be the set of vertices that have degree at least 3 in G'' . As G'' is a forest with at most $2k$ leaves, the size of D is at most $2k - 1$. We also define $X = R \cup D$, and we have $|X| \leq 4k - 1$.

Let \mathcal{P}' contain all paths in G'' with both endpoints in X and all internal vertices outside X . In other words, the paths in \mathcal{P}' correspond to the edges of the forest obtained from G'' by contracting all vertices of degree 2. As there are at most $4k - 1$ vertices in X and, as G'' is a forest, the number of paths in \mathcal{P}' is at most $4k - 2$. Finally, we define $\mathcal{P} = \mathcal{P}' \cup F$; we clearly have $|\mathcal{P}| \leq 5k - 2$.

Types of paths and their costs. Consider an optimum solution C^* , i.e., a shortest (in terms of the sum of weights) closed walk that visits every vertex from W in G . Among all such solutions, pick one with the minimum number of edges.

Consider a path $P \in \mathcal{P}$ with endpoints $u, v \in X$. A *visit* is an inclusion-wise maximal set of orientations of edges from P that appear consecutively in C^* . A visit is a *pass* if it contains for every edge of P exactly one of its orientations and it contains it exactly once, i.e., it corresponds to traversing P once, either from u to v or from v to u .

Notice that a path P might be of one of the following types:

- $(u \curvearrowright v)$ there is at least one visit and every visit is a pass from u to v ,
- $(u \curvearrowleft v)$ there is at least one visit and every visit is a pass from v to u ,
- $(u \rightleftarrows v)$ every visit is a pass and there is at least one pass in each direction,
- $(u \leftarrow\!\!\!\rightarrow v)$ none of the visits is a pass, there is a visit that starts in u and ends in u , but no visit that starts in v and ends in v ,
- $(u \rightarrow\!\!\!\leftarrow v)$ none of the visits is a pass, there is a visit that starts in v and ends in v , but no visit that starts in u and ends in u ,
- $(u \circ v)$ none of the visits is a pass, there is a visit that starts in u and ends in u , and a visit that that starts in v and ends in v ,
- $(u \cdots v)$ the path P is not visited at all by C^* .

▷ **Claim 8 (★).** The types listed above are mutually exclusive and cover all possibilities.

Let us make some further observations about possible visits in P by C^* .

If no internal vertex of P is in W , then P is of type $u \curvearrowleft v$, $u \curvearrowright v$, $u \rightleftarrows v$, or $u \cdots v$. Indeed, any non-pass visit can be removed from C^* , contradicting the minimality.

Now consider a path P of type $u \leftarrow\!\!\!\rightarrow v$ (the case $u \rightarrow\!\!\!\leftarrow v$ is symmetric). We may safely assume that P is visited exactly once by C^* , as one of the visits (the one reaching furthest towards v) covers all vertices covered by all the visits altogether. Further, we may assume that this single visit starts in u , reaches the internal vertex of P that belongs to W and is closest to v , and then goes back to u . Thus, if P is of type $u \leftarrow\!\!\!\rightarrow v$ or $u \rightarrow\!\!\!\leftarrow v$, then the visit in P contributes to the total cost of the solution by a fixed amount that can be computed in advance. We denote this cost by $\text{cost}(P, u \leftarrow\!\!\!\rightarrow v)$ or $\text{cost}(P, v \rightarrow\!\!\!\leftarrow u)$, respectively.

Now consider a path P of type $u \circ v$. Similarly as before, we can assume that there are at least two vertices from W among internal vertices of P , as otherwise we can shorten C^* by removing one of the visits. Furthermore, we can assume that P is visited exactly twice: one visit starts at u , reaches some vertex $x_u \in W$ and goes back to u , and the other visit starts at v , reaches some vertex $x_v \in W$, and goes back to v . In particular, no edge between x_u and x_v is traversed by C^* . As before, in polynomial time we can find an optimal pair of visits that cover all vertices from W that are on P , and compute its contribution to the cost of the solution (here note that we need to be careful as some edges might be only available in one direction). We denote this cost by $\text{cost}(P, u \circ v)$.

Finally, each pass in a fixed direction has a fixed cost, equal to the sum of weights of edges in G corresponding to the edges of P , where we only choose edges in the correct direction. We denote it by $\text{cost}(P, u \curvearrowright v)$ or $\text{cost}(P, v \curvearrowleft u)$, depending on the direction. We also define $\text{capacity}(P, u \curvearrowright v)$ (and, symmetrically, $\text{capacity}(P, u \curvearrowleft v)$) to the minimum capacity of an edge of G corresponding to an edge of P , in the direction from u to v (from v to u), resp.

Note that due to the directions and capacities of edges, some types might not be available for P . We will define the cost associated with such an unavailable type to be ∞ .

Building compressed instances. Now we perform the branching. For each path P in \mathcal{P} with endvertices u, v , we guess its type $u \sim v$, where $\sim \in \{\curvearrowleft, \curvearrowright, \rightleftarrows, \leftarrow\!\!\!\rightarrow, \circ, \rightarrow\!\!\!\leftarrow, \cdots\}$. This results in $7^{|\mathcal{P}|} = 2^{O(k)}$ branches.

Consider one such branch. Now we build an instance $(G', W', \omega', \kappa', b')$ of DWRP. Initialize $b' = b$. We start with including all vertices from X to G' , and all vertices from $X \cap W$ to W' . Now let P be a path in \mathcal{P} with endvertices $u, v \in X$. We proceed, depending on \sim . If the type of P is $u \cdots v$, we do nothing. If the type of P is $u \leftarrow\!\!\!\rightarrow v$ ($\rightarrow\!\!\!\leftarrow$), we add u (v) to W' , and decrease b' by $\text{cost}(P, u \leftarrow\!\!\!\rightarrow v)$ ($\text{cost}(P, v \rightarrow\!\!\!\leftarrow u)$), respectively. If the type of P is $u \circ v$, we add both u and v to W' , and decrease b' by $\text{cost}(P, u \circ v)$.

Now consider the case that the type of P is $u \curvearrowright v$ (the case $u \curvearrowleft v$ is symmetric). We add a new vertex $x_{u \curvearrowright v}^P$ and edges $(u, x_{u \curvearrowright v}^P)$ and $(x_{u \curvearrowright v}^P, v)$. Vertices u, v , and $x_{u \curvearrowright v}^P$ are all included into W' . We define $\omega'((u, x_{u \curvearrowright v}^P)) = \text{cost}(P, u \curvearrowright v)$ and $\omega'((x_{u \curvearrowright v}^P, v)) = 0$. Furthermore, we set $\kappa'((u, x_{u \curvearrowright v}^P)) = \kappa'((x_{u \curvearrowright v}^P, v)) = \text{capacity}(P, u \curvearrowright v)$. If the type of P is $u \leftrightarrow$, we proceed as if was of both types $u \curvearrowright v$ and $u \curvearrowleft v$.

This concludes the construction of the instance. If $b' < 0$, we reject it immediately. Clearly, the number of vertices of G' is at most $|X| + 2|\mathcal{P}| \leq 14k$, and the number of edges is at most $4|\mathcal{P}| \leq 20k$. The discussion about correctness is postponed to appendix. ◀

4 DWRP is FPT Parameterized by Vertex Integrity

In this section we prove Theorem 3.

► **Theorem 3.** *DWRP can be solved in $k^{O(k^6)} \cdot n \cdot \log^2 n (\log n + \log U)$ time, where k is the vertex integrity of the input graph and U is the maximum weight of an edge in G .*

Proof. An undirected graph G with vertex integrity k has a k -modulator $M \subseteq V(G)$ of size at most k such that each connected component of $G - M$ is of size at most k . We assume that G is given along with M as it can be found in $k^{O(k)}n$ time [28, 31]. Thus we assume that we are given an instance $(G, W, \omega, \kappa, b)$ of DWRP, with a set M of size at most k , such that each weakly connected component of $G - M$ has size at most k . For brevity, we will speak about ‘components’ instead of ‘weakly connected components.’

In what follows, we start with introducing some notions and establishing some properties of an optimum solution. Then, after some branching, we build an instance of ILP that corresponds to a solution with the required properties.

Segments. A *segment* is a walk in G with at least two vertices, that starts and ends in a vertex of M and all internal vertices are not in M . We remark that there might be no internal vertices – such a segment is just an arc contained in M . We call such a segment *trivial*. However, if a segment is non-trivial (i.e., not trivial), all its internal vertices are contained in a single component of $G - M$; the segment is *associated* with the component.

Note that the endpoints of a non-trivial segment need not to be distinct. A segment s from u to v is *minimal* if there is no segment from u to v that uses only the arcs of s , contains all vertices of s that are in W , and is shorter than s .

Segments in an Optimum Solution. Let C^* be an optimum solution, i.e., a shortest (w.r.t. the sum of weights of edges) closed directed walk in G that visits every vertex of W . We treat C^* as a sequence of vertices. Clearly C^* can be decomposed into segments (where two consecutive segments overlap on corresponding endpoints); let \mathcal{S}^* be the collection of these segments. We emphasize that \mathcal{S}^* is a multiset, as some segments might be used more than once. We observe that we can assume that each segment in \mathcal{S}^* is minimal, for otherwise we can replace it with a shorter one with the same endpoints, still covering all of W .

Some segments in \mathcal{S}^* might not be needed to cover W , but are required to keep the connectivity of the solution. We call them *connectors*. We can identify connectors with the following procedure. We initialize $\mathcal{S}^\circ = \emptyset$ and proceed iteratively as follows. If there is a segment s in $\mathcal{S}^* \setminus \mathcal{S}^\circ$ such that the segments $\mathcal{S}^* \setminus \mathcal{S}^\circ \setminus \{s\}$ cover all vertices of W , we include s into \mathcal{S}° . We emphasize that here we work with multisets, so the ‘\’ operation respects the multiplicity of elements. The set \mathcal{S}° contains connectors. Note that \mathcal{S}° is not defined uniquely and it might depend on the ordering in which segments were selected.

Consider a connector, i.e., a segment s from u to v that is in \mathcal{S}° . We observe that we can safely assume that it is a u - v -path, as otherwise we can replace it with such a path obtaining a solution of no larger length. We call such segments *simple*.

Traversals. Now let us consider a component C of $G - M$. A *traversal* of C is a collection \mathcal{S}_C of segments associated with C that:

- together cover every vertex from $W \cap V(C)$,
- for every $s \in \mathcal{S}_C$ there is a vertex in $W \cap V(C)$ not covered by segments in $\mathcal{S}_C \setminus \{s\}$,
- every segment in \mathcal{S}_C is minimal,
- every edge e is used at most $\kappa(e)$ times by \mathcal{S}_C .

Note that the segments from $\mathcal{S}^* \setminus \mathcal{S}^\circ$ associated with C form a traversal. Indeed, the first condition follows since C^* covers all vertices from W and we moved segments to \mathcal{S}° if they were not required to cover $W \cap V(C)$. The second condition again follows from the definition of \mathcal{S}° . Finally, recall that we already established the third condition without loss of generality.

▷ **Claim 9** (\star). Every traversal has at most $k(k+2)$ vertices. Moreover, for any component C of $G - M$, there are at most $k^{k(k+2)}$ traversals of C . Furthermore, each traversal uses each edge at most k times.

Skeleton of the Solution. Let us define two directed graphs H_0 and H_1 as follows. Both have the same vertex set, M .

For distinct $u, v \in M$, the arc uv exists in H_0 if such an arc appears in C^* (i.e., C^* contains u immediately followed by v). On the other hand, the arc uv exists in H_1 if there is a non-trivial segment in \mathcal{S}^* starting in u and ending in v . We emphasize that neither H_0 nor H_1 has parallel arcs nor loops.

Let H be the directed graph with vertex set M and the arc set being the union of arc sets of H_0 and H_1 (again, with no parallel arcs). Note that H has one strongly connected component H' containing all vertices from $W \cap M$, and vertices that are not in H' are isolated in H . These vertices do not appear on C^* at all.

For $u, v \in M$, we say that a segment from u to v is *bad* (for a particular choice of H_0, H_1) if u or v are not in H' , or the arc uv does not appear in H . A traversal is *bad* if it contains a bad segment.

We say that a collection \mathcal{S} of segments is compatible with H_0 and H_1 if, for every arc uv of H_0 , \mathcal{S} contains at least one copy of the trivial segment uv , and for every arc uv of H_1 , \mathcal{S} contains at least one non-trivial segment starting in u and ending in v .

The Algorithm. We proceed to the description of the algorithm. We exhaustively guess the graphs H_0 and H_1 described above, considering all possible graphs on the vertex set M , so that the union of these two graphs has one strongly connected component and the remaining vertices isolated. This gives us $2^{O(k^2)}$ possible guesses. We are going to look for a set of segments that are compatible with H_0 and H_1 , and the structure of H_0, H_1 will be used to ensure that the solution is connected.

Suppose we are in a branch corresponding to one guess. We will build an instance of generalized N -fold ILP [29, 45]. In our case, we have a brick corresponding to each component C of $G - M$ and some further bricks that do not have any local constraints.

For each traversal of C that is not bad, we introduce one binary variable; recall that there are at most $k^{k(k+2)}$ such variables. The value of this variable indicates whether the vertices from $W \cap V(C)$ are covered by a particular traversal. Furthermore, we introduce a local constraint to ensure that exactly one traversal is chosen.

For each simple segment associated with C which is a u - v -path for some $u, v \in M$, we introduce a variable with possible values from 0 to n . The intended meaning of this variable is the number of times this segment was used as a connector from u to v by the solution. There are at most k^{k+2} such variables.

For each edge e with at least one endpoint in C we introduce a (local) constraint ensuring that this edge is used by the solution at most as many times as the capacity permits. Namely, we sum over all simple segments and over all traversals that use edge e (multiplying by how many times they use e) and require that this sum is at most $\kappa(e)$.

Now, let us move to variables not related to components. For each ordered pair u, v of distinct vertices, we introduce a variable with possible values from 0 to $\kappa((u, v))$. The intended meaning of this variable is the number of times the trivial segment is used as a connector from u to v by the solution. This gives us at most k^2 variables. We can put each of this variable to a brick for itself. As the edges within M are not used by any traversals or non-trivial segments, the box constraints ensure that the capacities are adhered for these arcs, i.e., there will be no local constraints for these bricks.

Next, we add global constraints that ensure that the selected set of segments is compatible with H_0 and H_1 . More precisely, for each arc uv of H_0 (H_1), we add a constraint that ensures that we choose the trivial segment uv (at least one non-trivial segment that starts in u and ends in v), respectively. We sum it over all selected traversals and connectors. These constraints are responsible for making the solution connected.

Finally, we need to make sure that for each $v \in M$, the number of segments that start in v is equal to the number of segments that end in v . Again, we can force it by summing the number of particular segments, including trivial, connectors, and ones belonging to traversals. These are also global constraints.

We minimize the total length (i.e., the sum of weights of edges) in all selected segments. Note that for some guesses we might get an infeasible instance of ILP – this means that, e.g., the segments that are required by H_0 and H_1 do not exist in G or all traversals of some component are bad. If all branches give infeasible instances, we return that there is no solution. However, every solution found (in any branch) is feasible and we can return the shortest one. From the description above it follows that if an optimum solution C^* exist, we will find it (or, more precisely, a solution of the same total length) in the branch corresponding to the actual H_0 and H_1 .

Solving N -fold ILP. Recall that generalized N -fold ILP can be solved (see [30, Cor. 97] and also [29, 45]) in $(rs\Delta)^{O(r^2s+rs^2)}d \log d \log D \log val_{\max}$ time, where

- r is the number of global constraints (at most $k(k-1) + k = k^2$ in our case),
- s is the number of local constraints for each brick ($O(k^2)$ in our case),
- Δ is the largest coefficient in the constraints (k in our case),
- d is the total number of variables (at most $k^2 + n \cdot k^{k(k+2)} + n \cdot k^{k+2}$ in our case)
- D is the maximum range of a variable (at most n in our case),
- and val_{\max} is the maximum value of the objective function (by Observation 7, in our case it is at most $n^2 \cdot U$, where U is the maximum weight of an arc in G).

Including the branching to guess H_0 and H_1 , we obtain the final running time bound of $k^{O(k^6)} \cdot n \cdot \log^2 n (\log n + \log U)$. This completes the proof. ◀

5 W[1]-hardness with Respect to Distance to Constant Treedepth

In this section we prove Theorem 5.

► **Theorem 5.** *DWRP is $W[1]$ -hard with respect to modulator to constant treedepth, even if all weights are 1.*

We provide a reduction from CAPACITATED DOMINATING SET (CDS), which is defined as follows. We are given an undirected graph G and a *capacity function* $c: V(G) \rightarrow \mathbb{N}$, where $1 \leq c(u) \leq \deg_G(u)$ for every $u \in V(G)$. A set of vertices $S \subseteq V(G)$ is a *capacitated dominating set* if there exists a *domination mapping* $f: (V(G) \setminus S) \rightarrow S$ such that for every $v \in (V(G) \setminus S)$ we have $\{f(v), v\} \in E(G)$ and each vertex $s \in S$ dominates at most $c(s)$ vertices, i.e., $|f^{-1}(s)| \leq c(s)$. Given a $k \in \mathbb{N}$, the problem is to find a capacitated dominating set of size k . CDS was shown to be $W[1]$ -hard with respect to treewidth in [27]. In order to prove Theorem 5, we actually need to observe a slight improvement of the construction proposed in [27]: we require that CDS is $W[1]$ -hard with respect to the distance to constant treedepth. This $W[1]$ -hardness is obtained by carefully considering the structure of the graph used in the reduction in [27]. Formally:


► **Theorem 10 (★).** *CDS is $W[1]$ -hard with respect to the distance to stars.*

We are now ready to proceed with the sketch of the proof of Theorem 5. A more detailed version can be found in the appendix.

