

EGERVÁRY RESEARCH GROUP
ON COMBINATORIAL OPTIMIZATION



TECHNICAL REPORTS

TR-2020-19. Published by the Egerváry Research Group, Pázmány P. sétány 1/C,
H-1117, Budapest, Hungary. Web site: www.cs.elte.hu/egres. ISSN 1587-4451.

A $3/2$ -Approximation for the Metric Many-visits Path TSP

Kristóf Bérczi, Matthias Mnich, and Roland Vincze

October 2020

A $3/2$ -Approximation for the Metric Many-visits Path TSP¹

Kristóf Bérczi², Matthias Mnich³, and Roland Vincze⁴

Abstract

In the MANY-VISITS PATH TSP, we are given a set of n cities along with their pairwise distances (or cost) $c(uv)$, and moreover each city v comes with an associated positive integer request $r(v)$. The goal is to find a minimum-cost path, starting at city s and ending at city t , that visits each city v exactly $r(v)$ times.

We present a $3/2$ -approximation algorithm for the metric MANY-VISITS PATH TSP, that runs in time polynomial in n and *poly-logarithmic* in the requests $r(v)$. Our algorithm can be seen as a far-reaching generalization of the $3/2$ -approximation algorithm for PATH TSP by Zenklusen (SODA 2019), which answered a long-standing open problem by providing an efficient algorithm which matches the approximation guarantee of Christofides' algorithm from 1976 for metric TSP.

One of the key components of our approach is a polynomial-time algorithm to compute a connected, degree bounded multigraph of minimum cost. We tackle this problem by generalizing a fundamental result of Király, Lau and Singh (Combinatorica, 2012) on the MINIMUM BOUNDED DEGREE MATROID BASIS problem, and devise such an algorithm for general polymatroids, even allowing element multiplicities.

Our result directly yields a $3/2$ -approximation to the metric MANY-VISITS TSP, as well as a $3/2$ -approximation for the problem of scheduling classes of jobs with sequence-dependent setup times on a single machine so as to minimize the makespan.

Keywords: Traveling salesman problem, Degree constraints, Generalized polymatroids

¹Supported by DAAD with funds of the Bundesministerium für Bildung und Forschung (BMBF) and by DFG project MN 59/4-1.

²MTA-ELTE Egerváry Research Group, Department of Operations Research, Eötvös Loránd University, Budapest, Hungary. berkri@cs.elte.hu.

³TU Hamburg, Hamburg, Germany. matthias.mnich@tuhh.de.

⁴TU Hamburg, Hamburg, Germany. roland.vincze@tuhh.de.

1 Introduction

The traveling salesman problem (TSP) is one of the cornerstones of combinatorial optimization. Given a set V of n cities with non-negative costs $c(uv)$ for each cities u and v , the objective is to find a minimum cost closed walk visiting each city. TSP is well-known to be NP-hard even in the case of metric costs, i.e. when the cost function c satisfies the triangle inequality. For metric costs, the best known approximation ratio that can be obtained in polynomial time is $3/2$, discovered independently by Christofides [9] and Serdyukov [50].

In the traveling salesman path problem, or PATH TSP, two distinguished vertices s and t are given, and the goal is to find a minimum cost walk from s to t visiting each city. Approximating the metric PATH TSP has a long history, from the first $5/3$ -approximation by Hoogeveen [27], through subsequent improvements [2, 22, 48, 49, 57] to the recent breakthroughs. The latest results eventually closed the gap between the metric TSP and the metric PATH TSP: Traub and Vygen [53] provided a $(3/2 + \varepsilon)$ -approximation for any $\varepsilon > 0$, Zenklusen [58] provided a $3/2$ -approximation and finally the three authors showed a reduction from the PATH TSP to the TSP [55].

We consider a far-reaching generalization of the metric PATH TSP, the metric MANY-VISITS PATH TSP, where in addition to the costs c on the edges, a requirement $r(v)$ is given for each city v . The aim is to find a minimum cost walk from s to t that visits each city v exactly $r(v)$ times. The cycle version of this problem, where $s = t$, is known as MANY-VISITS TSP and was first considered in 1966 by Rothkopf [45]. Psaraftis [44] proposed a dynamic programming approach that solves the problem in time $(r/n)^n$ for $r = \sum_{v \in V} r(v)$. Later, Cosmadakis and Papadimitriou [10] gave the first algorithm for MANY-VISITS TSP with logarithmic dependence on r , though the space and time requirements of their algorithm were still superexponential in n . Recently, Berger et al. [5] simultaneously improved the run time to $2^{\mathcal{O}(n)} \cdot \log r$ and reduced the space complexity to polynomial. (The algorithm by Berger et al. [5] can be slightly modified to solve the path version as well.) Lately, Kowalik et al. [37] made further fine-grained time complexity improvements. To the best of our knowledge, no constant-factor approximation algorithms for the metric MANY-VISITS TSP¹ or metric MANY-VISITS PATH TSP are currently known.

Besides being of scientific interest in itself, the MANY-VISITS PATH TSP can be used for modeling various problems. The *aircraft sequencing problem* or *aircraft landing problem* is one of the most referred applications in the literature [3, 6, 40, 44], where the goal is to find a schedule of departing and/or landing airplanes that minimizes an objective function and satisfies certain constraints. The aircraft are categorized into a small number of classes, and for each pair of classes a non-negative lower bound is given denoting the minimum amount of time needed to pass between the take off/landing of two planes from the given classes. The problem can be embedded in the MANY-VISITS PATH TSP model by considering the classes to be cities and the

¹At the Hausdorff Workshop on Combinatorial Optimization in 2018, Rico Zenklusen brought up the topic of approximation algorithms for the metric version of MANY-VISITS TSP in the context of iterative relaxation techniques; he suggested an approach to obtain a 1.5-approximation, which is unpublished.

separation times to be costs between them, while the number of airplanes in a class corresponds to the number of visits of a city.