Sketch of Proof of Theorem 5. Let $I = (G = (V, E), c, k)$ be an instance of CDS. We describe the construction of an instance $I' = (G', W, \omega, \kappa, b)$ of DWRP that is equivalent to I such that G' has a bounded modulator to constant treedepth.

The construction. First, we describe the construction of the used gadgets. Each gadget has dedicated *connecting vertices* marked with ‘in’ (e.g. u_{in}), ‘out’, or ‘io’. Every gadget of the final G' will be connected to other gadgets by edges incident only to its connecting vertices: ‘in’ (‘out’ resp.) vertices can only be incident to arcs that are coming ‘into’ (going ‘out’ from resp.) the gadget, and ‘io’ can be incident to both types of arcs. If this is true then we say that the graph *contains* the gadget. We use u^X to denote the vertex u that belongs to a gadget X . All edges have weight 1. Unless stated otherwise, every edge has capacity n ; this is always sufficient by Observation 7. Crucially, getting from one terminal to another requires the traversal of exactly two non-terminal vertices. As the final budget is set to exactly $3 \cdot (\text{number of terminals})$, it follows that every terminal is visited exactly once.


We begin with the *force- p -traversals-gadget*. It consists of three non-terminal vertices $u_{\text{in}}, u_{\text{out}}, w$ and terminal vertices v_1, \dots, v_p joined with edges $(u_{\text{in}}, v_i), (v_i, w)$ for every $i \in [p]$ and (w, u_{out}) of capacity p , see Figure 2.



■ **Figure 2** Force-5-traversals-gadget. Black and white vertices represent terminals and non-terminals, respectively.

▷ **Claim 11.** Let G' be a graph that contains a force- p -traversals-gadget X . Then every solution to DWRP in G' traverses from u_{in}^X to u_{out}^X exactly p times.

Next, we need the *cover- z -gadget*, which is illustrated in Figure 3. Let R_1 be the directed path that starts at s_{in} , goes to f , then to b , continues to x_{in} , traverses all the w s until y_{out} , goes to c and finishes in t_{out} . Also, let R_2 and P be the red and blue paths depicted in Figure 3. The idea here is that this gadget has two possible behaviors: either its terminals are covered by the combination of the R_2 and P paths, or all of its terminals except from z_{io} are covered by the R_1 path, and z is covered by some external path that visits this gadget. Formally:



■ **Figure 3** The cover- z -gadget. Black and white vertices represent terminals and non-terminals, respectively. The path R_2 is colored red, while the path P is colored blue.

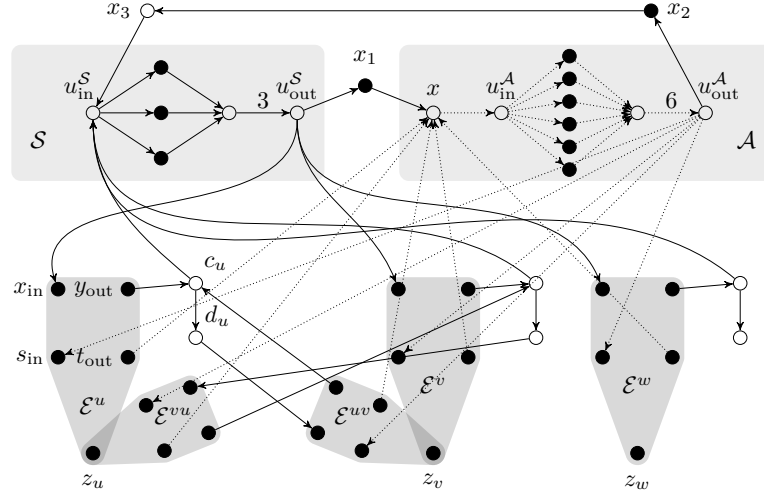
▷ **Claim 12.** Let G be a graph that contains the cover- z -gadget X . If S is a minimum solution to DWRP that traverses every terminal vertex of X exactly once, then S includes either the path R_1 , or the paths P and R_2 as continuous subpaths.

We now present the construction of $(G', W, \omega, \kappa, b)$, see Figure 4. Let \mathcal{S} be a force- $(k+1)$ -traversals-gadget that we call *selection-gadget* and let \mathcal{A} be a force- $(2|E| + |V| + 1)$ -traversals-gadget. We create an auxiliary non-terminal vertex x and add an arc (x, u_{in}^A) . Let x together with \mathcal{A} be called the *auxiliary-gadget*. We connect the selection- and auxiliary-gadgets by creating terminal vertices x_1, x_2 , non-terminal vertex x_3 , and arcs (u_{out}^A, x_2) , (x_2, x_3) , (x_3, u_{in}^S) , (u_{out}^S, x_1) , (x_1, x) . These two gadgets compose the *upper part* of G' .

Finally, we encode the vertices and edges of G into $G' = (V', E')$. For every vertex $u \in V$ we add vertices z_u, c_u , and d_u to V' , then we create a cover- z_u -gadget denoted \mathcal{E}^u . We add arcs $(u_{\text{out}}^S, x_{\text{in}}^{\mathcal{E}^u})$, $(y_{\text{out}}^{\mathcal{E}^u}, c_u)$, (c_u, d_u) , (c_u, u_{in}^S) , $(u_{\text{out}}^A, s_{\text{in}}^{\mathcal{E}^u})$, $(t_{\text{out}}^{\mathcal{E}^u}, x)$ to E' , and we set the capacity of (c_u, d_u) to $c(u)$. For every arc (u, v) such that $\{u, v\} \in E$, we create a cover- z_v -gadget denoted $\mathcal{E}^{u,v}$. Then we add arcs $(d_u, x_{\text{in}}^{\mathcal{E}^{u,v}})$, $(y_{\text{out}}^{\mathcal{E}^{u,v}}, c_v)$, $(u_{\text{out}}^A, s_{\text{in}}^{\mathcal{E}^{u,v}})$, $(t_{\text{out}}^{\mathcal{E}^{u,v}}, x)$. Note that in the above we created two gadgets, $\mathcal{E}^{u,v}$ and $\mathcal{E}^{v,u}$, for each $\{u, v\} \in E$. Finally, we set the budget $b = 132|E| + 69|V| + 3k + 12$. This concludes construction of the DWRP instance.

The instances are equivalent. On a high level, we show that any solution C of the instance $(G', W, \omega, \kappa, b)$ obeying the above budget has a specific behavior. First, the budget is set so that each terminal can be visited at most once. Let $S \subseteq V$ be a solution of I . The walk C is separated into two stages. In the first stage, for each $u \in S$ the walk C starts from u_{in}^S , goes through \mathcal{S} to u_{out}^S , from there to $x_{\text{in}}^{\mathcal{E}^u}$, using the P path (visiting z_u) of the gadget \mathcal{E}^u to $y_{\text{out}}^{\mathcal{E}^u}$ and to c_u . Then for each vertex v that is dominated by u , it goes from c_u to d_u , uses the P path of gadget $\mathcal{E}^{u,v}$ to visit z_v and return to c_u . Crucially, this can be done at most $c(u)$ times. Finally it returns to u_{in}^S .

Once all the vertices of S are dealt with in this way the first stage is done, and C proceeds to \mathcal{A} . For each cover-gadget \mathcal{E} (which are equal in number to the terminals of \mathcal{A} minus one),



■ **Figure 4** An example of a completely constructed DWRP instance for graph $G = (\{u, v, w\}, \{\{u, v\}\})$ and $k = 2$ (capacities omitted). The top half contains the upper part of G' . The bottom half corresponds to the vertices and edges of G . Gray boxes hide the inner parts of the corresponding cover-gadgets, see Figure 3. Dotted arrows signify arcs that belong to the auxiliary-gadget or connect cover-gadgets to the auxiliary-gadget. Black and white vertices represent terminals and non-terminals, respectively.

the walk C goes from x through \mathcal{A} to $u_{\text{out}}^{\mathcal{A}}$ and then through \mathcal{E} to cover the “remaining” vertices of \mathcal{E} (i.e., the vertices that were not visited during the first stage); this is done by employing either the R_1 or the R_2 path, according to whether \mathcal{E} was traversed during the first step. We then show that C exists and is of the correct budget if and only if I is a yes-instance of CDS.

The rest of the proof focuses on proving that G' has a modulator of size $7|M| + 7$ to treedepth 86, using that $G' \setminus M$ has treedepth 2 (it is a disjoint union of stars). ◀

6 Conclusions and Open Problems

We have shown that DWRP is FPT w.r.t. vertex integrity as well as w.r.t. feedback edge set number, while it is W[1]-hard w.r.t. distance to constant treedepth even in unweighted graphs. DTSP is W[1]-hard w.r.t. distance to pathwidth 3, also even in unweighted graphs and we conjecture that the above hardness w.r.t. distance to constant treedepth holds also for this problem. An immediate open problem stemming from our research is to close the gap by considering distances to even smaller constant values of the parameters. Namely, we would like to know whether the problems are FPT w.r.t. feedback vertex set (distance to treewidth 1), distance to disjoint paths (pathwidth 1), or distance to stars (treedepth 2).

We have also shown, that the problems are FPT parameterized by treewidth combined with the number of visits. Given that all the hardness reductions leverage vertices of high degrees, we are wondering whether a combination of treewidth with the maximum degree of the input graph might lead to an fpt-algorithm.

In this paper, we focused on structural parameters of the underlying undirected graph. Nevertheless, there are many width measure actually designated for directed graphs [34] and the tractability w.r.t. them is open. In fact, currently we do not even know whether DTSP is polynomial time solvable or NP-hard on directed graphs that become acyclic (DAGs) after

deleting a single vertex or even a single edge.

References

- 1 Noga Alon, Raphael Yuster, and Uri Zwick. Color-coding. *J. ACM*, 42(4):844–856, 1995. doi:10.1145/210332.210337.
- 2 Saeed Akhoondian Amiri, Klaus-Tycho Foerster, Riko Jacob, Mahmoud Parham, and Stefan Schmid. Waypoint routing in special networks. In Claudio Casetti, Fernando A. Kuipers, James P. G. Sterbenz, and Burkhard Stiller, editors, *2018 IFIP Networking Conference and Workshops, Networking 2018*, pages 496–504. IFIP, 2018. doi:10.23919/IFIPNetworking.2018.8696560.
- 3 Saeed Akhoondian Amiri, Klaus-Tycho Foerster, Riko Jacob, and Stefan Schmid. Charting the algorithmic complexity of waypoint routing. *Comput. Commun. Rev.*, 48(1):42–48, 2018. doi:10.1145/3211852.3211859.
- 4 Saeed Akhoondian Amiri, Klaus-Tycho Foerster, and Stefan Schmid. Walking through waypoints. *Algorithmica*, 82(7):1784–1812, 2020. doi:10.1007/s00453-020-00672-z.
- 5 Nima Anari and Shayan Oveis Gharan. Effective-resistance-reducing flows, spectrally thin trees, and asymmetric TSP. In Venkatesan Guruswami, editor, *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015*, pages 20–39. IEEE Computer Society, 2015. doi:10.1109/FOCS.2015.11.
- 6 David L. Applegate, Robert E. Bixby, Vašek Chvátal, and William J. Cook. *The Traveling Salesman Problem: A Computational Study*, volume 17 of *Princeton Series in Applied Mathematics*. Princeton University Press, Princeton, 2011. doi:10.1515/9781400841103.
- 7 Arash Asadpour, Michel X. Goemans, Aleksander Madry, Shayan Oveis Gharan, and Amin Saberi. An $O(\log n / \log \log n)$ -approximation algorithm for the asymmetric traveling salesman problem. *Oper. Res.*, 65(4):1043–1061, 2017. doi:10.1287/OPRE.2017.1603.
- 8 Jørgen Bang-Jensen and Gregory Z. Gutin. *Digraphs – Theory, Algorithms and Applications, Second Edition*. Springer Monographs in Mathematics. Springer, 2009.
- 9 Lukas Behrendt, Katrin Casel, Tobias Friedrich, J. A. Gregor Lagodzinski, Alexander Löser, and Marcus Wilhelm. From symmetry to asymmetry: Generalizing TSP approximations by parametrization. *J. Comput. Syst. Sci.*, 136:157–170, 2023. doi:10.1016/J.JCSS.2023.03.007.
- 10 Richard Bellman. Dynamic programming treatment of the travelling salesman problem. *Journal of the ACM*, 9(1):61–63, 1962. doi:10.1145/321105.321111.
- 11 Mark de Berg, Kevin Buchin, Bart M. P. Jansen, and Gerhard Woeginger. Fine-grained complexity analysis of two classic TSP variants. *ACM Transactions on Algorithms*, 17(1):5:1–5:29, 2021. doi:10.1145/3414845.
- 12 André Berger, László Kozma, Matthias Mnich, and Roland Vincze. Time- and space-optimal algorithm for the many-visits TSP. *ACM Trans. Algorithms*, 16(3):35:1–35:22, 2020. doi:10.1145/3382038.
- 13 René van Bevern and Daniel A. Skachkov. A quadratic-order problem kernel for the traveling salesman problem parameterized by the vertex cover number. *Operations Research Letters*, 52:107065, 2024. doi:10.1016/j.orl.2024.107065.
- 14 Václav Blažej, Pratibha Choudhary, Dušan Knop, Šimon Schierreich, Ondřej Suchý, and Tomáš Valla. On polynomial kernels for traveling salesperson problem and its generalizations. In Shiri Chechik, Gonzalo Navarro, Eva Rotenberg, and Grzegorz Herman, editors, *30th Annual European Symposium on Algorithms, ESA 2022*, volume 244 of *LIPICs*, pages 22:1–22:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.ESA.2022.22.
- 15 Václav Blažej, Andreas Emil Feldmann, Foivos Fioravantes, Pawel Rzażewski, and Ondřej Suchý. Parameterized complexity of directed traveling salesman problem. *CoRR*, abs/2506.22127, 2025. doi:10.48550/ARXIV.2506.22127.
- 16 Hans-Joachim Böckenhauer, Juraj Hromkovic, Joachim Kneis, and Joachim Kupke. The parameterized approximability of TSP with deadlines. *Theory Comput. Syst.*, 41(3):431–444, 2007. doi:10.1007/S00224-007-1347-X.

- 17 Hans L. Bodlaender, Marek Cygan, Stefan Kratsch, and Jesper Nederlof. Deterministic single exponential time algorithms for connectivity problems parameterized by treewidth. *Inf. Comput.*, 243:86–111, 2015. doi:10.1016/j.ic.2014.12.008.
- 18 Édouard Bonnet, Yoichi Iwata, Bart M. P. Jansen, and Łukasz Kowalik. Fine-grained complexity of k -OPT in bounded-degree graphs for solving TSP. In Michael A. Bender, Ola Svensson, and Grzegorz Herman, editors, *Proceedings of the 27th Annual European Symposium on Algorithms, ESA '19*, volume 144 of *Leibniz International Proceedings in Informatics*, pages 23:1–23:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.ESA.2019.23.
- 19 Édouard Bonnet, Michael Lampis, and Vangelis Th. Paschos. Time-approximation trade-offs for inapproximable problems. *J. Comput. Syst. Sci.*, 92:171–180, 2018. doi:10.1016/J.JCSS.2017.09.009.
- 20 Nicos Christofides. Worst-case analysis of a new heuristic for the travelling salesman problem. *Operations Research Forum*, 3:20, 2022. doi:10.1007/s43069-021-00101-z.
- 21 W. Cook. In *Pursuit of the Traveling Salesman: Mathematics at the Limits of Computation*. Princeton University Press, 2011. doi:10.1515/9781400839599.
- 22 Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms, 3rd Edition*. MIT Press, 2009. URL: <http://mitpress.mit.edu/books/introduction-algorithms>.
- 23 Marek Cygan, Fedor V. Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer International Publishing, Cham, 2015. doi:10.1007/978-3-319-21275-3.
- 24 Marek Cygan, Łukasz Kowalik, and Arkadiusz Socała. Improving TSP tours using dynamic programming over tree decompositions. *ACM Transactions on Algorithms*, 15(4):54:1–54:19, 2019. doi:10.1145/3341730.
- 25 Marek Cygan, Stefan Kratsch, and Jesper Nederlof. Fast hamiltonicity checking via bases of perfect matchings. *J. ACM*, 65(3):12:1–12:46, 2018. doi:10.1145/3148227.
- 26 Reinhard Diestel. *Graph Theory, 6th Edition*, volume 173 of *Graduate texts in mathematics*. Springer, 2025.
- 27 Michael Dom, Daniel Lokshtanov, Saket Saurabh, and Yngve Villanger. Capacitated domination and covering: A parameterized perspective. In Martin Grohe and Rolf Niedermeier, editors, *Parameterized and Exact Computation, Third International Workshop, IWPEC 2008*, volume 5018 of *Lecture Notes in Computer Science*, pages 78–90. Springer, 2008. doi:10.1007/978-3-540-79723-4_9.
- 28 Pål Grønås Drange, Markus S. Dregi, and Pim van 't Hof. On the computational complexity of vertex integrity and component order connectivity. *Algorithmica*, 76(4):1181–1202, 2016. doi:10.1007/S00453-016-0127-X.
- 29 Friedrich Eisenbrand, Christoph Hunkenschröder, and Kim-Manuel Klein. Faster algorithms for integer programs with block structure. In Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, editors, *45th International Colloquium on Automata, Languages, and Programming, ICALP '18*, volume 107 of *Leibniz International Proceedings in Informatics*, pages 49:1–49:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.ICALP.2018.49.
- 30 Friedrich Eisenbrand, Christoph Hunkenschröder, Kim-Manuel Klein, Martin Koutecký, Asaf Levin, and Shmuel Onn. An algorithmic theory of integer programming. *CoRR*, abs/1904.01361, 2019. arXiv:1904.01361.
- 31 Michael R Fellows and Sam Stueckle. The immersion order, forbidden subgraphs and the complexity of network integrity. *J. Combin. Math. Combin. Comput*, 6(1):23–32, 1989.
- 32 Fedor V. Fomin, Petr A. Golovach, Daniel Lokshtanov, and Saket Saurabh. Intractability of clique-width parameterizations. *SIAM Journal on Computing*, 39(5):1941–1956, 2010. doi:10.1137/080742270.