As another illustrious example, the MANY-VISITS PATH TSP is equivalent to the high-multiplicity job scheduling problem $1|HM, s_{ij}, p_j| \sum C_j$, where each class j of jobs has a processing time p_j and there is a setup time s_{ij} between processing two jobs of different classes. There is only a handful of constant-factor approximation algorithms for scheduling problems with setup times [1], see for example the results of Jansen et al. [31] or Deppert and Jansen [11] that consider sequence-independent batch setup times, or van der Veen et al. [56] that considers sequence-dependent setup times with a special structure. An approximation algorithm for the MANY-VISITS PATH TSP would further extend the list of such results.

A different kind of application comes from geometric approximation. Recently, Kozma and Mömke provided an EPTAS for the MAXIMUM SCATTER TSP [38]. Their approach involved grouping certain input points together and thus reducing the input size. The reduced problem is exactly the MANY-VISITS TSP. The same problem arises as a subproblem in the fixed-parameter algorithm for the HAMILTONIAN CYCLE problem on graphs with bounded neighborhood diversity [39].

Our work relies on a polymatroidal optimization problem with degree constraints. An illustrious example of such problem is the MINIMUM BOUNDED DEGREE SPANNING TREE problem, where the goal is to find a minimum cost spanning tree in a graph with lower and upper bounds on the degree of each vertex. Checking feasibility of a degree-bounded spanning tree contains the NP-hard HAMILTONIAN PATH problem, and several algorithms were given that were balancing between the cost of the spanning tree and the violation of the degree bounds [7, 8, 19, 21, 35, 36]. Based on an iterative rounding approach [30] combined with a relaxation step, Singh and Lau [51] provided a polynomial-time algorithm that finds a spanning tree of cost at most the optimum value violating each degree bound by at most 1. Király et al. [33] later showed that similar results can be obtained for the more general MINIMUM BOUNDED DEGREE MATROID BASIS PROBLEM.

Our results

In this paper we provide the first efficient constant-factor approximation algorithm for the metric MANY-VISITS PATH TSP. Formally, a graph $G = (V, E)$ is given with a positive integer $r(v)$ for each $v \in V$, and a non-negative cost $c(uv)$ for every pair of vertices u, v ; finally, a departure city s and an arrival city t are specified. We seek a minimum cost s - t -walk that visits each city v exactly $r(v)$ times, where leaving city s as well as arriving to city t counts as one visit.

The cost function $c : E \rightarrow \mathbb{R}_{\geq 0}$ is assumed to be metric. Besides the triangle inequality $c(uw) \leq c(uv) + c(vw)$ for every triplet u, v, w this implies that the cost of a self-loop $c(vv)$ at vertex v is at most the cost of leaving city v to any other city u and returning, that is:

$$c(vv) \leq 2 \cdot \min_{u \in V-v} c(uv) \quad \text{for all } v \in V .$$

The assumption of metric costs is necessary, as the TSP, and therefore the MANY-VISITS TSP, does not admit any non-trivial approximation for unrestricted cost functions assuming that $P \neq NP$ (see e.g. Theorem 6.13 in the book of Garey and Johnson [20]).

We start with a simple approximation idea, that leads to a constant factor approximation in strongly polynomial time:

Theorem 1. *There is a polynomial-time $5/2$ -approximation for the metric MANY-VISITS PATH TSP, that runs in time polynomial in n and $\log r$.*

The approximation factor $5/2$ in Theorem 1 still leaves a gap to the best-known factor $3/2$ for the metric PATH TSP, which is due to Zenklusen [58]. His recent $3/2$ -approximation for the metric PATH TSP uses a Christofides-Serdyukov-like construction that combines a spanning tree and a matching, with the key difference that it calculates a constrained spanning tree in order to bound the costs of the tree and the matching by $3/2$ times the optimal value.

Our main algorithmic result matches this approximation ratio for the metric MANY-VISITS PATH TSP.

Theorem 2. *There is a polynomial-time $3/2$ -approximation for the metric MANY-VISITS PATH TSP. The algorithm runs in time polynomial in n and $\log r$.*

As a direct consequence of Theorem 2, we obtain the following:

Corollary 3. *There is a $3/2$ -approximation for the metric MANY-VISITS TSP that runs in time polynomial in n and $\log r$.*

Our approach follows the main steps of Zenklusen's work [58]. However, the presence of requests $r(v)$ makes the problem significantly more difficult and several new ideas are needed to design an algorithm which returns a tour with the correct number of visits and still runs in polynomial time. For instance, whereas the backbone of both Christofides and Zenklusen's algorithm is a spanning tree (with certain properties), the possibly exponentially large number of (parallel) edges in a many-visits TSP solution requires us to work with a structure that is more general than spanning trees. We therefore consider the problem of finding a minimum cost connected multigraph with lower bounds ρ on the degree of vertices, and lower and upper bounds L and U , respectively, on the number of occurrences of the edges. We call this task the MINIMUM BOUNDED DEGREE CONNECTED MULTIGRAPH WITH EDGE BOUNDS problem, and show the following:

Theorem 4. *There is an algorithm for the MINIMUM BOUNDED DEGREE CONNECTED MULTIGRAPH WITH EDGE BOUNDS problem that, in time polynomial in n and $\log \sum_v \rho(v)$, returns a connected multigraph T with $\rho(V)/2$ edges, where each vertex v has degree at least $\rho(v) - 1$ and the cost of T is at most the cost of $\min\{c^T x \mid x \in P^{CG}(\rho, L, U)\}$, where*

$$(1) \quad P^{CG}(\rho, L, U) := \left\{ x \in \mathbb{R}_{\geq 0}^E \left| \begin{array}{ll} \text{supp}(x) \text{ is connected} & \\ x(E) = \sum_v \rho(v)/2 & \\ x(\delta(v)) \geq \rho(v) & \forall v \in V \\ L(vw) \leq x(vw) \leq U(vw) & \forall v, w \in V \end{array} \right. \right\}.$$

Note that an optimal solution x^* to the MINIMUM BOUNDED DEGREE CONNECTED MULTIGRAPH WITH EDGE BOUNDS problem is a minimum cost integral vector from the polytope P^{CG} . We use the result of [Theorem 4](#) to obtain a multigraph that serves a key role in our approximation algorithm for the metric MANY-VISITS PATH TSP; the values ρ , L and U depend on the instance and the details are given in [Section 4](#).

The MINIMUM BOUNDED DEGREE CONNECTED MULTIGRAPH WITH EDGE BOUNDS problem shows a lot of similarities to the MINIMUM BOUNDED DEGREE SPANNING TREE problem. However, neither the result of Singh and Lau [[51](#)] nor the more general approach by Király et al. [[33](#)] applies to our setting, due to the presence of parallel edges and self-loops in a multigraph.

One of our key contributions is therefore an extension of the result of Király et al. [[33](#)] to generalized polymatroids, which might be of independent combinatorial interest. Formally, the BOUNDED DEGREE G-POLYMATROID ELEMENT WITH MULTIPLICITIES problem takes as input a g-polymatroid $Q(p, b)$ defined by a paramodular pair $p, b : 2^S \rightarrow \mathbb{R}$, a cost function $c : S \rightarrow \mathbb{R}$, a hypergraph $H = (S, \mathcal{E})$ with lower and upper bounds $f, g : \mathcal{E} \rightarrow \mathbb{Z}_{\geq 0}$ and multiplicity vectors $m_\varepsilon : S \rightarrow \mathbb{Z}_{\geq 0}$ for $\varepsilon \in \mathcal{E}$ satisfying $m_\varepsilon(s) = 0$ for $s \in S - \varepsilon$. The objective is to find a minimum-cost integral element x of $Q(p, b)$ such that $f(\varepsilon) \leq \sum_{s \in \varepsilon} m_\varepsilon(s)x(s) \leq g(\varepsilon)$ for each $\varepsilon \in \mathcal{E}$. We give a polynomial-time algorithm for finding a solution of cost at most the optimum value with bounds on the violations of the degree prescriptions.

Theorem 5. *There is an algorithm for the BOUNDED DEGREE G-POLYMATROID ELEMENT WITH MULTIPLICITIES problem which returns an integral element x of $Q(p, b)$ of cost at most the optimum value such that $f(\varepsilon) - 2\Delta + 1 \leq \sum_{s \in \varepsilon} m_\varepsilon(s)x(s) \leq g(\varepsilon) + 2\Delta - 1$ for each $\varepsilon \in \mathcal{E}$, where $\Delta = \max_{s \in S} \{\sum_{\varepsilon \in \mathcal{E}: s \in \varepsilon} m_\varepsilon(s)\}$. The run time of the algorithm is polynomial in n and $\log \sum_\varepsilon (f(\varepsilon) + g(\varepsilon))$.*

When only lower bounds (or only upper bounds) are present, we call the problem LOWER (UPPER) BOUNDED DEGREE G-POLYMATROID ELEMENT WITH MULTIPLICITIES. Similarly to Király et al. [[33](#)], we obtain an improved bound on the degree violations when only lower or upper bounds are present: ²

Theorem 6. *There is an algorithm for LOWER BOUNDED DEGREE G-POLYMATROID ELEMENT WITH MULTIPLICITIES which returns an integral element x of $Q(p, b)$ of cost at most the optimum value such that $f(\varepsilon) - \Delta + 1 \leq \sum_{s \in \varepsilon} m_\varepsilon(s)x(s)$ for each $\varepsilon \in \mathcal{E}$. An analogous result holds for UPPER BOUNDED DEGREE G-POLYMATROID ELEMENT, where $\sum_{s \in \varepsilon} m_\varepsilon(s)x(s) \leq g(\varepsilon) + \Delta - 1$. The run time of these algorithms is polynomial in n and $\log \sum_\varepsilon f(\varepsilon)$ or $\log \sum_\varepsilon g(\varepsilon)$, respectively.*

²The results in [Theorem 5](#), [Theorem 6](#) and [Corollary 3](#) appeared in an unpublished work [[4](#)] by a superset of the authors. In order to make the paper self-contained, we include all the details and proofs in this paper as well.