- 33 Fedor V. Fomin, Petr A. Golovach, Daniel Lokshtanov, Saket Saurabh, and Meirav Zehavi. Clique-width III: Hamiltonian cycle and the odd case of graph coloring. *ACM Trans. Algorithms*, 15(1), 2018. doi:10.1145/3280824.
- 34 Robert Ganian, Petr Hlinený, Joachim Kneis, Daniel Meister, Jan Obdržálek, Peter Rossmanith, and Somnath Sikdar. Are there any good digraph width measures? *J. Comb. Theory B*, 116:250–286, 2016. doi:10.1016/J.JCTB.2015.09.001.
- 35 M. R. Garey, David S. Johnson, and Robert Endre Tarjan. The planar hamiltonian circuit problem is NP-complete. *SIAM J. Comput.*, 5(4):704–714, 1976. doi:10.1137/0205049.
- 36 Shayan Oveis Gharan, Amin Saberi, and Mohit Singh. A randomized rounding approach to the traveling salesman problem. In Rafail Ostrovsky, editor, *Proceedings of the 52nd IEEE Annual Symposium on Foundations of Computer Science, FOCS '11*, pages 550–559. IEEE Computer Society, 2011. doi:10.1109/FOCS.2011.80.
- 37 Jiong Guo, Sepp Hartung, Rolf Niedermeier, and Ondřej Suchý. The parameterized complexity of local search for TSP, more refined. *Algorithmica*, 67(1):89–110, 2013. doi:10.1007/s00453-012-9685-8.
- 38 Gregory Gutin, Mark Jones, and Magnus Wahlström. The mixed chinese postman problem parameterized by pathwidth and treedepth. *SIAM Journal on Discrete Mathematics*, 30(4):2177–2205, 2016. doi:10.1137/15M1034337.
- 39 Gregory Gutin and Abraham P. Punnen, editors. *The Traveling Salesman Problem and Its Variations*. Springer, 2007. doi:10.1007/b101971.
- 40 Michael Held and Richard M. Karp. A dynamic programming approach to sequencing problems. *Journal of the Society for Industrial and Applied Mathematics*, 10(1):196–210, 1962. doi:10.1137/0110015.
- 41 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k -SAT. *Journal of Computer and System Sciences*, 62(2):367–375, 2001. doi:10.1006/jcss.2000.1727.
- 42 Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001. doi:10.1006/JCSS.2001.1774.
- 43 Marek Karpinski, Michael Lampis, and Richard Schmieđ. New inapproximability bounds for TSP. *J. Comput. Syst. Sci.*, 81(8):1665–1677, 2015. doi:10.1016/J.JCSS.2015.06.003.
- 44 Philip N. Klein and Dániel Marx. A subexponential parameterized algorithm for subset TSP on planar graphs. In Chandra Chekuri, editor, *Proceedings of the 2014 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '14*, pages 1812–1830. SIAM, 2014. doi:10.1137/1.9781611973402.131.
- 45 Martin Koutecký, Asaf Levin, and Shmuel Onn. A parameterized strongly polynomial algorithm for block structured integer programs. In Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, editors, *45th International Colloquium on Automata, Languages, and Programming, ICALP '18*, volume 107 of *Leibniz International Proceedings in Informatics*, pages 85:1–85:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.ICALP.2018.85.
- 46 Lukasz Kowalik, Shaohua Li, Wojciech Nadara, Marcin Smulewicz, and Magnus Wahlström. Many-visits TSP revisited. *J. Comput. Syst. Sci.*, 124:112–128, 2022. doi:10.1016/J.JCSS.2021.09.007.
- 47 Michael Lampis. Algorithmic meta-theorems for restrictions of treewidth. *Algorithmica*, 64(1):19–37, 2012. doi:10.1007/S00453-011-9554-X.
- 48 Giuseppe Lancia and Marcello Dalpasso. Algorithmic strategies for a fast exploration of the TSP 4-opt neighborhood. *J. Heuristics*, 30(3-4):109–144, 2024. URL: <https://doi.org/10.1007/s10732-023-09523-w>, doi:10.1007/S10732-023-09523-w.
- 49 Isja Mannens, Jesper Nederlof, Céline M. F. Swennenhuis, and Krisztina Szilágyi. On the parameterized complexity of the connected flow and many visits TSP problem. In Lukasz Kowalik, Michal Pilipczuk, and Pawel Rzażewski, editors, *Graph-Theoretic Concepts in*

- Computer Science - 47th International Workshop, WG 2021*, volume 12911 of *Lecture Notes in Computer Science*, pages 52–79. Springer, 2021. doi:10.1007/978-3-030-86838-3_5.
- 50 Dániel Marx. Searching the k -change neighborhood for TSP is $W[1]$ -hard. *Operations Research Letters*, 36(1):31–36, 2008. doi:10.1016/j.orl.2007.02.008.
- 51 Dániel Marx, Marcin Pilipczuk, and Michal Pilipczuk. A subexponential parameterized algorithm for directed subset traveling salesman problem on planar graphs. *SIAM J. Comput.*, 51(2):254–289, 2022. doi:10.1137/19M1304088.
- 52 Dániel Marx, Ario Salmasi, and Anastasios Sidiropoulos. Constant-factor approximations for asymmetric TSP on nearly-embeddable graphs. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2016*, volume 60 of *LIPICs*, pages 16:1–16:54. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPICs.APPROX-RANDOM.2016.16.
- 53 Tobias Mömke and Ola Svensson. Removing and adding edges for the traveling salesman problem. *Journal of the ACM*, 63(1):2:1–2:28, 2016. doi:10.1145/2739008.
- 54 Marcin Mucha. 13/9-approximation for graphic TSP. *Theory of Computing Systems*, 55:640–657, 2014. doi:10.1007/s00224-012-9439-7.
- 55 Moni Naor, Leonard J. Schulman, and Aravind Srinivasan. Splitters and near-optimal derandomization. In *36th Annual Symposium on Foundations of Computer Science, Milwaukee, Wisconsin, USA, 23-25 October 1995*, pages 182–191. IEEE Computer Society, 1995. doi:10.1109/SFCS.1995.492475.
- 56 Šimon Schierreich and Ondřej Suchý. Waypoint routing on bounded treewidth graphs. *Information Processing Letters*, 173:106165, 2022. doi:10.1016/j.ipl.2021.106165.
- 57 András Sebő and Jens Vygen. Shorter tours by nicer ears: 7/5-approximation for the graph-TSP, 3/2 for the path version, and 4/3 for two-edge-connected subgraphs. *Combinatorica*, 34(5):597–629, 2014. doi:10.1007/s00493-014-2960-3.
- 58 D.B. Shmoys, J.K. Lenstra, A.H.G. Rinnooy Kan, and E.L. Lawler. *The Traveling Salesman Problem*, volume 12 of *Wiley Series in Discrete Mathematics & Optimization*. John Wiley & Sons, Incorporated, 1985. URL: <https://www.wiley.com/en-ie/The+Traveling+Salesman+Problem%3A+A+Guided+Tour+of+Combinatorial+Optimization-p-9780471904137>.
- 59 Ola Svensson. Approximating ATSP by relaxing connectivity. In Venkatesan Guruswami, editor, *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015*, pages 1–19. IEEE Computer Society, 2015. doi:10.1109/FOCS.2015.10.
- 60 Ola Svensson, Jakub Tarnawski, and László A. Végh. A constant-factor approximation algorithm for the asymmetric traveling salesman problem. *J. ACM*, 67(6):37:1–37:53, 2020. doi:10.1145/3424306.
- 61 Vera Traub and Jens Vygen. An improved approximation algorithm for the asymmetric traveling salesman problem. *SIAM J. Comput.*, 51(1):139–173, 2022. doi:10.1137/20M1339313.

A Omitted Material from Preliminaries

► **Definition 13** (Tree decomposition). A tree decomposition of a graph G is a triple (T, β, r) , where T is a tree rooted at node r and $\beta: V(T) \rightarrow 2^{V(G)}$ is a mapping that assigns to each node x of the tree its bag $\beta(x)$ and satisfies:

1. $\bigcup_{x \in V(T)} \beta(x) = V(G)$;
2. For every $\{u, v\} \in E(G)$ there exists a node $x \in V(T)$, such that $u, v \in \beta(x)$;
3. For every $u \in V(G)$ the nodes $\{x \in V(T) \mid u \in \beta(x)\}$ form a connected subtree of T .

To distinguish between the vertices of a tree decomposition and the vertices of the underlying graph, we use the term *node* for the vertices of a given tree decomposition.

The *width* of a tree decomposition (T, β, r) is $\max_{x \in V(T)} |\beta(x)| - 1$. The *treewidth* of a graph G (denoted $\text{tw}(G)$) is the minimum width of a tree decomposition of G over all such decompositions.

For a tree decomposition (T, β, r) and a node $x \in V(T)$, we denote by $\gamma(x)$ the union of vertices in $\beta(x)$ and in $\beta(y)$ for all descendants y of x in T . By E_x , we denote the set of edges introduced in the subtree of T rooted in x . Altogether, we denote by G_x the graph $(\gamma(x), E_x)$.

► **Definition 14** (Nice tree decomposition [23, p. 168]). A tree decomposition of a graph G is nice if $\beta(r) = \emptyset$, and each node $x \in V(T)$ is of one of the following five types:

- Leaf node— x has no children and $\beta(x) = \emptyset$;
- Introduce vertex node— x has exactly one child y and $\beta(x) = \beta(y) \cup \{u\}$ for some $u \in V(G) \setminus \beta(y)$;
- Introduce edge node— x has exactly one child y , $\beta(x) = \beta(y)$, and an edge $\{u, v\} \in E(G)$ for $u, v \in \beta(x)$ is introduced (i.e., $E_x = E_y \cup \{\{u, v\}\}$);
- Forget node— x has exactly one child y and $\beta(x) = \beta(y) \setminus \{u\}$ for some $u \in \beta(y)$;
- Join node— x has exactly two children y, z and $\beta(x) = \beta(y) = \beta(z)$, and each edge $e \in E(G)$ is introduced exactly once.

We note that any given tree decomposition of width t can be transformed into a nice one of the same width in $O(t^2 \cdot n)$ time [23, Lemma 7.4, see also the discussion on page 168].

Proof of Observation 7. Consider the decomposition of the optimal solution S to cycles as guaranteed by Lemma 6, denote it $\mathcal{C} = \{C_1, \dots, C_q\}$. Consider an undirected bipartite graph B with $V(B) = \mathcal{C} \cup V(G)$ and an edge connecting $u \in V(G)$ with $C \in \mathcal{C}$ if and only if $u \in V(C)$. Obviously $\deg_B v$ equals the number of cycles v appears in, which in turn equals the number of times v is visited by S . Let T be a BFS tree for the BFS started from the vertex v in graph B . It might not contain all vertices of $V(G)$, it only contains vertices visited by S , and in particular it contains all vertices in W . Obviously $\deg_T v = \deg_B v$.

If v is a leaf of T , then we are done. Otherwise, let T' be a tree obtained from T by iteratively removing leaves that are not vertices of W . Let \mathcal{C}' be the set of cycles that appear in T' and H be a multigraph obtained as the edge disjoint union of cycles in \mathcal{C}' . As H is an edge disjoint union of cycles, it is balanced. We claim that it is also connected. In particular, there is a path from any vertex $u \in V(H)$ to v in H . Indeed, u appears on some cycle $C \in \mathcal{C}'$ and there is a path $P_{T'}$ from C to v in T' . By replacing each cycle on $P_{T'}$ by a part of the cycle between the appropriate vertices, we obtain a path from u to v in H . Thus there is an Eulerian trail S' in H . Also, since T' contains all vertices in W , H contains all vertices in W , and S' visits all vertices in W . Moreover it traverses each edge at most as many times as S . Thus it is a solution. Therefore, as S is an optimal solution, we have $\mathcal{C}' = \mathcal{C}$.

This, in particular, implies that $v \in V(T')$ and $\deg_{T'} v = \deg_T v$. However, as each leaf of T' is a vertex of $W \setminus \{v\}$, T' has at most $|W \setminus \{v\}|$ leaves, which implies $\deg_{T'} v \leq |W \setminus \{v\}|$. \blacktriangleleft

B FPT Algorithm with Respect to Solution Size

► **Theorem 1** (\star). *In $e^k 4^k k^{O(\log k)} \cdot m \log m$ time we can decide whether the given m -edge instance of DWRP has a solution consisting of at most k edge occurrences.*

Our proof is based on the color coding technique introduced by Alon et al. [1], hence we need the following proposition.

► **Proposition 15** (Naor et al. [55]). *Let m, k be any positive integers. There exists a family \mathcal{F} of colorings $f: [m] \rightarrow [k]$ of size $e^k k^{O(\log k)} \log m$ constructible in $e^k k^{O(\log k)} m \log m$ time such that for any $F \subseteq [m]$ with $|F| \leq k$ there exists $f \in \mathcal{F}$ such that $f|_F$ is injective.*

Proof of Theorem 1. Since a solution which consist of at most k edge occurrences can use each edge at most k times, we assume that each capacity is at most k . Now we turn the input graph into a multigraph by replacing each edge e with $\kappa(e)$ edges, each with the same tail and head as e , weight $\omega(e)$, and capacity 1. It is easy to see that there is a solution of weight b in this new multigraph if and only if there is a solution of weight b in the original graph. For simplicity, we denote this new multigraph $G = (V, E)$, the new weight function ω and κ is the capacity function assigning 1 to every edge.

Since there must be an arc entering each of the waypoints, if $|W| > k$, then there is no solution consisting of at most k edge occurrences and we can answer NO. Hence, let us assume that $|W| \leq k$ for the rest of the proof.

Now we construct the family \mathcal{F} of colorings as guaranteed by Proposition 15 to color E with k colors. For each coloring $f \in \mathcal{F}$ we search for a *colorful* solution to our problem, that is, a solution that uses at most one edge of each color. Obviously, any colorful solution is a solution for the original problem. In particular, it cannot use any edge more than once, as this would mean that it used more than one edge of the corresponding color. Thus, it obeys the capacities. Moreover, by the properties of f , any solution consisting of at most k edges becomes colorful under at least one coloring from \mathcal{F} . Hence, the minimum weight colorful solution over all colorings from \mathcal{F} is also a minimum weight solution using at most k edges, and it is enough to compare the weight of this solution with the budget b .

To find a minimum weight colorful solution under a fixed coloring $f \in \mathcal{F}$ we use dynamic programming. First fix a waypoint $w_0 \in W$ and let $\widehat{W} = W \setminus \{w_0\}$. Now for every $W' \subseteq \widehat{W}$, every $C \subseteq [k]$ and every $v \in V$ we create a table cell $T[W', C, v]$ with the following semantics. Cell $T[W', C, v]$ stores the minimum weight of a walk that starts in w_0 , visits each waypoint in W' , uses exactly one edge of each color in C , and ends in v . If there is no such walk, then it stores ∞ .

We fill the table in order of increasing $|C|$. For $C = \emptyset$ we let $T[\emptyset, \emptyset, w_0] = 0$ and $T[W', \emptyset, v] = \infty$ for any $W' \neq \emptyset$ or $v \neq w_0$. For $|C| \geq 1$ we let

$$T[W', C, v] = \min \{T[W' \setminus \{v\}, C \setminus \{f(u, v)\}, u] + \omega(u, v) \mid (u, v) \in E, f(u, v) \in C\}.$$

Once the table was filled for all C , the minimum weight of a colorful solution under f can be obtained as $\min \{T[\widehat{W}, C, w_0] \mid C \subseteq [k]\}$.

As for the running time, for a fixed coloring f and a fixed vertex $v \in V$, there are $2^{|\widehat{W}|} \cdot 2^k$ table cells of the form $T[W', C, v]$ and each such cell can be computed in $O(\deg^{in} v)$ time

Summing over all vertices $v \in V$ we obtain that the tables can be filled in $2^{|\widehat{W}|} \cdot 2^k \cdot m'$ time, where $m' = |E|$. As $|\widehat{W}| \leq k-1$, this gives $O(4^k \cdot m')$ time to fill the cells. The answer can be found in the same running time. Summing over all $f \in \mathcal{F}$, this gives $e^k 4^k k^{O(\log k)} \cdot m' \log m'$ time in total, including the time to construct f . Since the capacity of each edge in the original graph was at most k , the multigraph has at most k times more edges than the original graph, i.e., $m' \leq k \cdot m$. Therefore we get the claimed $e^k 4^k k^{O(\log k)} \cdot m \log m$ running time.

Towards correctness, we first prove that the table is filled according to the desired semantics. We proceed by induction on $|C|$. If the walk is not allowed to use edge of any color, then it has to be trivial, and, thus, must end in w_0 and cannot visit any waypoint in \widehat{W} . As the trivial walk is of weight 0, for $C = \emptyset$ the table is filled correctly.

Let us now assume that $|C| \geq 1$ and the table is filled correctly for all smaller sets of colors. On the one hand, for every $(u, v) \in E$ with $f(u, v) \in C$, by induction hypothesis, if $T[W' \setminus \{v\}, C \setminus \{f(u, v)\}, u] \neq \infty$, then there is a walk P of weight $T[W' \setminus \{v\}, C \setminus \{f(u, v)\}, u]$ that starts in w_0 , visits each waypoint in $W' \setminus \{v\}$, uses exactly one edge of each color in $C \setminus \{f(u, v)\}$, and ends in u . Attaching the edge (u, v) to the end of walk P , we obtain a walk of weight $T[W' \setminus \{v\}, C \setminus \{f(u, v)\}, u] + \omega(u, v)$ that starts in w_0 , visits each waypoint in W' , uses exactly one edge of each color in C , and ends in v . Therefore, the value stored in $T[W', C, v]$ is at least the minimum weight of such a walk. On the other hand, if P is a walk of weight h that starts in w_0 , visits each waypoint in W' , uses exactly one edge of each color in C , and ends in v , then let (u, v) be the last edge on that walk (there is at least one edge since $|C| \geq 1$). Let P' be obtained from P by removing the last edge. Then P' is a walk of weight $h - \omega(u, v)$ that starts in w_0 , visits each waypoint in $W' \setminus \{v\}$, uses exactly one edge of each color in $C \setminus \{f(u, v)\}$, and ends in u . Thus $T[W' \setminus \{v\}, C \setminus \{f(u, v)\}, u] \leq h - \omega(u, v)$ and

$$T[W', C, v] \leq T[W' \setminus \{v\}, C \setminus \{f(u, v)\}, u] + \omega(u, v) \leq h - \omega(u, v) + \omega(u, v) = h.$$

I.e., as P was arbitrary, $T[W', C, v]$ is at most the weight of any such walk. Together with the previous one, it shows that the value $T[W', C, v]$ is computed correctly.

Any walk that starts and ends in w_0 , visits each waypoint in \widehat{W} and uses exactly one edge of each color in C for some $C \subseteq [k]$ gives a colorful solution to the problem, as it visits each waypoint in $\widehat{W} \cup \{w_0\} = W$. Thus any cell $T[\widehat{W}, C, w_0]$ corresponds to a colorful solution. Conversely, any colorful solution must visit w_0 . Therefore, it can be interpreted as a walk that starts and ends in w_0 , visits each waypoint in \widehat{W} and uses exactly one edge of each color in C for some $C \subseteq [k]$. The minimum weight of such a walk is stored in $T[\widehat{W}, C, w_0]$. Thus, our answer is correct. ◀

► **Corollary 16.** *DWRP is FPT with respect to budget b .*

Proof. As the weight of each edge is at least one, a solution of weight b cannot use more than b edges. ◀

► **Corollary 17.** *In $2^{O(\widehat{k}^2)} \cdot m \log m$ time we can decide whether the given instance of DWRP has a solution using at most \widehat{k} distinct edges (ignoring multiplicities).*

Proof. Consider a solution using at most \widehat{k} distinct edges. By Observation 7, each of the edges is used at most $|W| \leq \widehat{k}$ times. This implies that the solution uses at most \widehat{k}^2 edge occurrences in total and the result follows from Theorem 1. ◀

C Omitted Material from Section 3

Proof of Claim 8. First, suppose that there is a visit that starts at, say, u , ends in v , and traverses some edge more than once (i.e., is not a pass). In such a case one can easily remove some edges from this visit, trimming it to a pass. Note that the obtained walk visits the same vertices as C^* and is shorter, contradicting the minimality of C^* . Consequently, every non-pass visit starts and ends in the same vertex, either u or v .

Second, suppose that there is at least one visit that is a pass and at least one non-pass visit. Again, we can remove all edges corresponding to the non-pass visit, obtaining a better solution – it is still a closed walk, as the removed visit starts and ends in the same vertex. This justifies the claim. \triangleleft

Correctness of the Algorithm from Theorem 2. To prove correctness, we need to argue that $(G, W, \omega, \kappa, b)$ is a yes-instance of DWRP if and only if there exists a branch for which we created a yes-instance $(G', W', \omega', \kappa', b')$ of DWRP.

Suppose first that $(G, W, \omega, \kappa, b)$ is a yes-instance and let C^* be an optimum solution of cost at most b . As argued before, we may assume that the type of each path $P \in \mathcal{P}$ with respect to C^* is as listed above. Consider the branch in which all types are guessed correctly, according to C^* and the corresponding instance $(G', W', \omega', \kappa', b')$. Define a closed walk C' in G' as follows. Clearly, C^* can be decomposed into visits in paths from \mathcal{P} . Each visit that is not a pass is removed from C^* . Each visit in a path P that is a pass, say from u to v , is replaced by two consecutive edges $(u, x_{u \curvearrowright v}^P), (x_{u \curvearrowright v}^P, v)$ in G' . Call the obtained closed walk (in G') C' . It is straightforward to verify that C' visits all vertices from W' , does not exceed edge capacities, and its total cost is at most b' . Thus, $(G', \omega', \kappa', b')$ is a yes-instance of DWRP.

Now suppose that there is a branch in which the corresponding instance $(G', \omega', \kappa', b')$ is a yes-instance. Let C' be an optimum solution for this instance. We modify it as follows. For each subpath of the form $u, x_{u \curvearrowright v}^P, v$ for some $P \in \mathcal{P}$, we replace it by traversing the path P from u to v . Note that as $x_{u \curvearrowright v}^P \in W'$, all edges of types $u \curvearrowright v, u \curvearrowleft v, u \rightleftharpoons v$ are correctly passed by the constructed walk.

Now consider a path P from u to v , of type $u \leftrightarrow v$ (the cases $u \nleftrightarrow v$ and $u \circ v$ are treated analogously). Recall that $u \in W'$, so it appears somewhere in C' . Immediately after an arbitrary occurrence of u , we include a precomputed optimal visit in P that starts and ends at u , and visits all vertices of W that appear as internal vertices of P . Note that this increases the cost of the solution by $\text{cost}(P, u \leftrightarrow v)$. Let us denote the walk obtained by all these modifications by C^* . Again, it is straightforward to verify that C^* is a feasible solution to the instance $(G, W, \omega, \kappa, b)$ of DWRP.

D Omitted Proofs from Section 4

Proof of Claim 9. Note that each traversal can be decomposed into paths with both endvertices in M and internal vertices in C and cycles with at most one vertex in M and the remaining ones in C .

Note that each such path has at most $k + 2$ vertices and each such cycle has at most $k + 1$ vertices. Moreover, by the second and third condition in the definition of a traversal, the total number of paths and cycles does not exceed $|W \cap V(C)| \leq |V(C)| \leq k$. Consequently, the total number of vertices of a traversal is at most $k(k + 2)$ and each edge is used at most k times.

The bound on the number of traversals follows from the fact that C and M both have at most k vertices. \triangleleft

E DWRP is XP with Respect to Treewidth

In this section we present our algorithm parameterized by treewidth. As the algorithm uses standard techniques, it is rather similar to the algorithm of Mannens et al. [49]. The following theorem is a more detailed version of Theorem 4.

► **Theorem 18.** *There is an algorithm that given an instance $(G = (V, E), W, \omega, \kappa, b)$ of DIRECTED WAYPOINT ROUTING PROBLEM, a nice tree decomposition (T, β) of G of width t and an integer ν , computes in $O(((t+1)(\nu+1)^2)^{2(t+1)} \cdot t^2 \cdot n)$ (FPT wrt $t + \nu$) time the length of a shortest closed walk C that traverses each vertex in W at least once and each vertex in V at most ν times and such that for each edge e the number of times C traverses e is at most $\kappa(e)$ if there is any walk satisfying these conditions.*

► **Corollary 19.** *DIRECTED WAYPOINT ROUTING PROBLEM is in XP with respect to treewidth and FPT with respect to treewidth combined with the number of waypoints.*

Proof of Corollary 19. Follows directly from the theorem and Observation 7. \blacktriangleleft

Proof of Theorem 18. We use a dynamic programming over the given decomposition (T, β) . Let $r \in V(T)$ be the root of the decomposition. Let $w_0 \in W$ be an arbitrary waypoint. Remove the node where w_0 is forgotten from the decomposition and add w_0 to all the ancestors of that bag. With a suitable choice of the root while turning the tree decomposition into a nice one, this could be done without increasing the width of the decomposition.

We create a table A indexed by nodes $x \in V(T)$, partitions \mathcal{P} of $\beta(x)$, and pairs of functions $in, out: \beta(x) \rightarrow \{0, \dots, \nu\}$.

We call a sequence $\mathcal{C} = C_1, \dots, C_q$ of walks a *partial solution* at x if

- i) each C_i is a non-trivial walk in G_x that starts and ends in a vertex of $\beta(x)$,
- ii) each vertex in $W \cap (\gamma(x) \setminus \beta(x))$ is traversed by at least one of the walks,
- iii) no vertex is traversed more than ν times (in particular, no vertex has in-degree or out-degree more than ν with respect to the walks), summing over all the walks,
- iv) and each edge e is traversed at most $\kappa(e)$ times, summing over all the walks.

The *length* of a partial solution is the sum of the length of its walks.

A partial solution \mathcal{C} is *compatible with* (\mathcal{P}, in, out) at x , if for each $v \in \beta(x)$ vertex v is entered $in(v)$ times and exited $out(v)$ times summing over all walks in \mathcal{C} and moreover the following holds. Let H be the subgraph of G_x formed by the edges used by the walks in \mathcal{C} and H' be the underlying undirected graph of H . Then for each non-trivial connected component K of H' we have that $K \cap \beta(x) \neq \emptyset$ and $K \cap \beta(x) \in \mathcal{P}$ and for each $v \in \beta(x)$ isolated in H' we have $\{v\} \in \mathcal{P}$.

For each $x \in V(T)$, partition \mathcal{P} of $\beta(x)$, and a pair of functions $in, out: \beta(x) \rightarrow \{0, \dots, \nu\}$ we simply store in $A[x, \mathcal{P}, in, out]$ the length of a shortest partial solution compatible with (\mathcal{P}, in, out) at x , if there is any, or $+\infty$ otherwise.

We compute the values of $A[x, \mathcal{P}, in, out]$ in a bottom-up manner distinguishing the type of the node x .

Leaf Node If x is a leaf node, then, as $\beta(x) = \emptyset$, the only possible partition is \emptyset and the only possible function $\beta(x) \rightarrow \{0, \dots, \nu\}$ is also \emptyset . We let $A[x, \emptyset, \emptyset, \emptyset] = 0$.

Introduce Vertex Node If x is a node introducing vertex v and y is the only child of x , then v is isolated in G_x . Hence, there is no way to enter it or leave it. Thus, if $in(v) \neq 0$ or $out(v) \neq 0$, or $\{v\} \notin \mathcal{P}$, then we let $A[x, \mathcal{P}, in, out] = +\infty$. Otherwise, we let $\mathcal{P}' = \mathcal{P} \setminus \{\{v\}\}$, and in' and out' be in and out restricted to $\beta(y) = \beta(x) \setminus \{v\}$, respectively. We let $A[x, \mathcal{P}, in, out] = A[y, \mathcal{P}', in', out']$.

Introduce Edge Node If x is a node introducing an edge $\{u, v\}$ and y is the only child of x , then we have to guess how many times was the edge (u, v) and how many times the edge (v, u) is used by the solution. Let $P_u \in \mathcal{P}$ be such that $u \in P_u$. If $v \notin P_u$, then we let $A[x, \mathcal{P}, in, out] = A[y, \mathcal{P}, in, out]$.

Otherwise, let $a_{uv} = 0$ if the edge (u, v) is not in G and let $a_{uv} = \min\{\kappa((u, v)), out(u), in(v)\}$ otherwise. Similarly, let $a_{vu} = 0$ if the edge (v, u) is not part of G and let $a_{vu} = \min\{\kappa((v, u)), out(v), in(u)\}$ otherwise. Let

$$A[x, \mathcal{P}, in, out] = \min\left\{ \min_{\substack{0 \leq b_{uv} \leq a_{uv} \\ 0 \leq b_{vu} \leq a_{vu}}} (A[y, \mathcal{P}, in', out'] + b_{uv} \cdot \omega((u, v)) + b_{vu} \cdot \omega((v, u))), \right. \\ \left. \min_{\substack{0 \leq b_{uv} \leq a_{uv} \\ 0 \leq b_{vu} \leq a_{vu} \\ b_{uv} + b_{vu} \geq 1 \\ \mathcal{P}'}} (A[y, \mathcal{P}', in', out'] + b_{uv} \cdot \omega((u, v)) + b_{vu} \cdot \omega((v, u))) \right\},$$

where

$$in'(h) = \begin{cases} in(h) & h \notin \{u, v\} \\ in(v) - b_{uv} & h = v \\ in(u) - b_{vu} & h = u \end{cases},$$

$$out'(h) = \begin{cases} out(h) & h \notin \{u, v\} \\ out(u) - b_{uv} & h = u \\ out(v) - b_{vu} & h = v \end{cases},$$

and \mathcal{P}' contains sets P'_u and P'_v such that $u \in P'_u$, $v \in P'_v$ and $\mathcal{P} = (\mathcal{P}' \setminus \{P'_u, P'_v\}) \cup \{P_u\}$ (possibly $P'_u = P'_v = P_u$).

Forget Node If x is a node forgetting vertex v and y is the only child of x , then we distinguish cases based on whether v is a waypoint or not. In either case, let $\mathcal{P} = \{P_1, \dots, P_s\}$. To assure that the final solution is connected, if v is traversed, then v must share a connected component with a vertex in $\beta(x)$.

If v is a waypoint, then we need to make sure that it is traversed by the solution. We let

$$A[x, \mathcal{P}, in, out] = \min_{\substack{1 \leq i \leq \nu \\ 1 \leq j \leq s}} A[y, \mathcal{P}_j, in_i, out_i],$$

where $\mathcal{P}_j = (\mathcal{P} \setminus \{P_j\}) \cup \{P_j \cup \{v\}\}$ (i.e., we add v to the set P_j of \mathcal{P}),

$$in_i(u) = \begin{cases} in(u) & u \in \beta(x) \\ i & u = v \end{cases},$$

and

$$out_i(u) = \begin{cases} out(u) & u \in \beta(x) \\ i & u = v \end{cases}.$$

If v is not a waypoint, then it might not be part of the solution. We let

$$A[x, \mathcal{P}, in, out] = \min\{A[y, \mathcal{P} \cup \{\{v\}\}, in_0, out_0], \min_{\substack{1 \leq i \leq \nu \\ 1 \leq j \leq s}} A[y, \mathcal{P}_j, in_i, out_i]\},$$

where the meaning of \mathcal{P}_j , in_i , and out_i is the same as in the previous case.

Join Node Finally, let x be a join node with two children y and z . We simply let

$$A[x, \mathcal{P}, in, out] = \min_{\substack{\mathcal{P}_y, \mathcal{P}_z \\ in_y, in_z \\ out_y, out_z}} A[y, \mathcal{P}_y, in_y, out_y] + A[z, \mathcal{P}_z, in_z, out_z],$$

where the minimum is taken over all pairs of $(\mathcal{P}_y, in_y, out_y)$ and $(\mathcal{P}_z, in_z, out_z)$ such that \mathcal{P} is the finest common coarsening of \mathcal{P}_y and \mathcal{P}_z and for every $v \in \beta(x) = \beta(y) = \beta(z)$ we have $in(v) = in_y(v) + in_z(v)$ and $out(v) = out_y(v) + out_z(v)$. Note that we can simply iterate over the tables of the children, gradually updating the table of x .

Once the computation is finished, the desired length of a closed walk satisfying the conditions is obtained as $\min_{1 \leq i \leq \nu} A[r, \{\{w_0\}\}, w_0 \mapsto i, w_0 \mapsto i]$.

Time Complexity Since each bag is of size at most $t + 1$, there are at most $(t + 1)^{(t+1)}$ partitions of it.² Also there are at most $(\nu + 1)^{(t+1)}$ functions $\beta(x) \rightarrow \{0, \dots, \nu\}$. Hence, for each x the table $A[x, \bullet, \bullet, \bullet]$ has at most $((t + 1)(\nu + 1)^2)^{(t+1)}$ entries.

The computation in leaf nodes takes a constant time. Similarly, the computation of each entry in an introduce vertex node takes a constant time. In introduce edge node, for each entry we take a minimum over at most $(\nu + 1)^2 + (\nu + 1)^2 \cdot 2^{(t-1)}$ entries of the child. In forget node, for each entry we take a minimum over at most $1 + (\nu)t$ entries of the child. Finally, in a join node, we make a $O(t)$ time computation for each pair of entries of the children. Therefore, the computation in each node can be done in $O(((t + 1)(\nu + 1)^2)^{2(t+1)} \cdot t)$ time. As we can assume that there are $O(tn)$ nodes in the given decomposition [23, Lemma 7.4, see also p. 168], the algorithm runs in $O(((t + 1)(\nu + 1)^2)^{2(t+1)} \cdot t^2 \cdot n)$ time.

Correctness We first prove that the number stored in the table satisfies the desired semantics, that is, $A[x, \mathcal{P}, in, out]$ is the length of a shortest partial solution compatible with (\mathcal{P}, in, out) at x , if there is any, or $+\infty$ otherwise.

We start by proving that if $A[x, \mathcal{P}, in, out]$ is finite, then there is a compatible partial solution of length $A[x, \mathcal{P}, in, out]$. We proceed by induction on the height of the subtree of the tree decomposition rooted at x . For leaf nodes, the empty sequence (of walks) forms such a solution of length 0, proving the claim.