2 Preliminaries

Basic notation. Throughout the paper, we let $G = (V, E)$ be a finite, undirected complete graph on n vertices, whose edge set E also contains a self-loop at every vertex $v \in V$. For a subset $F \subseteq E$ of edges, the *set of vertices covered by F* is denoted by $V(F)$. The *number of connected components* of the graph $(V(F), F)$ is denoted by $\text{comp}(F)$. For a subset $X \subseteq V$ of vertices, the *set of edges spanned by X* is denoted by $E(X)$. Given a multiset F of edges (that is, F might contain several copies of the same edge), the multiset of edges leaving the vertex set $C \subseteq V(F)$ is denoted by $\delta_F(C)$. Similarly, denote the multiset of regular edges (i.e. excluding self-loops) in F incident to a vertex $v \in V$ is denoted by $\delta_F(v)$. Denote the multiset of all edges (i.e. including self-loops) in F incident to a vertex $v \in V$ by $\dot{\delta}_F(v)$, then the *degree* of v in F is denoted by $\deg_F(v) := |\dot{\delta}_F(v)|$, where every copy of the self-loop at v in F is counted twice. We will omit the subscript when F contains all the edges of G , that is, $F = E$. For a vector $x \in \mathbb{R}^E$, we denote the sum of the x -values on the edges incident to v by $x(\dot{\delta}(v))$. Note that the x -value of the self-loop at v is counted twice in $x(\dot{\delta}(v))$. Let us denote the set of edges between two disjoint vertex sets A and B by $\delta(A, B)$. Given two graphs or multigraphs H_1, H_2 on the same vertex set, $H_1 + H_2$ denotes the multigraph obtained by taking the union of the edge sets of H_1 and H_2 .

Given a vector $x \in \mathbb{R}^S$ and a set $Z \subseteq S$, we use $x(Z) = \sum_{s \in Z} x(s)$. The *lower integer part* of x is denoted by $\lfloor x \rfloor$, so $\lfloor x \rfloor(s) = \lfloor x(s) \rfloor$ for every $s \in S$. This notation extends to sets, so by $\lfloor x \rfloor(Z)$ we mean $\sum_{s \in Z} \lfloor x \rfloor(s)$. The *support* of x is denoted by $\text{supp}(x)$, that is, $\text{supp}(x) = \{s \in S \mid x(s) \neq 0\}$. The *difference of set B from set A* is denoted by $A - B = \{s \in A \mid s \notin B\}$. We denote a single-element set $\{s\}$ by s , and with a slight abuse of notation, we write $A - s$ to indicate $A - \{s\}$. Let us denote the *symmetric difference* of two sets A and B by $A \Delta B := (A - B) \cup (B - A)$ and the *characteristic vector* of a set A by χ_A .

For a collection \mathcal{T} of subsets of S , we call $\mathcal{L} \subseteq \mathcal{T}$ an *independent laminar system* if for any pair $X, Y \in \mathcal{L}$: (i) they do not properly intersect, i.e. either $X \subseteq Y$, $Y \subseteq X$ or $X \cap Y = \emptyset$, and (ii) the characteristic vectors χ_Z of the sets $Z \in \mathcal{L}$ are independent over the real numbers. A *maximal* independent laminar system \mathcal{L} with respect to \mathcal{T} is an independent laminar system in \mathcal{T} such that for any $Y \in \mathcal{T} - \mathcal{L}$ the system $\mathcal{L} \cup \{Y\}$ is not independent laminar. In other words, if we include any set Y from $\mathcal{T} - \mathcal{L}$, it will intersect at least one set Z from \mathcal{L} , or χ_Y can be given as a linear combination of $\{\chi_Z \mid Z \in \mathcal{L}\}$. Given a laminar system \mathcal{L} and a set $X \subseteq S$, the set of maximal members of \mathcal{L} lying inside X is denoted by $\mathcal{L}^{\max}(X)$, that is, $\mathcal{L}^{\max}(X) = \{Y \in \mathcal{L} \mid Y \subseteq X, \nexists Y' \in \mathcal{L} \text{ s.t. } Y \subsetneq Y' \subseteq X\}$.

Many-visits Path TSP. Recall that in the MANY-VISITS PATH TSP, we seek for a minimum cost s - t -walk P such that P visits each vertex $v \in V$ exactly $r(v)$ times. Let $r(V) = \sum_{v \in V} r(v)$. The sequence of the edges of P has length $r(V) - 1$, which is exponential in the size of the input, as the values $r(v)$ are stored using $\log r(V)$ space. For this reason, instead of explicitly listing the edges in a walk (or tour) we always consider *compact representations* of the solution and the multigraphs that arise in our

algorithms. That is, rather than storing an $(r(V) - 1)$ -long sequence of edges, for every edge e we store its multiplicity $z(e)$ in the solution. As there are at most n^2 different edges in the solution each having multiplicity at most $\max_{v \in V} r(v)$, the space needed to store a feasible solution is $\mathcal{O}(n^2 \log r(V))$. Therefore, a vector $z \in \mathbb{Z}_{\geq 0}^E$ represents a feasible tour if $\text{supp}(z)$ is a connected subgraph of G and $\deg_z(v) = 2 \cdot r(v)$ holds for all $v \in V - \{s, t\}$ and $\deg_z(v) = 2 \cdot r(v) - 1$ for $v \in \{s, t\}$. (Note that each self-loop vv contributes 2 in the value $\deg(v) = |\delta(v)|$.)