Next we assume that the claim already holds for the children of node x by induction hypothesis. For introduce vertex node x with child y , it is easy to show that the partial solution compatible with $(\mathcal{P}', in', out')$ at y is also compatible with (\mathcal{P}, in, out) at x .

For introduce edge node x with child y and the edge $\{u, v\}$ being introduced, take the solution \mathcal{C}' compatible with $(\mathcal{P}', in', out')$ at y , which exists by the induction hypothesis, and add to it b_{uv} walks from u to v and b_{vu} walks from v to u , each consisting of a single arc, to produce a sequence \mathcal{C} . It is straightforward to check that \mathcal{C} is a partial solution compatible with (\mathcal{P}, in, out) at x .

² We could reduce the number of partitions considered to $2^{O(t)}$ by the now standard technique of representative sets [17], but we refrain from that here to simplify the presentation.

For forget node, it is easy, if $A[x, \mathcal{P}, in, out] = A[y, \mathcal{P} \cup \{\{v\}\}, in_0, out_0]$ —simply take the same partial solution. If $A[x, \mathcal{P}, in, out] = A[y, \mathcal{P}_j, in_i, out_i]$ for some $i \in \{1, \dots, \nu\}$ and $j \in \{1, \dots, s\}$, then let C_1, \dots, C_q be the partial solution for $A[y, \mathcal{P}_j, in_i, out_i]$, which exists by the induction hypothesis. Furthermore, assume that walks C_1, \dots, C_a are closed walks that start and end in v , walks C_{a+1}, \dots, C_b start in v , but end in a different vertex of $\beta(x)$ and walks C_{b+1}, \dots, C_c start in a different vertex of $\beta(x)$ and end in v (possibly, $a = 0, b = a$, or $c = b$) and no other walks start or end in v . Note that $b - a = c - b$. For $d \in \{1, \dots, b - a\}$ we connect C_{b+d} with C_{a+d} to make a walk that neither starts nor ends in v . If $a = 0$, then we are done, as no walk starts or ends in v and we can use the solution for x . Otherwise, we connect all walks C_1, \dots, C_a into a single closed walk C' . Note that C' is now the only walk that starts or ends in v . As $P_j \neq \emptyset$ and $(P_j \cup \{v\}) \in \mathcal{P}_j$, either C' contains a vertex $u \in \beta(x)$ different from v , or it shares a vertex u with another walk C'' that starts and ends in a vertex of $\beta(x)$ different from v . As C' is a closed walk, we can “rotate” it, to start and end in u . In the former case we are done, the walk now ends in starts and ends in a vertex of $\beta(x)$ different from v . In the latter case, we break C'' at u and insert C' into it. Again, after that, no walk starts or ends in v and we can use this solution for x . As we did not change how many times edges are traversed, the other conditions are easy to verify.

Finally, for the join node, it is enough to take the union of the partial solutions for the two children. Since each edge is introduced exactly once, no edge can appear both in G_y and G_z and the capacities are adhered. Also each vertex of $\gamma(x) \setminus \beta(x)$ is traversed by at most one of the solutions and hence at most ν times and each vertex of $W \cap (\gamma(x) \setminus \beta(x))$ is traversed by one of the solutions.

For the converse direction, suppose that there is partial solution $\mathcal{C} = \{C_1, \dots, C_q\}$ compatible with (\mathcal{P}, in, out) at x of length h . We want to show that $A[x, \mathcal{P}, in, out] \leq h$. We again proceed by induction on the height of the subtree of the tree decomposition rooted at x . For leaves this is clear, as the only solution is \emptyset with length 0.

Next we assume that the claim already holds for the children of the node x by induction hypothesis. For introduce vertex node, v is isolated in G_x and hence must be also isolated in the solution. As the solution is compatible with (\mathcal{P}, in, out) at x , this implies that $in(v) = 0$, $out(v) = 0$, and $\{v\} \in \mathcal{P}$. Thus the algorithm set $A[x, \mathcal{P}, in, out] = A[y, \mathcal{P}', in', out']$. But we can use the same partial solution for the child y and the claim follows from the claim for the child.

Let now x be an introduce edge node, y its child, and uv the edge being introduced. Let b_{uv}^* be the sum of the number of times the edge (u, v) is traversed over all walks in \mathcal{C} and similarly with b_{vu}^* and the edge (v, u) . Obviously, if (u, v) is not in G , then we have $b_{uv}^* = 0$ and otherwise we have $b_{uv}^* \leq \kappa(u, v)$, and by compatibility also $b_{uv}^* \leq out(u)$ and $b_{uv}^* \leq in(v)$. In either case $b_{uv}^* \leq a_{uv}$ a similarly $b_{vu}^* \leq a_{vu}$.

If $b_{uv}^* = b_{vu}^* = 0$, then \mathcal{C} is also compatible with (\mathcal{P}, in, out) at y and, by induction hypothesis, we have $A[y, \mathcal{P}, in, out] \leq h$. If $v \notin P_u$, then $A[x, \mathcal{P}, in, out] = A[y, \mathcal{P}, in, out] \leq h$. Otherwise we have $A[x, \mathcal{P}, in, out] \leq \min_{\substack{0 \leq b_{uv} \leq a_{uv} \\ 0 \leq b_{vu} \leq a_{vu}}} (A[y, \mathcal{P}, in', out'] + b_{uv} \cdot \omega((u, v)) + b_{vu} \cdot \omega((v, u))) \leq A[y, \mathcal{P}, in, out] \leq h$ (note that $in' = in$ and $out' = out$ if $b_{uv} = b_{vu} = 0$).

If $b_{uv}^* + b_{vu}^* \geq 1$, then we remove all occurrences of (u, v) and (v, u) from the walks in \mathcal{C} , splitting them into more walks or shortening them. Call the resulting set of walks $\widehat{\mathcal{C}}$. As both u and v are in $\beta(x)$, each new walk starts and ends in $\beta(x) = \beta(y)$ and each vertex in $W \cap (\gamma(x) \setminus \beta(x))$ is traversed exactly as many times by $\widehat{\mathcal{C}}$ as by \mathcal{C} , i.e., at least once. Furthermore, each vertex and each edge is traversed at most as many times by $\widehat{\mathcal{C}}$ as by \mathcal{C} . Hence $\widehat{\mathcal{C}}$ is a partial solution of length $h - b_{uv}^* \cdot \omega((u, v)) - b_{vu}^* \cdot \omega((v, u))$.

Let H be the graph constructed for \mathcal{C} and \widehat{H} be obtained from H by removing the edge

$\{u, v\}$. If u and v are in the same connected component of \widehat{H} , then \widehat{C} is compatible with (\mathcal{P}, in', out') , where in' and out' are obtained by letting $b_{uv} = b_{uv}^*$ and $b_{vu} = b_{vu}^*$. Therefore $A[y, \mathcal{P}, in', out'] \leq h - b_{uv}^* \cdot \omega((u, v)) - b_{vu}^* \cdot \omega((v, u))$ and $A[x, \mathcal{P}, in, out] \leq h$.

If u and v are in different connected components of \widehat{H} , then let K_u be the component containing u and K_v be the component containing v . Note that, if K is the component containing u and v in H , then $V(K) = V(K_u) \cup V(K_v)$, since H and \widehat{H} only differ by the edge $\{u, v\}$. Let $P'_u = \beta(x) \cap K_u$, $P'_v = \beta(x) \cap K_v$, and $\mathcal{P}' = (\mathcal{P} \setminus \{P_u\}) \cup \{P'_u, P'_v\}$. Then \widehat{C} is compatible with $(\mathcal{P}', in', out')$, where in' and out' are again obtained by letting $b_{uv} = b_{uv}^*$ and $b_{vu} = b_{vu}^*$. Therefore $A[y, \mathcal{P}', in', out'] \leq h - b_{uv}^* \cdot \omega((u, v)) - b_{vu}^* \cdot \omega((v, u))$ and $A[x, \mathcal{P}, in, out] \leq h$.

Now suppose that x is a forget node, y its child and v the vertex being forgotten. Let $\mathcal{P} = \{P_1, \dots, P_s\}$. Let i be the number of times that v is traversed by \mathcal{C} . If v is waypoint, then $i \geq 1$ by the conditions on a partial solution. If v is traversed, then it is not isolated in H and therefore its connected component has a non-empty intersection with $\beta(x)$, which is in \mathcal{P} , say P_j . Then the intersection of this component with $\beta(y)$ is $P_j \cup \{v\}$ and \mathcal{C} is compatible with (P_j, in_i, out_i) at y . Thus $A[y, P_j, in_i, out_i] \leq h$ by induction hypothesis and also $A[x, \mathcal{P}, in, out] \leq h$.

If v is not traversed (which can only occur if it is not a waypoint), then it is isolated in H and \mathcal{C} is compatible with $(\mathcal{P} \cup \{\{v\}\}, in_0, out_0)$ at y . Then $A[y, \mathcal{P} \cup \{\{v\}\}, in_0, out_0] \leq h$ by induction hypothesis and also $A[x, \mathcal{P}, in, out] \leq h$.

Finally, let x be a join node and y and z its children. Let \mathcal{C}_y and \mathcal{C}_z be initially empty. For each walk $C \in \mathcal{C}$, take each maximal continuous part of walk C using only arcs of G_y and put it into \mathcal{C}_y and put each maximal continuous part only using arcs of G_z and put it into \mathcal{C}_z . Since every arc is introduced exactly once, graphs G_y and G_z do not share any arcs and therefore, the above maximal parts are not overlapping, apart from their endpoints and the multiset of arcs appearing in \mathcal{C} is a disjoint union of multisets of arcs of \mathcal{C}_y and \mathcal{C}_z . Thus the sum of their lengths h_y and h_z equals the length h of \mathcal{C} . Furthermore, only vertices of $\beta(x)$ might be incident with arcs of both G_y and G_z . Therefore, all the walks in \mathcal{C}_y and \mathcal{C}_z start and end in $\beta(x)$.

Each vertex u in $\gamma(x) \setminus \beta(x)$ is either in $\gamma(y) \setminus \beta(x)$ or $\gamma(z) \setminus \beta(x)$. Thus each such waypoint is traversed by at least one walk either of \mathcal{C}_y or of \mathcal{C}_z . No vertex and no edge is traversed more times by \mathcal{C}_y or by \mathcal{C}_z . Hence, \mathcal{C}_y and \mathcal{C}_z form a partial solution for y and z , respectively.

Let H_y be the subgraph of G_y formed by the edges used by the walks in \mathcal{C}_y and H'_y be the underlying undirected graph of H_y . Let $\mathcal{P}_y = \{K \cap \beta(x) \mid K \text{ is a connected component of } H_y \wedge K \cap \beta(x) \neq \emptyset\}$. Similarly define H_z, H'_z , and \mathcal{P}_z based on \mathcal{C}_z and H, H', \mathcal{P} based on \mathcal{C} . By assumption \mathcal{C} is compatible with (\mathcal{P}, in, out) , where for each $v \in \beta(x)$ vertex v is entered $in(v)$ times and exited $out(v)$ times summing over all walks in \mathcal{C} .

Since $E(H') = E(H'_y) \cup E(H'_z)$, partition \mathcal{P} is the finest common coarsening of \mathcal{P}_y and \mathcal{P}_z . Let $in_y, out_y, in_z, out_z: \beta(x) \rightarrow \{0, \dots, \nu\}$ be such that for each $v \in \beta(x)$ vertex v is entered $in_y(v)$ times and exited $out_y(v)$ times by the walks in \mathcal{C}_y and entered $in_z(v)$ times and exited $out_z(v)$ times by the walks in \mathcal{C}_z . By definition \mathcal{C}_y is compatible with $(\mathcal{P}_y, in_y, out_y)$ at y and \mathcal{C}_z is compatible with $(\mathcal{P}_z, in_z, out_z)$ at z . Therefore, by induction hypothesis, $A[y, \mathcal{P}_y, in_y, out_y] \leq h_y$ and $A[z, \mathcal{P}_z, in_z, out_z] \leq h_z$.

Since the walks in \mathcal{C}_y and \mathcal{C}_z form a partition of the walks in \mathcal{C} , we have for every $v \in \beta(x)$ that $in(v) = in_y(v) + in_z(v)$ and $out(v) = out_y(v) + out_z(v)$. Therefore $A[x, \mathcal{P}, in, out] \leq A[y, \mathcal{P}_y, in_y, out_y] + A[z, \mathcal{P}_z, in_z, out_z] \leq h_y + h_z = h$, finishing this implication.

Now we know that $A[r, \{\{w_0\}\}, w_0 \mapsto i, w_0 \mapsto i]$ is the shortest length of a sequence \mathcal{C} of

walks in $G_r = G$ that start and end in w_0 , each vertex in $W \cap (\gamma(r) \setminus \beta(r)) = W \setminus \{w_0\}$ is traversed at least once and not more than ν times, w_0 is traversed exactly $i \leq \nu$ times and each arc is traversed at most $\kappa(e)$ times. Note that a closed walk which is a solution to the instance forms such a sequence of one walk for an appropriate $i \in \{1, \dots, \nu\}$. Conversely, any such sequence can be turned into a single closed walk by concatenating the walks and the resulting walk represent a solution to the instance. Therefore, the answer of the algorithm is correct. ◀

F Omitted Details and Proofs from Section 5

CAPACITATED DOMINATING SET (CDS)

Input: An (undirected) graph G , a capacity function c , and positive integer k .

Question: Does there exist a capacitated dominating set $S \subseteq V(G)$ of size $|S| \leq k$?

CDS was shown to be $W[1]$ -hard with respect to treewidth by Dom, Lokshantov, Saurabh, and Villanger [27]. They prove that the graph has bounded treewidth by having a parameter-sized modulator to a forest. To prove that DIRECTED WAYPOINT ROUTING is $W[1]$ -hard by treedepth we adapt this reduction of Fomin et al. by (in essence) replacing high order bipartite subgraphs with a construction that exploits edges of bounded capacities. However, to prove that the resulting graph has bounded distance to constant treedepth, we require that CDS is $W[1]$ -hard for the distance to a constant treedepth. We get this $W[1]$ -hardness, in the theorem below, we carefully consider the structure of the modulator to forest from [27] to conclude that the forest is a disjoint union of stars. Hence, the reduction in their proof creates a graph with bounded distance to stars. This parameter implies not only bounded treewidth but also bounded treedepth, and more importantly for us, bounded distance to treedepth 2.

Proof of Theorem 10. $W[1]$ -hardness of CDS with respect to the treewidth of the input graph is shown by a reduction [27] from the MULTICOLORED-CLIQUE problem. To argue that the result has a bounded distance to stars, let us revisit the constructed graph. Say we have a MULTICOLORED-CLIQUE instance (G, k) where the graph has a k -partition V_1, \dots, V_k and our task is to find a clique K_k which has exactly one vertex in each of the parts. It can be assumed that $|V_i| = |V_j|$ for each $i, j \in [k]$, let $N = |V_i|$ be the size of each part. Let E_{ij} denote all edges between V_i and V_j . We reiterate how the authors reduce from (G, k) to a CDS instance (H, c, k') . We argue just that the H has bounded distance to stars and omit the proof of correctness which follows from the source paper [27].


The reduction uses the following two simplifying notions. *Marked vertex* is one which at the end of the reduction gets additional $k' + 1$ leaves and $k' + 1$ capacity – this forces the solution to contain such a vertex. Second notion is an (s, ℓ) -arrow from u to v which stands for adding s once-subdivided edges between u and v , and then adding ℓ leaves to v .

Let $k' = 7k(k - 1) + 2k$. Let c_x and c'_x be positive integers that depend on $x \in G$, their precise value is irrelevant to this proof. Now we list all the vertices and edges of H , see Figure 5:

- For each $i \in [k]$ there is a marked vertex \hat{x}_i ,
- for each $v \in V_i$ there is a vertex \bar{v} that neighbors \hat{x}_i ,
- for each $i \in [k]$ there are $k' - 1$ vertices S_i where each is adjacent to every vertex \bar{v} where $v \in V_i$,

- for each $i, j \in [k], i \neq j$ there are vertices \hat{y}_{ij} and \hat{z}_{ij} and for every $v \in V_i$ there is an (c_v, c'_v) -arrow from \bar{v} to \hat{y}_{ij} and an (c'_v, c_v) -arrow from \bar{v} to \hat{z}_{ij} ,
- for each $i, j \in [k], i < j$ there are five marked vertices $\hat{x}_{ij}, \hat{p}_{ij}, \hat{p}_{ji}, \hat{q}_{ij}, \hat{q}_{ji}$,
- for each $e = \{u, v\} \in E_{ij}$ where $u \in V_i$ and $v \in V_j$ there is a vertex \bar{e} that neighbors \hat{x}_{ij} , also there is a (c'_u, c_u) -arrow from \bar{e} to p_{ij} , a (c_u, c'_u) -arrow from \bar{e} to q_{ij} , a (c'_v, c_v) -arrow from \bar{e} to p_{ji} , and a (c_v, c'_v) -arrow from \bar{e} to q_{ji} , and there are $k' + 1$ vertices S_{ij} where each is adjacent to every vertex \bar{e} where $e \in E_{ij}$,
- last, for each $i, j \in [k], i \neq j$ there are marked vertices \hat{r}_{ij} and \hat{s}_{ij} with $(2N, 0)$ -arrows from \hat{y}_{ij} to \hat{r}_{ij} , the same arrows from \hat{p}_{ij} to \hat{r}_{ij} , from \hat{z}_{ij} to \hat{s}_{ij} , and from \hat{q}_{ij} to \hat{s}_{ij} .