Denote by $\mathbf{P}_{c,r,s,t}^*$ an optimal solution for an instance (G, c, r, s, t) of the MANY-VISITS PATH TSP. Let us denote by $\mathbf{P}_{c,1,s,t}^*$ an optimal solution for the single-visit counterpart of the problem, i.e. when $r(v) = 1$ for each $v \in V$. Relaxing the connectivity requirement for solutions of the MANY-VISITS PATH TSP yields Hitchcock's transportation problem [26], where supply and demand vertices $\{a_v\}_{v \in V}$ and $\{b_v\}_{v \in V}$ are given. The supplies for $v \in V - s$ are then defined by $r(v)$, the supply of s by $r(s) - 1$; the demand of each vertex $v \in V - t$ by $r(v)$ and the demand of t by $r(t) - 1$. Finally, by setting the transportation costs between a_u and b_v as $c(uv)$, the objective is to fulfill the supply and demand requirements by transporting goods from vertices $\{a_v\}_{v \in V}$ to vertices in $\{b_v\}_{v \in V}$, while keeping the total cost minimal. The transportation problem is solvable in polynomial time using a minimum cost flow algorithm [13] and we denote an optimal solution by $\mathbf{TP}_{c,r,s,t}^*$, where s and t denote the special vertices with decreased supply and demand value, respectively.

Lemma 7. *Let $\mathbf{TP}_{c,r,s,t}^*$ be an optimal solution to the Hitchcock transportation problem, where $\text{supply}(v) + \text{demand}(v)$ is odd for $v \in \{s, t\}$ and it is even otherwise. Then $\mathbf{TP}_{c,r,s,t}^*$ can be decomposed into cycles and exactly one s - t -path.*

Proof. Any solution X to the transportation problem is essentially a multigraph that has an even degree for vertices $v \in V - \{s, t\}$, and an odd degree for $v \in \{s, t\}$. Hence, because of a parity argument, there has to be an s - t -path U in X , possibly covering other vertices $W \subset V - \{s, t\}$. Vertices $w \in W$ have an even degree in U . Therefore, deleting the edges of U from X , all vertices $v \in V$ will have an even degree in the modified multigraph X' . Thus X' can be decomposed into a union of (not necessarily distinct) cycles, and the lemma follows. \square

The decomposition provided by the lemma is called a path-cycle representation. Such a representation can be stored as a path P_0 and a collection \mathcal{C} of pairs (C, μ_C) , where each C is a simple closed walk (cycle) and μ_C is the corresponding integer denoting the number of copies of C . Below we show that one can always calculate a path-cycle decomposition in polynomial time, and such a decomposition takes polynomial space.

Lemma 8. *Let $\mathbf{P}_{c,r,s,t}$ be a many-visits TSP path with endpoints s, t , and $\mathbf{TP}_{c,r,s,t}$ be a transportation problem solution with special vertices s, t . There is a path-cycle representation of $\mathbf{P}_{c,r,s,t}$ and $\mathbf{TP}_{c,r,s,t}$, both of which take space polynomial in n and $\log r(V)$, and can be computed in time polynomial in n and $\log r(V)$.*

Proof. We first show the proof for a many-visits TSP path $\mathbf{P}_{c,r,s,t}$. Let us first add an edge ts to $\mathbf{P}_{c,r,s,t}$, and denote the resulting multigraph by T . Observe that T is a many-visits TSP tour with the same number of visits, since it is connected and the degree of

every vertex v in T is $2 \cdot r(v)$. We can now use the procedure **ConvertToSequence** by Grigoriev and van de Klundert [23], which takes the edge multiplicities of T , denoted by $\{x_{uv}\}_{u,v \in V}$ as input, and outputs a collection \mathcal{C} of pairs (C, μ_C) . Here, C is a simple closed walk, and μ_C is the corresponding integer denoting the number of copies of the walk C in P . Lastly, choose an arbitrary cycle C , such that $ts \in C$, and transform one copy of C into a path as follows. Let $C_0 := C$, and remove the edge ts from C_0 , resulting in an s - t -path P_0 . Update $\mu_C := \mu_C - 1$. Now (P_0, \mathcal{C}) is a compact path-cycle representation of $\mathbf{P}_{c,r,s,t}$.

In every iteration, the procedure **ConvertToSequence** looks for a cycle C and removes each of its occurrences from $\{x_{uv}\}_{u,v \in V}$. The procedure stops when $\{x_{uv}\}_{u,v \in V}$ represents a graph without edges. This demonstrates that the input need not represent a connected graph in the first place, as the edge removals possibly make it disconnected during the process. Note that the only structural difference between $\mathbf{TP}_{c,r,s,t}$ and $\mathbf{P}_{c,r,s,t}$ is that the underlying multigraph of $\mathbf{TP}_{c,r,s,t}$ might be disconnected. This means that the procedure **ConvertToSequence** can be applied to obtain a compact path-cycle representation of $\mathbf{TP}_{c,r,s,t}$ the same way as in the case of $\mathbf{P}_{c,r,s,t}$.

Finally, the number of cycles C in \mathcal{C} can be bounded by $\mathcal{O}(n^2)$ (as removing all occurrences of a cycle C sets at least one variable x_{uv} to zero), and the algorithm has a time complexity of $\mathcal{O}(n^4)$ [23]. The edge insertion and deletion, and other graph operations during the process, can also be implemented efficiently. This concludes the proof. \square

From now on, we assume that the path-cycle decompositions appearing in this paper are stored in space polynomial in n and $\log r(V)$.