Each of the vertices is assigned a capacity through the capacity function, which we ignore as it is unnecessary to this proof. Let M_1 be the union of $S_i \cup \{\hat{x}_i\}$ for each $i \in [k]$, M_2 be the union of $S_{ij} \cup \{\hat{x}_{ij}, \hat{y}_{ij}, \hat{z}_{ij}, \hat{p}_{ij}, \hat{q}_{ij}, \hat{r}_{ij}, \hat{s}_{ij}\}$ for each $i, j \in [k], i \neq j$. Set M_1 has size $(k' + 1) \cdot k$ while M_2 has size $(k' + 1 + 7) \cdot \binom{k}{2}$. Let $M = M_1 \cup M_2$. We see that $|M| = 7k^4/2 + 6k^3 - 7k^2/2 - 2k$, i.e., it is bounded in terms of k , it remains to show that M is a modulator to stars.



■ **Figure 5** Part of the graph H described in the proof of Theorem 10 from [27] that models an edge $e = uv$. Marked vertices are depicted as empty, thick arrows stand for (c, c') -arrows for some integers c, c' .


Observe that adding an arrow from u and v creates only edges that are incident to either u or v . Let C be the set that contains \bar{v} for each $v \in V_i$, \bar{e} for each $e \in E_{ij}$. We claim that $H - M - C$ contains no edges and adding back the independent set C results in the graph $H - M$ being a disjoint union of stars. First, note that all arrows have either both endpoints in M or one endpoint in M and the second in C , implying that none of the edges added due to arrows are present in $H - M - C$. The (non-arrow) edges that were added to H are $\{o_i, \bar{v}\}$ for each $i \in [k], o_i \in \{\hat{x}_i\} \cup S_i, v \in V_i$, and similarly $\{o_{ij}, \bar{e}\}$ for each

$i, j \in [k], o_{ij} \in \{\hat{x}_{ij}\} \cup S_{ij}, \bar{e} \in E_{ij}$, which are edges between M and C . Finally, all the marked vertices are in M , therefore all leaves attached to the marked vertices at the end of the construction are isolated in $H - M$. ◀

Note that Fomin et al. [32, Section 5] showed W[1]-hardness of HAMILTONIAN CYCLE with respect to clique-width by a reduction from CDS parameterized by treewidth. The improved analysis of Theorem 10 does not strengthen this result as the construction still creates high order bipartite subgraphs.

Proof of Theorem 5. Let (G, c, k) be an instance of CDS where G is a graph, $c: V(G) \rightarrow \mathbb{N}_0$ is the capacity function, and k is a maximum size of the solution; let M be the modulator to stars of minimum cardinality in G . We build an instance $(G', W, \omega, \kappa, b)$ of DWRP that has modulator of size $7 \cdot |M| + 7$ to treedepth 86. Let us first show a construction of gadgets one by one while noting intuition for how they work. Each gadget has dedicated connecting vertices which we mark with ‘in’ (e.g. u_{in}), ‘out’, or ‘io’. In the final instance, within G' every gadget will be connected to other gadgets by edges incident only to its connecting vertices; notably ‘in’ vertices can be incident to arcs that are coming into the gadget, ‘out’ vertices can be incident to outgoing arcs, and ‘io’ can be incident to both types of arcs. If the above assertions about arcs incident to gadget’s vertices in a graph is true then we say that the graph *contains* the gadget. When a gadget X contains in its construction a vertex u we use u^X to denote this vertex in the resulting graph. All the edges have weight 1. Unless stated otherwise, every edge has capacity n ; this is always sufficient by Observation 7. Last before describing the gadgets, let us note that the crucial property of the construction is that to get from one terminal to another we have to traverse exactly two non-terminal vertices. As the final budget is set to exactly $3 \cdot (\text{number of terminals})$ it follows that every terminal is visited exactly once.

To force $p \in \mathbb{N}^+$ traversals we define *force- p -traversals-gadget*. It consists of three non-terminal vertices $u_{\text{in}}, u_{\text{out}}, w$ and terminal vertices v_1, \dots, v_p joined with edges $(u_{\text{in}}, v_i), (v_i, w)$ for every $i \in [p]$ and (w, u_{out}) of capacity p , see Figure 2.




■ **Figure 6** Force-5-traversals-gadget

▷ **Claim 20.** Let G' be a graph that contains a force- p -traversals-gadget X . Then every solution to DWRP in G' traverses from u_{in}^X to u_{out}^X exactly p times.

Proof. Every terminal vertex v_i^X for $i \in [p]$ requires one (unique) traversal from u_{in}^X to u_{out}^X . The capacity of (w, u_{out}^X) upper bounds the number of traversals from u_{in}^X to u_{out}^X by p . ◀


In order to define the second gadget, called *cover- z -gadget*, we first introduce an undirected *raw-cover- z -gadget* (see Figure 7) that was used for HAMILTONIAN CYCLE (gadget is called L_2 in [32, Lemma 9]). For a fixed vertex z let a *raw-cover- z -gadget* be an undirected graph that consist of a path R_1 on terminals $s_{\text{in}}, e, f, b, a, x_{\text{in}}, w_1, w_2, \dots, w_9, y_{\text{out}}, d, c, h, g, t_{\text{out}}$ (in this order) together with additional edges $\{f, w_3\}, \{w_1, w_6\}, \{w_4, w_9\}$, and $\{w_7, h\}$. Moreover,

the vertex z is also considered to be part of the raw-cover- z -gadget; in this context we call it z_{io} and it is connected via edges $\{c, z_{io}\}$ and $\{z_{io}, b\}$. Let R_2 be the path $s_{in}, e, f, w_3, w_2, w_1, w_6, w_5, w_4, w_9, w_8, w_7, h, g, t_{out}$ and P be the path $x_{in}, a, b, z_{io}, c, d, y_{out}$.



■ **Figure 7** A raw-cover- z -gadget [32] and its two partial solutions for HAMILTONIAN CYCLE problem.

For a fixed vertex z , we create *cover- z -gadget* by first creating a raw-cover- z -gadget, then changing the following edges into arcs (s_{in}, e) , (e, f) , (f, b) , (b, z) , (z, c) , (c, h) , (h, g) , (g, t_{out}) , (x_{in}, w_1) , (w_9, y_{out}) , (w_1, w_6) , (w_4, w_9) , (w_3, w_4) , and (w_6, w_7) , then subdividing every edge and arc with two non-terminal vertices, and last replacing every $\{u, v\}$ edge with arcs uv and vu , see Figure 8. Let the paths P , R_1 , and R_2 be directed paths in the cover- z -gadget that were created by subdivision and orientation (in the direction of order of vertices from their definition) of their analogues in the raw-cover- z -gadget.



■ **Figure 8** Cover- z -gadget is created via selective orientation and subdivision of the raw-cover- z -gadget.

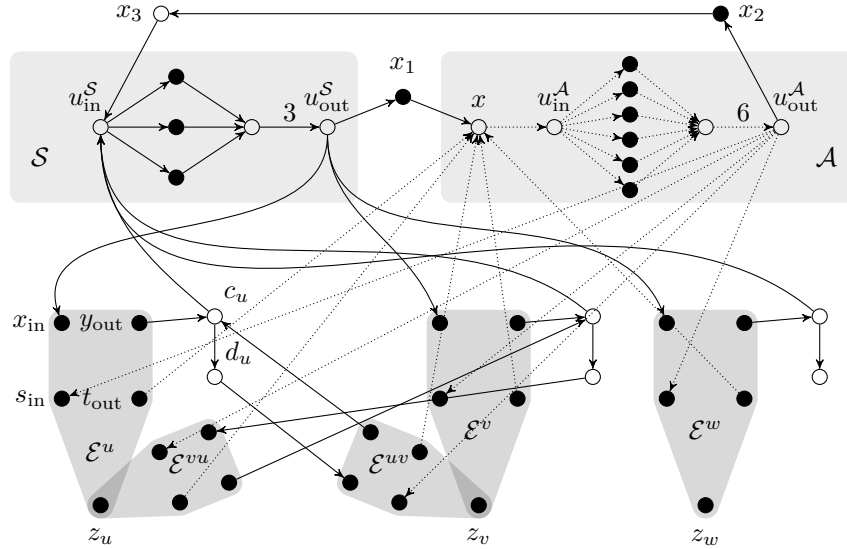
▷ **Claim 21.** Let G be a graph that contains the cover- z -gadget X on $G[X]$, i.e., $G[X]$ is isomorphic to the cover- z -gadget and all arcs in $E(G) \setminus E(G[X])$ that end in X are incident with the vertices x_{in}^X , s_{in}^X , or z_{io}^X and the arcs in $E(G) \setminus E(G[X])$ that start in X are incident with the vertices y_{out}^X , t_{out}^X , or z_{io}^X in X . If S is a minimal solution to DWRP that traverses every terminal vertex of X exactly once, then the solution includes either the path R_1 , or two paths P and R_2 as continuous subpaths.

Proof. Let S be any minimal solution that traverses the terminal vertices of the cover- z -gadget X exactly once. Note that once S enters a non-terminal vertex that was created by subdividing edges, S has to follow through to the other terminal, as the initial terminal vertex cannot be visited again. Terminal vertices e , a , w_2 , w_5 , w_8 , d , g have only two terminal “neighbors”. Hence, they must be visited immediately after their first terminal neighbor is visited. Solution traverses from s_{in} through e to f where it has two choices. If S continues from f to b , then it is forced through a to x_{in} , by the above argument, and then the only way to continue is to w_1 and through w_2 to w_3 . After that the only option is w_4 , forced to visit w_5 continues to w_6 and w_7 . Then S must visit w_8 , continues to w_9 and y_{out} . Again, S must visit d and continues to c, h, g, t_{out} as no other options remain. In the other case,

where S initially goes from s_{in} through e to f and continues to w_3 one observes that then it is forced through w_2 to w_1 where no other option remains than to go to w_6 , by similar logic we have $w_5, w_4, w_9, w_8, w_7, h$ after which, again, the only option is to continue through g to t_{out} . This traversal leaves out $x_{\text{in}}, a, b, z, c, d, y_{\text{out}}$ where the only way to visit these terminals by S is exactly in this order. \triangleleft

Now we present the construction, see Figure 9; recall that we are reducing from an instance (G, c, k) of CDS and we aim to create an instance $(G', W, \omega, \kappa, b)$ of DWRP. Let S be a force- $(k+1)$ -traversals-gadget that we call *selection-gadget* and let \mathcal{A} be a force- $(2|E(G)| + |V(G)| + 1)$ -traversals-gadget. We create an auxiliary non-terminal vertex named x and add an arc $(x, u_{\text{in}}^{\mathcal{A}})$. Let x together with \mathcal{A} be called *auxiliary-gadget*. We mutually connect the selection-gadget and the auxiliary-gadget by creating terminal vertices x_1, x_2 , non-terminal vertex x_3 , and arcs $(u_{\text{out}}^{\mathcal{A}}, x_2), (x_2, x_3), (x_3, u_{\text{in}}^{\mathcal{S}}), (u_{\text{out}}^{\mathcal{S}}, x_1), (x_1, x)$.

For every vertex $u \in V(G)$ we add vertices z_u, c_u , and d_u to $V(G')$, then we create a cover- z_u -gadget denoted \mathcal{E}^u . We add arcs $(u_{\text{out}}^{\mathcal{S}}, x_{\text{in}}^{\mathcal{E}^u}), (y_{\text{out}}^{\mathcal{E}^u}, c_u), (c_u, d_u), (c_u, u_{\text{in}}^{\mathcal{S}}), (u_{\text{out}}^{\mathcal{A}}, s_{\text{in}}^{\mathcal{E}^u}), (t_{\text{out}}^{\mathcal{E}^u}, x)$ to $E(G')$, and we set the capacity of (c_u, d_u) to $c(u)$ (the capacity of u). For every arc (u, v) such that $\{u, v\} \in E(G)$ we create a cover- z_v -gadget denoted $\mathcal{E}^{u,v}$. Then we add arcs $(d_u, x_{\text{in}}^{\mathcal{E}^{u,v}}), (y_{\text{out}}^{\mathcal{E}^{u,v}}, c_u), (u_{\text{out}}^{\mathcal{A}}, s_{\text{in}}^{\mathcal{E}^{u,v}}), (t_{\text{out}}^{\mathcal{E}^{u,v}}, x)$. Note that in the above we created two gadgets, $\mathcal{E}^{u,v}$ and $\mathcal{E}^{v,u}$, for each $\{u, v\} \in E(G)$. Finally, we set the budget $b = 132|E(G)| + 69|V(G)| + 3k + 12$ which concludes construction of the DWRP instance.



■ **Figure 9** An example of a complete constructed DWRP instance for graph $G = (\{u, v, w\}, \{\{u, v\}\})$, $k = 2$, (capacities omitted). Top half contains the selection- and auxiliary-gadget along with their connections through x_1, x_2, x_3 . Bottom half corresponds to the graph G . Gray boxes hide inner parts of the cover-gadgets, see Figure 8. Dotted arrows signify arcs that make-up the auxiliary-gadget or connect cover-gadgets to the auxiliary-gadget.

▷ **Claim 22.** If (G, c, k) is a yes-instance of CDS, then the instance $(G', W, \omega, \kappa, b)$ of DIRECTED WAYPOINT ROUTING PROBLEM with budget $b = 132|E(G)| + 69|V(G)| + 3k + 12$ has a solution.

Proof. Let S ($|S| \leq k$) be the set of vertices in the solution of the instance (G, c, k) of CDS and $f: (V(G) \setminus S) \rightarrow S$ be the domination mapping. Solution to DIRECTED WAYPOINT

ROUTING is presented as two vertex-disjoint walks, one starting in the selection-gadget and one in the auxiliary-gadget. These two walks are then joined via the auxiliary vertices x_1 , x_2 , and x_3 .

First walk: For each $s \in S$ we traverse the selection-gadget once, covering one of its terminals, and then through the P path of \mathcal{E}^s , covering z_s , then to c_s . For each $x \in f^{-1}(s)$ the walk continues from c_s to d_s , through the P path in $\mathcal{E}^{s,x}$ covering z_x , returning back to c_s . Finally, the walk jumps back to the selection-gadget.

Second walk: Let E be a cover- z -gadget within G' . The second walk goes through the auxiliary-gadget while covering one of its terminals, then it continues through a path of E . More precisely, if the path P of E was traversed by the first walk then the second walk goes through R_2 , otherwise it traverses R_1 . Last, the walk goes to x , returning back to the auxiliary-gadget.

These two walks are joined into a single walk by adding cycle $u_{\text{out}}^S, x_1, x, u_{\text{in}}^A, v_1^A, w^A, u_{\text{out}}^A, x_2, x_3, u_{\text{in}}^S, v_j^S, w^S, u_{\text{out}}^S$ for some i and j so that uncovered terminals of the selection-gadget and the auxiliary-gadget are traversed by this cycle.

To compute the walk weight, note that every terminal is traversed exactly once and that distance of every pair of consecutive terminals along the walk is 3, hence, we used the budget of $3 \cdot T$ where T is the number of terminals. There are $k + 1$ terminals in the selection gadget, $2|E(G)| + |V(G)| + 1$ terminals in the auxiliary-gadget, terminals x_1 and x_2 , then in every cover- z -gadget we have 21 terminals (note that z is not counted here), and last we have terminals z_u for every $u \in V(G)$. In total this is $3 \cdot (k + 1 + 2|E(G)| + |V(G)| + 1 + 2 + 21 \cdot (2|E(G)| + |V(G)|) + |V(G)|)$ which simplifies to $132|E(G)| + 69|V(G)| + 3k + 12$ which is equal to the budget b . \triangleleft

\triangleright Claim 23. If DWRP instance $(G', W, \omega, \kappa, b)$ with budget $b = 132|E(G)| + 69|V(G)| + 3k + 12$ is a yes-instance, then the instance (G, c, k) of CDS is a yes-instance.

Proof. Let us identify partial walks that must be present in any solution and argue that they compose exactly into a walk shown in the proof of Claim 22. Observe that the budget is exactly so that every terminal can be assigned budget of 3, and as each pair of terminals is at distance at least 3 we conclude that the solution traverses every terminal exactly once.

First, by Claim 20 the selection-gadget is traversed exactly $k + 1$ times and the auxiliary-gadget is traversed $2|E(G)| + |V(G)| + 1$ times. As x_1 and x_2 need to be covered the paths $u_{\text{out}}^S, x_1, x, u_{\text{in}}^A$ and $u_{\text{out}}^A, x_2, x_3, u_{\text{in}}^S$ need to be traversed exactly once. These together with one traversal of the selection-gadget and one traversal of the auxiliary-gadget give us the cycle that joined the first and second walk.

Next, from the vertex u_{out}^S the remaining k walks must continue, and as all its out-neighbors are terminals the solution continues into k disjoint out-neighbors. These out-neighbors belong to some cover- z_u -gadgets for a set $u \in S$, consequently $|S| = k$. More precisely, these out-neighbors are x_{in}^u for $u \in S$. Let $\mathcal{E}^S = \{\mathcal{E}^u \mid u \in S\}$ and similarly for any vertex w let $w^{\mathcal{E}^S} = \{w^{\mathcal{E}^u} \mid u \in S\}$. As vertices $x^{\mathcal{E}^S}$ are beginnings of the respective P paths within \mathcal{E}^S we know by Claim 21 that \mathcal{E}^S are traversed by their P and R_2 paths, covering z_u for each $u \in S$ in the process. After the traversal of P the walks continue to c_u for each $u \in S$. From there each walk can either return back to the selection-gadget or go to d_u , but as (c_u, d_u) has capacity $c(u)$ such traversal can happen at most $c(u)$ many times. From d_u the walk continues through $\mathcal{E}^{u,w}$ for any $w \in N(u)$, covering z_w in the process, and returning to c_u .