Let (P_0, \mathcal{C}) be a compact path-cycle representation of a many-visits TSP path $\mathbf{P}_{c,r,s,t}$. One can obtain the explicit order of the vertices from (P_0, \mathcal{C}) the following way: traverse the s - t -path P_0 , and whenever a vertex u is reached for the first time, traverse μ_C copies of every cycle C containing u . Note that while the size of \mathcal{C} is polynomial in n , the size of the explicit order of the vertices is exponential, hence the approaches presented in this paper consider symbolic rather than literal traversals of many-visits TSP paths and tours.

3 A Simple $5/2$ -Approximation for Metric Many-visits Path TSP

In this section we give a simple $5/2$ -approximation algorithm for the metric MANY-VISITS PATH TSP that runs in polynomial time. The algorithm is as follows:

Theorem 1. *There is a polynomial-time $5/2$ -approximation for the metric MANY-VISITS PATH TSP, that runs in time polynomial in n and $\log r$.*

Proof. The algorithm is presented as [Algorithm 1](#). Since $\mathbf{P}_{c,1,s,t}^\alpha$ is connected, and P contains all the edges of $\mathbf{P}_{c,1,s,t}^\alpha$, P is also connected. Let (P_0, \mathcal{C}) be the compact path-cycle decomposition of $\mathbf{TP}_{c,r,s,t}^*$. The graph P thus consists of $\mathbf{P}_{c,1,s,t}^\alpha$ and the cycles of \mathcal{C} . The edges of $\mathbf{P}_{c,1,s,t}^\alpha$ contribute a degree of 1 in case of s and t , and 2

Algorithm 1 A polynomial-time $(\alpha + 1)$ -approximation for metric MANY-VISITS PATH TSP.

Input: A complete undirected graph $G = (V, E)$, costs $c : E \rightarrow \mathbb{R}_{\geq 0}$ satisfying the triangle inequality, requests $r : V \rightarrow \mathbb{Z}_{\geq 1}$, distinct vertices $s, t \in V$.

Output: An s - t -path that visits each $v \in V$ exactly $r(v)$ times.

- 1: Calculate an α -approximate solution $P_{c,1,s,t}^\alpha$ for the single-visit metric PATH TSP instance $(G, c, 1, s, t)$.
 - 2: Calculate an optimal solution $TP_{c,r,s,t}^*$ for the corresponding transportation problem, together with a compact path-cycle decomposition (P_0, \mathcal{C}) , where \mathcal{C} is a collection of pairs (C, μ_C) .
 - 3: Let P be the union of $P_{c,1,s,t}^\alpha$ and μ_C copies of every cycle $C \in \mathcal{C}$.
 - 4: Do shortcuts in P and obtain a solution P' , such that P' visits every city v exactly $r(v)$ times (that is, $\deg_{P'}(v) = 2 \cdot r(v)$ for every vertex $v \in V - \{s, t\}$, and $\deg_{P'}(v) = 2 \cdot r(v) - 1$ otherwise).
 - 5: **return** P' .
-

for $v \in V - \{s, t\}$; the cycles of \mathcal{C} contribute degrees of $2 \cdot r(v) - 2$ for $v \in \{s, t\}$, and degrees of $2 \cdot r(v)$ or $2 \cdot r(v) - 2$ for $v \in V - \{s, t\}$. Let us denote the latter set by W , matching the notation in the proof of [Lemma 7](#). The total degree of v in P is:

$$\begin{aligned} 2 \cdot r(v) - 1 & \quad \text{for } v \in \{s, t\}, \\ 2 \cdot r(v) & \quad \text{for } v \in W, \text{ and} \\ 2 \cdot r(v) + 2 & \quad \text{for the remaining vertices in } V - (W \cup \{s, t\}). \end{aligned}$$

As a direct consequence of the degrees and connectivity, P is an open walk that starts in s , visits every vertex $v \in V$ either $r(v)$ or $r(v) + 1$ times, and ends in t . Since the edge costs are metric, we can use shortcuts at the vertices $w \in V - (W \cup \{s, t\})$ to reduce their degrees by 2. We describe the procedure below.

Shortcutting. At [Step 3](#), $(P_{c,1,s,t}^\alpha, \mathcal{C})$ denotes the compact path-cycle representation of P . Let us construct an auxiliary multigraph A on the vertex set V by taking the edges of $P_{c,1,s,t}^\alpha$ and each cycle C from \mathcal{C} exactly once. Note that parallel edges appear in A if and only if an edge appears in multiple distinct cycles, or in the path $P_{c,1,s,t}^\alpha$ and at least one cycle C . Due to the construction, s and t have odd degree, while every other vertex has an even degree in A , which means that there exist an Eulerian trail in A . Moreover, there are $\mathcal{O}(n^2)$ cycles [\[23\]](#), hence the total number of edges in A is $\mathcal{O}(n^3)$. Consequently, using Hierholzer's algorithm, we can compute an Eulerian trail η in A in $\mathcal{O}(n^3)$ time [\[15, 25\]](#). The trail η covers the edges of each cycle C once. Now an implicit order of the vertices in the many-visits TSP path P is the following. Traverse the vertices of the Eulerian trail η in order. Every time a vertex u appears the first time, traverse all cycles C that contain the vertex μ_C times. Denote this trail by η' . It is easy to see that the sequence η' is a sequence of vertices that uses the edges of $P_{c,1,s,t}^\alpha$ once and the edges of each cycle C exactly μ_C times, meaning this is a feasible sequence of the vertices in the path P . Moreover, the order itself takes polynomial space, as it is enough to store indices of $\mathcal{O}(n^3)$ vertices and $\mathcal{O}(n^2)$ cycles.

Denote the surplus of visits of a vertex $w \in W$ by $\gamma(w) := \deg_P(w)/2 - r(w)$. In [Step 4](#), we remove the last $\gamma(w)$ occurrences of every vertex $w \in W$ from P by doing shortcuts: if an occurrence of w is preceded by u and superseded by v in P , replace the edges uw and wv by uv in the sequence. This can be done by traversing the compact representation of η' backwards, and removing the vertex w from the last $\gamma(w)$ cycles $C_{r(w)-\gamma(w)+1}^{(w)}, \dots, C_{r(w)}^{(w)}$. As $\sum_w \gamma(w)$ can be bounded by $\mathcal{O}(n)$, this operation makes $\mathcal{O}(n)$ new cycles, keeping the space required by the new sequence of vertices and cycles polynomial. Moreover, since the edge costs are metric, making shortcuts the way described above cannot increase the total cost of the edges in P . Finally, using a similar argument as in the algorithm of Christofides, the shortcutting does not make the trail disconnected. The resulting graph is therefore an s - t -walk P' that visits every vertex v exactly $r(v)$ times, that is, a feasible solution for the instance (G, c, r, s, t) .

Note that by construction, P is such that the surplus of visits $\gamma(w)$ equals to either 0 or 1. However, the same shortcutting procedure is used in [Algorithm 2](#) later in the paper, where $\gamma(w)$ can take higher values as well.

Costs and complexity. The cost of the path P constructed by [Algorithm 1](#) equals to $c(P') \leq c(P_{c,1,s,t}^\alpha) + c(\text{TP}_{c,r,s,t}^*)$. Since $c(\text{TP}_{c,r,s,t}^*)$ is an optimal solution to a relaxation of the MANY-VISITS PATH TSP, its cost is a lower bound to the cost of the corresponding optimal solution, $P_{c,r,s,t}^*$. Since the cost of $P_{c,1,s,t}^\alpha$ is at most α times the cost of an optimal single-visit TSP path $P_{c,1,s,t}^*$, and $c(P_{c,1,s,t}^*) \leq c(P_{c,r,s,t}^*)$ holds for any r , [Algorithm 1](#) provides an $(\alpha + 1)$ -approximation for the MANY-VISITS PATH TSP. Using Zenklusen's recent polynomial-time $3/2$ -approximation algorithm on the single-visit metric PATH TSP [\[58\]](#) in [Step 1](#) yields the approximation guarantee of $5/2$ stated in the theorem.

The transportation problem in [Step 2](#) can be solved in $\mathcal{O}(n^3 \log n)$ operations using the approach of Orlin [\[43\]](#) or its extension due to Kleinschmidt and Schannath [\[34\]](#). [Step 3](#) can also be performed in polynomial time [\[23\]](#), and the number of closed walks can be bounded by $\mathcal{O}(n^2)$. Moreover, the total surplus of degrees in P is at most $n - 2$, therefore the number of operations performed during shortcutting in [Step 4](#) is also bounded by $\mathcal{O}(n)$. This proves that the algorithm has a polynomial time complexity.³ \square

Remark 1. *The TSP, as well as the PATH TSP can also be formulated for directed graphs, where the costs c are asymmetric. (Note that c still has to satisfy the triangle inequality, which implies the following bound for the self-loops: $c(vv) \leq \max_{u \neq v} \{c(vu) + c(uv)\}$.) In a recent breakthrough, Svensson et al. [\[52\]](#) gave the first constant-factor approximation for the metric ATSP. In subsequent work, Traub and Vygen [\[54\]](#) improved the constant factor to $22 + \varepsilon$ for any $\varepsilon > 0$. Moreover, Feige and*

³One can obtain a $5/2$ -approximation for the metric MANY-VISITS TSP by simply running [Algorithm 1](#) for every pair $(u, v) \in V \times V$ and setting $s = u$ and $t = v$, then choosing a solution whose cost together with the cost of the edge uv is minimal. However, [Algorithm 1](#) can be simplified while maintaining the same approximation guarantee. This approach appeared in the unpublished manuscript [\[4\]](#) by a superset of the authors and has a simpler proof, as the algorithm does not involve making shortcuts.

Singh [14] proved that an α -approximation for the metric ATSP yields a $(2\alpha + \varepsilon)$ -approximation for the metric PATH-ATSP, for any $\varepsilon > 0$. By combining these results with a suitable modification of Algorithm 1, we can obtain a $(23 + \varepsilon)$ -approximation for the metric MANY-VISITS ATSP, and a $(45 + \varepsilon)$ -approximation for any $\varepsilon > 0$ for the metric MANY-VISITS PATH-ATSP in polynomial time.

4 A $3/2$ -Approximation for the Metric Many-visits Path TSP

In this section we show how to obtain a $3/2$ -approximation for the metric MANY-VISITS PATH TSP. Our approach follows the general strategy of Zenklusen [58], but we need to make several crucial modifications for the many-visits setting with exponentially large requests. This means that instead of calculating a constrained spanning tree, we use the result in Theorem 4 to obtain a connected multigraph P with a sufficiently large number of edges. Then compute a matching M so that all the degrees in $P + M$ have the correct parity, and the cost of $P + M$ is at most $3/2$ times the optimal cost. In order to ensure this bound, we have to enforce certain restrictions on P , similarly to the computation of the spanning tree in [58]. However, as we will show, the many-visits setting leads to further challenges.