For every other cover-gadget \mathcal{E}^v where $v \in V(G) \setminus S$ we know the walks do not enter the gadgets from u_{out}^S (the only in-neighbor of x_{in}^v) so these must be traversed by their R_1 paths. All of the R_1 and R_2 traversals of the cover-gadgets can only continue to the auxiliary

gadget, and vice-versa, all walks from the auxiliary gadget (except for the connecting cycle) continue only continue to the cover-gadgets paths R_1 or R_2 .

This forced structure of the solution is identical to the one described in Claim 22 and so there is a straight-forward translation back to the capacitated dominating set, namely, the set S is a capacitated dominating set of G . \triangleleft

\triangleright **Claim 24.** If G has modulator of size k to treedepth td , then G' has modulator of size $7 \cdot k + 7$ to treedepth $7 \cdot td + 72$.

Proof. Recall that G having treedepth bounded by td is witnessed by an *elimination forest* which is a family of rooted trees of height at most td where the trees are over vertices $V(G)$ but their edges are unrelated to $E(G)$; however, they satisfy a single property – for every edge $uv \in E(G)$ the vertices u and v are in the forest in an ancestor-descendant relationship; we call such edges *satisfied*. For example, if we place a set of vertices V_1 on path rooted in one of its endpoints, then all edges in $V_1 \times V_1$ have their endpoints in an ancestor-descendant relationship, however, the height of such a tree is $|V_1|$.

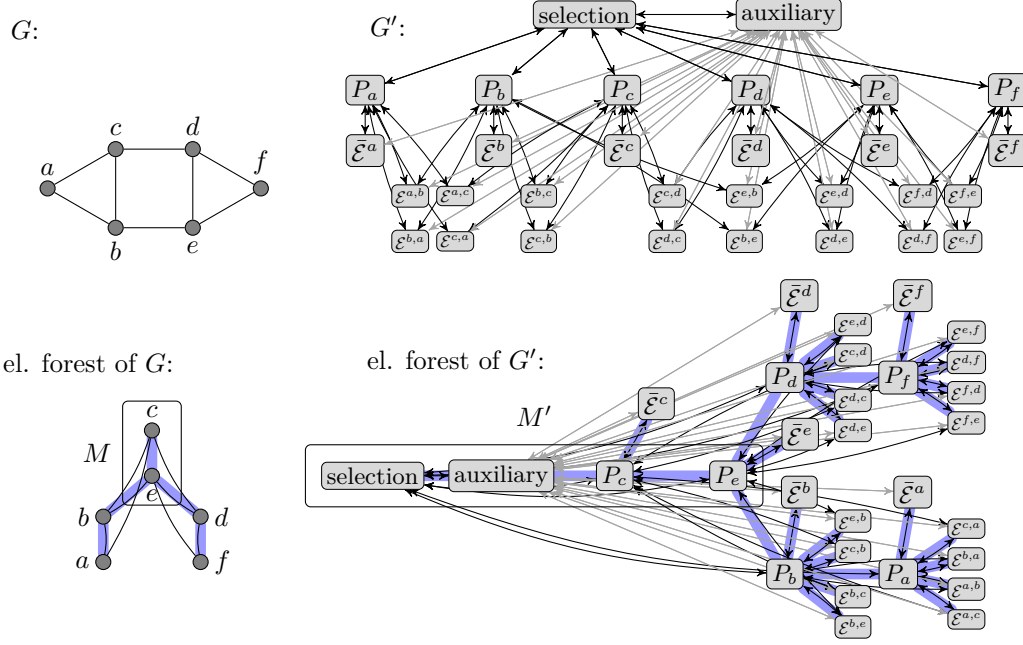
To show the claim we will proceed as follows. We take the set of stars created by removing the modulator M from G , and create an elimination forest where M forms a path from the root, and the stars are attached below – this naturally satisfies all the edges of G . Then, we replace each vertex of the graph with a constant number of vertices of G' , placing the new set on a path in-place of the original vertices, and arriving at an elimination forest of G' . As the sets are paths, they satisfy all the edges within each path separately; then we argue that all edges of G' are still in the ancestor-descendant relationship. We then observe that the vertices M' that replaced the vertices of M form a new path from the root and contains $7 \cdot k + 7$ vertices; and whose removal creates an elimination forest of constant depth.

Consider M to be the modulator of size k and \mathcal{F}_o to an elimination forest that consists of trees T_1, \dots, T_ℓ and witnesses that the treedepth of $G \setminus M$ is at most td . Let \mathcal{F} be an elimination forest that contains a single tree with M being its path from the root and with trees T_1, \dots, T_ℓ of \mathcal{F}_o attached below M . Observe that the height of \mathcal{F} is $k + td$. We aim to create an elimination forest \mathcal{F}' of G' such that when a path from the root M' of small size is removed, we get an elimination forest \mathcal{F}'_o of the desired size.

The elimination forest \mathcal{F}' consists of a single tree T' (see Figure 10). The tree T' contains a path R from the root containing vertices $x, u_{\text{in}}^A, u_{\text{out}}^A, w^A, u_{\text{in}}^S, u_{\text{out}}^S, w^S$, i.e., all non-terminal vertices of the selection-gadget and the auxiliary-gadget. Below R we attach a path on vertices x_1, x_2 , and x_3 . We also attach below R (independently) each terminal vertex of the selection-gadget and the auxiliary-gadget. Next, below R we attach the tree $T_i \in \mathcal{F}$ while replacing each vertex $u \in T_i$ (where $u \in V(G)$) by a path P_u containing vertices $z_u, c_u, d_u, s_{\text{in}}^{\mathcal{E}^u}, x_{\text{in}}^{\mathcal{E}^u}, y_{\text{out}}^{\mathcal{E}^u}, t_{\text{out}}^{\mathcal{E}^u}$. Below each P_u we attach all vertices of \mathcal{E}^u that are not in the decomposition yet. Recall that the elimination forest \mathcal{F} has the property that for every edge $\{u, v\} \in E(G)$ either u is the ancestor of v or v is the ancestor of u in \mathcal{F} . So for every edge $\{u, v\} \in E(G)$ we can find a leaf $\ell_{u,v}$ of \mathcal{F} such that the path from $\ell_{u,v}$ to the root contains both u and v . As the last step, for every edge $\{u, v\} \in E(G)$ we attach to the bottommost vertex of $P_{\ell_{u,v}}$ a path containing all vertices of $\mathcal{E}^{u,v}$.

To show that \mathcal{F}' is an elimination forest we need to show that all the edges of G' are in an ancestor-descendant relationship within \mathcal{F}' – let us call such edges *resolved*. Note that if vertices are attached as a path in \mathcal{F}' then all edges among these vertices are resolved. All edges incident to the 7 vertices in R are resolved as these vertices are placed as ancestors of all the other vertices of G' . Placing x_1, x_2, x_3 below R resolved all edges incident to them because they are adjacent only to R . Attaching all terminal vertices of the selection-gadget

and the auxiliary gadget below R resolves all their incident edges as they are also adjacent only to R . For any vertex $u \in V(G)$ the path P_u contains all vertices that have edges connecting to $G \setminus (V(\mathcal{E}^u) \cup R)$. By attaching \mathcal{E}^u as a path below P_u we resolved all edges between two vertices of $P_u \cup \mathcal{E}^u$. Last, all the edges within or incident to $\mathcal{E}^{u,v}$ for any $u, v \in E(G)$ are resolved by placing the vertices of $\mathcal{E}^{u,v}$ below $P_{\ell_{u,v}}$ as $\mathcal{E}^{u,v}$ has edges only to c_u, d_u , and z_v which are all ancestors of $\ell_{u,v}$ in \mathcal{F}' .



■ **Figure 10** An example input graph G with a respective result of the reduction G' depicted as a set of simplified gadgets and paths of Claim 24 together with their connections. We use \mathcal{E}^x for $\mathcal{E}^x \setminus P_x$. Bottom half: elimination forests of G and G' ; forest for G' is rotated to preserve space.

Now, let $M' = R \cup \bigcup_{u \in M} P_u$ and let \mathcal{F}'_o be \mathcal{F}' with M' removed. By definition $|M'| = |R| + \sum_{u \in M} |P_u| = 7 + |M| \cdot 7$. As M was a path from the root of the tree in \mathcal{F}_o we observe that the construction replaced each vertex of $u \in M$ by P_u which is a path, hence, their concatenation along with R creates a path M' from the root of \mathcal{F}'_o . Hence, removing M' from \mathcal{F}' results in a forest where each tree is rooted in its vertex that lied below M' in \mathcal{F}' . We observe that the trees of \mathcal{F}'_o consists of:

- x_1, x_2, x_3 of height 3,
- terminal vertices of the selection-gadget and the auxiliary-gadget of height 1,
- $\mathcal{E}^u = \mathcal{E}^u \setminus P_u$ for each $u \in M$ of height $|\mathcal{E}^u| - |P_u| = 72 - 7 = 65$,
- trees T'_1, \dots, T'_ℓ corresponding to T_1, \dots, T_ℓ in \mathcal{F}_o , height of which we argue below.

The trees T_i have height at most td . Each vertex $u \in V(T_i)$ was replaced with P_u to create T'_i , multiplying its height by 7. Moreover, for each $u \in V(T_i)$ we attached to P_u vertices in $V(\mathcal{E}^u) \setminus P_u$, which is of size 65 (as computed above). Independently of that, for each edge uv we attached entire \mathcal{E}^{uv} below either P_u or P_v , whichever was lower in the tree. As each \mathcal{E}^{uv} is added independently, their addition increases the height of T'_i at most once. As $|\mathcal{E}^{uv}| = 72$ we conclude that each T'_i has height at most $7 \cdot td + 72$. \triangleleft

The above claim together with the fact that disjoint union of stars has treedepth 2 gives us the following corollary.

► **Corollary 25.** *If G has modulator of size k to stars, then G' has modulator of size $7 \cdot k + 7$ to treedepth 86.*

As noted in Theorem 10 CDS is W[1]-hard w.r.t. distance to stars. Assuming G has a modulator to stars M we have that G' has a modulator of size at most $7 \cdot |M| + 7$ to treedepth 86 by Corollary 25. By Claims 22 and 23 the constructed instance is equivalent to the initial one, therefore, DIRECTED WAYPOINT ROUTING PROBLEM is W[1]-hard w.r.t. distance to constant treedepth. ◀

G Pathwidth Hardness for Unweighted Graphs

Marx et al. [52] showed the W[1]-hardness of DTSP w.r.t. pathwidth by a chain of reductions. In the last step, they ensured that some vertices are visited at most once by significantly increasing the weight of their incoming edges. An obvious idea to show the W[1]-hardness also for unweighted graphs would be to subdivide the edges of large weight. However, as this creates new vertices that might not be visited by the solution to the original instance, this only shows the W[1]-hardness for DsTSP. In this section we revisit the proof by Marx et al. [52] with two goals. First, we want to give precise bounds on the resulting pathwidth. Second, we want to modify the last step in order to prove the W[1]-hardness for unweighted graphs also for DTSP.

► **Theorem 26.** *DTSP is W[1]-hard w.r.t.*

a) *distance to pathwidth 2 and*

b) *distance to pathwidth 3 in unweighted graphs.*

Unless ETH fails, there is no algorithm for DTSP with running time

- i) $n^{o(\sqrt{d_2})}$, where d_2 is the distance to pathwidth 2,
- ii) $n^{o(d_3)}$, where d_3 is the distance to pathwidth 3,
- iii) $n^{o(\sqrt{d_3})}$ in unweighted graphs, where d_3 is the distance to pathwidth 3, or
- iv) $n^{o(d_4)}$ in unweighted graphs, where d_4 is the distance to pathwidth 4.

The chain of parameterized reductions given by Marx et al. [52] starts with CLIQUE and continues over MULTICOLORED BICLIQUE, through EDGE BALANCING and CONSTRAINED CLOSED WALK to DTSP (see the figure [52, page 46]). While they do not explicitly state W[1]-hardness, it follows from the fact that CLIQUE is W[1]-hard and all the reductions are parameterized reductions. To prove our theorem, we only revisit the last two steps of the chain, where we only analyze the structure obtained in the second to last step and only modify the last step. Let us introduce the problems.

EDGE BALANCING

Input: A directed graph D and a set X_e of positive integers for every edge $e \in E(D)$.


Question: Is it possible to select $\chi(e) \in X_e$ for every $e \in E(D)$ such that for every $v \in V(D)$ we have $\sum_{u \in N^{\text{in}}(v)} \chi((u, v)) = \sum_{u \in N^{\text{out}}(v)} \chi((v, u))$?

► **Lemma 27** (Marx et al. [52, Lemma 86]). *Assuming ETH, EDGE BALANCING has no $f(k)n^{o(k)}$ time algorithm for any computable function f , where k is the number of vertices of D . Moreover, the problem is W[1]-hard with respect to k .*


CONSTRAINED CLOSED WALK

Input: An unweighted directed graph G and a set $U \subseteq V(G)$.
Question: Is there a closed walk (of any length) that visits every vertex at least once and each vertex in U exactly once?

The reduction from EDGE BALANCING to CONSTRAINED CLOSED WALK uses two gadgets that we now present.



■ **Figure 11** Gadget H_s used in the reduction from EDGE BALANCING to CONSTRAINED CLOSED WALK.



■ **Figure 12** Gadget H_X for the set $X = \{x_1, x_2, x_3, x_4\}$. The gray rectangles schematically represent the internal parts of the gadgets H_{x_i} , while the vertices $a_{in}, a_{out}, v, c_{out}$ play the role of the vertices $a_{in}, a_{out}, b_{in}, b_{out}$ for them, respectively.

The gadgets use connecting vertices with subscripts 'in' or 'out' similarly as in Section 5, the other vertices are called *internal*.

First, for a positive integer s we construct a gadget H_s depicted on Figure 11 with 4 connecting vertices: $a_{in}, a_{out}, b_{in},$ and b_{out} . As we rely on the original proof concerning the correctness, the only important property of the gadget for us is that the pathwidth of the graph $H_s \setminus \{a_{in}, a_{out}, b_{in}, b_{out}\}$ is 2.

Based on this gadget, for a set X of positive integers, another gadget H_X is constructed as follows (see Figure 12). First we introduce 4 connecting vertices $a_{in}, a_{out}, c_{in},$ and c_{out} , an internal vertex v , and an edge $\{c_{in}, v\}$. Then for each $x \in X$ we introduce gadget H_x and identify its connecting vertices $a_{in}, a_{out}, b_{in},$ and b_{out} with $a_{in}, a_{out}, v,$ and c_{out} , respectively.

Again the only important property of gadget H_X for us is that the pathwidth of the graph $H_X \setminus \{a_{\text{in}}, a_{\text{out}}, c_{\text{in}}, c_{\text{out}}\}$ is 3, while if we also remove vertex v , the pathwidth drops to 2.

Now we are ready to describe the reduction from EDGE BALANCING to CONSTRAINED CLOSED WALK. Let $(D, (X_e)_{e \in E(D)})$ with $V(D) = \{w_1, \dots, w_k\}$ be an instance of EDGE BALANCING. We construct graph D^* as follows. First we introduce vertices of D into D^* together with two auxiliary vertices c_{in} and c_{out} . Let $Z = V(D) \cup \{c_{\text{in}}, c_{\text{out}}\}$. Then, for each edge $e = (w_i, w_j) \in E(D)$ we construct a copy of the gadget H_{X_e} , add it to D^* , and identify its vertices $a_{\text{in}}, a_{\text{out}}, c_{\text{in}}$, and c_{out} with vertices w_i, w_j, c_{in} , and c_{out} , respectively.

Then we add some more paths into D^* . Again, as we rely on the original paper concerning the correctness, the exact numbers of the paths are not important to us. We add directed paths of length 2 (the internal vertices of them will be newly introduced)

1. from c_{in} to w_i ,
2. from w_i to c_{out} ,
3. and from c_{out} to c_{in} .

We let $U = V(D^*) \setminus Z$ be the set of vertices that should be traversed exactly once. This finishes the construction of the instance (D^*, U) of CONSTRAINED CLOSED WALK. The correctness of the reduction follows from [52, Lemma 90].

Since each connected component of graph $D^* \setminus Z$ consist either of the internal vertices of a gadget H_{X_e} or of a single internal vertex of the introduced path, the pathwidth of $D^* \setminus Z$ is 3. Hence, the distance of D^* to pathwidth 3 is $O(k)$, where k is the number of vertices of D .


Moreover, let Z' be the set Z together with the vertex v from every gadget H_{X_e} introduced. Since we introduced gadget H_{X_e} for every edge of D and D is a simple graph, we have that $|Z'| = O(k^2)$. However, each connected component of graph $D^* \setminus Z'$ consist either of the internal vertices of a gadget H_x or of a single internal vertex of the introduced path, and thus, the pathwidth of $D^* \setminus Z'$ is 2. Therefore, the distance of D^* to pathwidth 2 is $O(k^2)$ and the reduction is parameterized reduction even to CONSTRAINED CLOSED WALK parameterized by distance to pathwidth 2.

Together we obtain the following.

► **Lemma 28.** *CONSTRAINED CLOSED WALK is $W[1]$ -hard w.r.t. distance d_2 to pathwidth 2. Moreover, unless ETH fails, there is no algorithm for the problem with running time $f(d_2)n^{o(\sqrt{d_2})}$ or $f(d_3)n^{o(d_3)}$ for any computable function, where d_3 is the distance to pathwidth 3.*

Proof of Theorem 26. To reduce CONSTRAINED CLOSED WALK to DTSP Marx et al. [52] suggested to use edge weights to ensure that all vertices in U are visited at most once. Note that if n is the total number of vertices in the instance of CONSTRAINED CLOSED WALK, then any solution walk is of length at most n^2 , as it consist of at most n cycles, each of length at most n . Hence, if we give weight $2n^2$ to every edge with head in a vertex of U and weight 1 to every other edge and set the budget to $b = 2n^2|U| + n^2$, then the obtained instance of DTSP is equivalent to the instance of CONSTRAINED CLOSED WALK. Indeed, any closed walk which visits each vertex of D^* and each vertex of U exactly once has weight at most b and any closed walk that visits every vertex can visit each vertex of U at most once, as otherwise it would exceed the budget. This proves parts **a)**, **i)**, and **ii)** of the theorem.

To prove the remaining parts of the theorem, that is, for unweighted graphs, we take a slightly different approach. We divide every vertex $u \in U$ into two vertices \tilde{u} and \hat{u} , redirect all the edges with head in u to \tilde{u} and all the edges with tail in u to \hat{u} . Then we connect \tilde{u} to \hat{u} by a path of length $2n^2$. Denote the modified graph \hat{D} and set the budget again to



■ **Figure 13** Modified gadget H_s . The “snake” paths represent directed paths of length $2n^2$ each.

$b = 2n^2|U| + n^2$, to obtain an instance $(\widehat{D}, \mathbf{1}, b)$ of DTSP. Figure 13 depicts the gadget H_s from Figure 11 after this modification.

Let us now prove the equivalence of the instances (D^*, U) of CONSTRAINED CLOSED WALK and $(\widehat{D}, \mathbf{1}, b)$ of DTSP. If C is a closed walk in D^* that visit each vertex at least once and each vertex of U exactly once, then by replacing each visit to a vertex $u \in U$ by the \tilde{u} - \hat{u} -path, we obtain a walk \widehat{C} in \widehat{D} of total length at most $|U| \cdot 2n^2 + n^2 = b$ visiting each vertex of \widehat{D} . Conversely, any closed walk \widehat{C} that visits each vertex in \widehat{D} must visit vertex \tilde{u} for each $u \in U$. Whenever it visits it, it must continue along the path of length $2n^2$ to \hat{u} . Hence, if it obeys the budget, it can visit each such vertex at most once. Thus, replacing each traverse of a \tilde{u} - \hat{u} -path with a visit to vertex u , we obtain a closed walk in D^* which visits each vertex and visit each vertex of U exactly once.

It is easy to see that if the pathwidth of (part of) D^* is p then the pathwidth of (the corresponding part of) \widehat{D} is at most $2p + 3$. We next prove more tight bounds.

Note that the modified gadget H_s (Figure 13), after removal of connection vertices, has pathwidth 3, as it is a subgraph of a $3 \times \ell$ grid for a suitable ℓ . Thus, each modified gadget H_{X_e} has pathwidth 4 after removal of the connection vertices, since we can put \hat{v} into every bag and the \tilde{v} - \hat{v} -path has pathwidth 1.

As $U = V(D^*) \setminus Z$ and we only modified vertices of U , we still have $Z \subseteq V(\widehat{D})$. Since each connected component of graph $\widehat{D} \setminus Z$ consist either of the internal vertices of a modified gadget H_{X_e} or of a single path resulting from a division of an internal vertex of the introduced path or of a vertex v , the pathwidth of $\widehat{D} \setminus Z$ is 4. Hence, the distance of \widehat{D} to pathwidth 4 is $O(k)$, where k is the number of vertices of D . This implies part **iv)** of the theorem.

Moreover, apart from vertices in Z only the vertices \hat{v} are adjacent to internal vertices of (modified) gadgets H_s . Hence, we let \widehat{Z} be the set Z together with the vertex \hat{v} from every gadget H_{X_e} introduced. We again have that $|\widehat{Z}| = O(k^2)$. Furthermore, each connected component of graph $\widehat{D} \setminus \widehat{Z}$ consist either of the internal vertices of a modified gadget H_x or of a single path obtained from dividing the internal vertex of a path or of vertex v . Thus, the pathwidth of $\widehat{D} \setminus \widehat{Z}$ is 3. Therefore, the distance of \widehat{D} to pathwidth 3 is $O(k^2)$. This implies parts **b)** and **iii)** of the theorem. ◀

H Alternative Proof of W[1]-hardness with Respect to Distance to Constant Pathwidth

In this section we present an independent alternative proof of the W[1]-hardness of the DIRECTED TRAVELING SALESMAN PROBLEM w.r.t. distance to pathwidth. Even though

this proof may be somewhat redundant in view of [52], we still chose to include it here as it presents a different approach towards this result, which may be of interest to the potential reader.

► **Theorem 29.** *DIRECTED TRAVELING SALESMAN PROBLEM is $W[1]$ -hard with respect to distance to pathwidth 19 even if all weights are 1.*

Gutin, Jones, and Wahlström [38] have shown hardness of an auxiliary intermediate problem called PROPERLY BALANCED SUBGRAPH. We use the same problem to prove Theorem 29. We first introduce this problem.

A directed multigraph is called *balanced* if the in-degree of each vertex equals its out-degree.

PROPERLY BALANCED SUBGRAPH (PBS)

Input: A directed multigraph $D = (V, A)$, a weight function $w : A \rightarrow \mathbb{Z}$, a set $X = \{(a_1, a'_1), \dots, (a_r, a'_r)\}$ of disjoint pairs of arcs (called *double arcs*), such that $a_i, a'_i \in A$ and a_i, a'_i have the same start vertex and end vertex, for each $(a_i, a'_i) \in X$; a set $Y = \{(b_1, b'_1), \dots, (b_s, b'_s)\}$ of disjoint pairs of arcs (called *forbidden pairs*), such that $b_i, b'_i \in A$ and b_i is the reverse of b'_i , for each $(b_i, b'_i) \in Y$. Double arcs are disjoint from forbidden pairs.

Question: Is there a balanced subgraph D' of D of negative weight such that $|A(D') \cap \{a_i, a'_i\}| \neq 1$ for each $(a_i, a'_i) \in X$, and $|A(D') \cap \{b_i, b'_i\}| \neq 2$ for each $(b_i, b'_i) \in Y$?

Note that in this problem we are considering subgraphs, i.e., each arc can be contained in D' at most as many times as it is contained in D .

► **Proposition 30** (Gutin et al. [38, Theorem 2]). *PBS is $W[1]$ -hard parameterized by distance to pathwidth 16, even under the following restrictions:*

- *There exists a single arc a^* of weight -1 ;*
- *a^* is not part of a double arc;*
- *All other arcs have weight 0;*
- *If there are at least two arcs between a pair of vertices, then they are exactly two and it is a double arc (in particular, there are no forbidden pairs).*

Gutin et al. prove the proposition by a reduction from MULTICOLORED CLIQUE. Although they formulate it for the parameterization by pathwidth, it is easy to observe (cf., “ D^* has pathwidth at most 16” [38, Page 10–11]) that, in fact, the constructed graph is 2ℓ vertices from being of pathwidth 16, where ℓ is the number of edges of the clique to be found. Also, it is easy to verify, that the reduction would also work starting from PARTITIONED SUBGRAPH ISOMORPHISM, giving an ETH lowerbound.

As a next step, Gutin et al. [38] provided a reduction from PBS (of the special form guaranteed by Proposition 30) to the following problem, only increasing the pathwidth bound by 1.

WEIGHTED MIXED CHINESE POSTMAN PROBLEM (WMCPP)

Input: A mixed multigraph $H = (V, A, E)$, where A are arcs and E are edges, a weight function $w : A \cup E \rightarrow \mathbb{N}$, a budget $b \in \mathbb{N}$

Question: Is there a closed walk in H of total weight at most b which traverses each arc and each edge at least once?

Since Gutin et al. do not prove the properties of this reduction separately but they include it as a part of a larger reduction, we repeat this part of the reduction here with some small modifications and simplifications for completeness.


► **Lemma 31.** *WMCPP is $W[1]$ -hard parameterized by distance to pathwidth 17, even if the resulting instance $(H = (V, A, E), w, b)$ satisfies the following restrictions:*

- Each weight is either 2, 3, or M , where $M \leq 3|A| + 12$;
- weight of each edge is M ;
- $b < M \cdot |w^{-1}(M)| + M - 2$.

Proof of Lemma 31. We provide a reduction from PBS with the special properties as guaranteed by Proposition 30. We follow the approach of Gutin et al. [38] with slight simplification and modification. Let (D, w, X, \emptyset) be an instance of PBS with the special properties as guaranteed by Proposition 30. Since there is only one weakly connected component containing the special arc of weight -1 , we may ignore any parts of the solution in any other component. I.e., we assume that D is weakly connected.

We replace each arc, or double arc by a gadget which allows either a *passive solution* (all arc and edges are traversed, but no imbalance is caused) or an *active solution* (all vertices are traversed and the only disbalanced vertices are the original vertices—this corresponds to using the arc or the double arc). The weights of these two solutions are the same for all gadgets except for the one replacing the special arc of negative weight, where the active solution is cheaper.

We present the gadgets in terms of normal arcs (weight 2), heavy arcs (weight M), and heavy edges (weight M). We also introduce one arc of weight 3. The value M will be set later. Due to the condition on b any solution traverses each heavy arc and each heavy edge exactly once. The gadgets are only connected by sharing the vertices of the original graph, i.e., all newly introduced vertices only have connections within their gadget. Let us now present the gadgets.



■ **Figure 14** Replacing a simple arc (u, v) of weight 0. Left to right: original arc, the replacement gadget, passive solution, active solution.


Arc (u, v) of weight 0 that is not part of a double arc: To construct the gadget, we remove the original arc, add a new vertex z_{uv} , add normal arcs (z_{uv}, u) and (z_{uv}, v) , two heavy arcs (u, z_{uv}) , and a heavy arc (v, z_{uv}) (see Figure 14).

In any solution that traverses each heavy arc exactly once we have three incoming arcs to z_{uv} so we have to use one of the outgoing normal arcs once and the other one twice to make z_{uv} balanced.

- In the *passive solution* we use (z_{uv}, u) twice, which makes all vertices of the gadget balanced.
- In the *active solution* we use (z_{uv}, v) twice. Then u has one more outgoing arc than incoming, while v has one more incoming than outgoing—which exactly corresponds to using the arc (u, v) .

Both solutions use 3 normal arcs and 3 heavy ones, i.e., they have the same weight $3M + 6$.

The arc (u, v) of weight -1 : We use the above gadget, but make the arc (z_{uv}, u) of weight 3. Then the weight of the active solution is $3M + 7$, whereas the weight of the passive one is $3M + 8$.



■ **Figure 15** Replacing a double arc (u, v) . Left to right: original arcs, the replacement gadget, passive solution, active solution.

Double arc from u to v : We remove the original arcs and introduce a heavy arc (u, v) and a heavy edge $\{u, v\}$ (see Figure 15). For a solution that traverses each of them exactly once, we only need to know, which direction the edge is traversed.

- In the *passive solution* we traverse it from v to u which makes both vertices balanced.
- In the *active solution* we traverse it from u to v , which creates two outgoing arcs from u and two incoming arcs to v , corresponding to using the double arc.

Both solutions use one heavy arc and one heavy edge, i.e., they have the same weight $2M$.

We replace each arc and double arc by the appropriate gadget and call the obtained mixed multigraph H and the corresponding weight function \hat{w} . Suppose that D had m_1 simple arcs of weight 0, m_2 double arcs, and the one arc of weight -1 . Then, taking the passive solution for each gadget, we obtain a solution of total weight

$$m_1 \cdot (3M + 6) + m_2 \cdot 2M + 3M + 8 = M \cdot (3m_1 + 2m_2 + 3) + 6m_1 + 8.$$

We let $M = 6m_1 + 12$ and $b = M \cdot (3m_1 + 2m_2 + 3) + 6m_1 + 7$. Note that there are $|\hat{w}^{-1}(M)| = 3m_1 + 2m_2 + 3$ heavy arcs and edges in H . Therefore, we have

$$b = M \cdot (3m_1 + 2m_2 + 3) + 6m_1 + 7 < M \cdot (3m_1 + 2m_2 + 3) + 6m_1 + 12 - 2 = M \cdot |\hat{w}^{-1}(M)| + M - 2.$$

This implies that every solution of weight at most b traverses each heavy arc and each heavy edge exactly once.

Let $Z \subseteq V(D)$ be such that $D \setminus Z$ has pathwidth at most 16. We claim, that the pathwidth of $H \setminus Z$ is at most 17. Indeed, given a path decomposition for $D \setminus Z$, to obtain a path decomposition of $H \setminus Z$, we only need to place the new vertices z_{uv} into the decomposition. To do so, for each arc (u, v) we find a bag of the original decomposition which contains $\{u, v\} \setminus Z$, make a copy of this bag right after the original bag and place z_{uv} into this new copy. This way every arc and every edge appears in a bag and the width is only increased by 1.

As the reduction can be clearly carried out in polynomial time, it remains to verify its correctness. Let D' be a balanced subgraph of D of negative weight such that $|A(D') \cap \{a_i, a'_i\}| \neq 1$ for each $(a_i, a'_i) \in X$. We construct a solution S for the constructed instance (H, w, b) of WMCPP by taking the passive solution in each gadget replacing an arc or double arc not being part of D' and the active solution for each arc or double arc included in D' . Obviously, S contains each edge and each arc at least once. Since the disbalance caused by S at each original vertex of D is exactly the same as the disbalance caused by the arcs of D' , S is a balanced subgraph. Since D was weakly connected and S traverses all arcs and edges of H , S is weakly connected. Therefore, it corresponds to a walk traversing each arc and

each edge at least once. Moreover, since D' is of negative weight, $a^* \in A(D')$ and the active solution is used in the gadget replacing arc a^* . Thus, the total weight of the solution is

$$m_1 \cdot (3M + 6) + m_2 \cdot 2M + 3M + 7 = M \cdot (3m_1 + 2m_2 + 3) + 6m_1 + 7 = b.$$

Hence, (H, w, b) is a yes-instance of WMCPP.

Conversely, if there is a solution S for (H, w, b) , then, due to the budget constraint, it traverses every heavy edge and every heavy arc exactly once. As vertices z_{uv} must be balanced with respect to S , it follows, that S either takes the passive or the active solution in each gadget. In particular, it must take the active one in the gadget replacing a^* , as otherwise it would exceed the budget. Now taking D' as the subgraph of D formed by the arcs and double arcs whose gadgets use the active solution in S , we obtain a balanced subgraph of D , since the imbalance for each vertex is the same in S as in D' . The subgraph D' also respects double arcs by construction. As $a^* \in D'$, D' has a negative weight. Therefore, (D, w, X, \emptyset) is a yes-instance, finishing the proof. \blacktriangleleft

Proof of Theorem 29. We provide a parameterized reduction from WMCPP parameterized by distance to pathwidth 17. Let (H, w, b) be an instance with the special properties guaranteed by Lemma 31.

To produce an instance of DTSP, we replace each arc of weight h by a directed path of length h . We replace a heavy edge of weight M by a bidirected path of length M , that is, we take a directed path of length M from one endpoint to the other and then add the opposite arc to every arc of the path. Note that there are no edges of weight other than M in H by assumption. We denote the resulting directed graph G .

The reduction can be clearly carried out in polynomial time, as $M \leq 3|A| + 12$.

Let $Z \subseteq V(H)$ be such that $H \setminus Z$ has pathwidth at most 17. Given a path decomposition for $H \setminus Z$ of width 17, we can obtain a path decomposition of $G \setminus Z$ of width at most 19 in the following way. For each arc or edge (u, v) of H of weight h find a bag in which $\{u, v\} \setminus Z$ appears. If $h = 2$, then introduce a copy of the bag right after the original bag, and add the inner vertex of the path replacing the arc (u, v) into that copy. Otherwise introduce $h - 2$ copies of the bag, subsequently containing adjacent pairs of inner vertices of the replacement path. It is easy to verify that we indeed obtain a path decomposition of $G \setminus Z$. Therefore, the pathwidth of $G \setminus Z$ is at most 2 more than the pathwidth of $H \setminus Z$ and the reduction preserves the parameter.

Finally, we claim that the instance (H, w, b) of WMCPP is equivalent to the instance $(G, \mathbf{1}, b)$ of DTSP, where $\mathbf{1}$ assigns weight 1 to every arc of G . If S is a solution for H , then replacing each traversal of an edge or an arc by the traversal of the corresponding path in G we obtain a walk of the same weight that visits every vertex at least once. This is the case for the vertices of H , since each of them is incident with at least one (traversed) arc or edge. For the inner vertices of the paths it follows from the fact that each path is traversed at least once (in at least one of the directions).

Conversely, suppose that we have a walk S' in G of weight at most b that visits all vertices of G . First, to visit an inner vertex of a directed path, the walk has to traverse the whole path at least once. Now, consider a bidirected path of length M . As each inner vertex has to be visited, S' contains for each such vertex w at least one arc with head in w and one arc with tail in w . To reach also other vertices, the walk must also contain an arc of the path with tail in one of the endpoints and an arc of the path with head in one of the endpoints. It follows that altogether S' contains at least M arcs from each such path. If neither of the two arcs between two consecutive vertices of the path would be contained in S' , then both arcs would have to be contained for all other pairs. Therefore, S' would contain

at least $2M - 2$ arcs from that path, i.e., $M - 2$ more than M , which cannot happen as $b < M \cdot |w^{-1}(M)| + M - 2$. Hence, S' traverses the path in at least one direction completely. As it cannot traverse it completely twice, we can omit any other visits to the path.

Now consider a walk S in H obtained from S' by replacing each traversal of a path by a traversal of the corresponding arc or edge of H . By the above argument S visit every edge and every arc of H and the weight of S is at most that of S' , which is at most b . Hence, (H, w, b) is a yes-instance, finishing the proof. ◀