For technical reasons, from now on we assume that the two endpoints s and t are different. Let us start by defining the Held-Karp relaxation of the MANY-VISITS PATH TSP as $\min\{c^\top x \mid x \in P_{\text{HK}}^{\text{MV}}\}$, where $P_{\text{HK}}^{\text{MV}}$ denotes the following polytope:

$$(2) \quad P_{\text{HK}}^{\text{MV}} := \left\{ x \in \mathbb{R}_{\geq 0}^E \mid \begin{array}{ll} x(\delta(C)) \geq 2 & \forall C \subset V, C \neq \emptyset, |C \cap \{s, t\}| \in \{0, 2\} \\ x(\delta(C)) \geq 1 & \forall C \subseteq V, |C \cap \{s, t\}| = 1 \\ x(\delta(v)) = 2 \cdot r(v) & \forall v \in V - \{s, t\} \\ x(\delta(s)) = 2 \cdot r(s) - 1 \\ x(\delta(t)) = 2 \cdot r(t) - 1 \end{array} \right\}$$

The Q -join polytope (where $Q \subseteq V$ is of even cardinality) is defined as follows:

$$(3) \quad P_{Q\text{-join}}^\uparrow := \{x \in \mathbb{R}_{\geq 0}^E \mid x(\delta(C)) \geq 1 \quad \forall Q\text{-cut } C \subset V\},$$

where a Q -cut is a set $C \subseteq V$ with $|C \cap Q|$ odd.

In the following we assume arbitrary but fixed parameters c, r and denote the optimal many-visits TSP path by $P^* = P_{c,r,s,t}^*$. Given a solution y of the linear program $\min\{c^\top x \mid x \in P_{\text{HK}}^{\text{MV}}\}$, the vector $y/2$ is not necessarily in $P_{Q\text{-join}}^\uparrow$ for every even set $Q \subseteq V$. Indeed, y only needs to have a load of 1 on s - t -cuts, therefore $y/2$ may violate some of the constraints of $P_{Q\text{-join}}^\uparrow$. This means that calculating a minimum cost perfect matching on an arbitrary even set $Q \subseteq V$ might lead to a matching M with higher cost than $c(P^*)/2$. Therefore, simply taking a solution P provided by Theorem 4 and a minimum cost matching M on the vertices with degrees having incorrect parity, then applying shortcuts would not lead to a $3/2$ -approximation.

To circumvent this problem, we would like to have a control over the vertices of P that take part in the perfect matching phase of the algorithm. Similarly to Zenklusen [58], we calculate a point q that is feasible for the Held-Karp relaxation of the MANY-VISITS PATH TSP, and that is only needed for the analysis of the algorithm. Let $\text{odd}(P)$ denote the vertices v with an odd degree in P . We need P and q to meet the following requirements:

$$(R1) \quad c(P) \leq c(P^*),$$

$$(R2) \quad c(q) \leq c(P^*), \text{ and}$$

$$(R3) \quad q/2 \in P_{Q_P\text{-join}}^\uparrow, \text{ where } Q_P := \text{odd}(P) \Delta \{s, t\},$$

where $c(q)$ stands for the cost of the vector q with respect to c , that is, $c(q) = \sum_{e \in E} c(e)q(e)$.

Adding a shortest Q_P -join J to the multigraph P results in a multigraph P' where every vertex $v \in V - \{s, t\}$ has an even degree at least $2 \cdot r(v)$, and every $v \in \{s, t\}$ has an odd degree at least $2 \cdot r(v) - 1$. Due to (R3), the cost of the shortest Q_P -join J satisfies $c(J) \leq c(q)/2$. Therefore, using Wolsey's analysis for Christofides' algorithm, the solution P'' obtained by taking the edges of P' and applying shortcuts has cost at most $3/2 \cdot c(P^*)$.

Let x^* be an optimal solution to the Held-Karp relaxation of the MANY-VISITS PATH TSP:

$$(4) \quad \min\{c^\top x \mid x \in P_{\text{HK}}^{\text{MV}}\} .$$

In order to obtain P and q that satisfy the conditions (R1)-(R3) above, we will calculate another solution $y \in P_{\text{HK}}^{\text{MV}}$ with $c(y) \leq c(P^*)$, and set q to be the midpoint between x^* and y , that is, $q = x^*/2 + y/2$. The construction of P also depends on y , the details are given in Algorithm 2 and the reasoning in the proof of Theorem 2. Being a convex combination of two points in $P_{\text{HK}}^{\text{MV}}$, q is in $P_{\text{HK}}^{\text{MV}}$ as well. We would like to ensure the existence of a multigraph P such that $q/2 \in P_{Q_P\text{-join}}^\uparrow$, therefore we need to construct y accordingly.

Let $Q \subseteq V$ be a set of even cardinality. Recall the definition of $P_{Q\text{-join}}^\uparrow$ at Equation (3), which requires that the load on Q -odd cuts is at least 1. Since q is in $P_{\text{HK}}^{\text{MV}}$, $q(\delta(C))/2 \geq 1$ holds for any non- s - t -cuts, i.e. any cuts $C \subset V, C \neq \emptyset$ with $|C \cap \{s, t\}| \in \{0, 2\}$. However, for s - t -cuts, the property $y \in P_{\text{HK}}^{\text{MV}}$ only implies $y(\delta(C)) \geq 1$. If in addition $x^*(\delta(C)) \geq 3$ holds, then we get $q(\delta(C))/2 \geq 1$ regardless of our choice of the multigraph P . If, however, $x^*(\delta(C)) < 3$ holds, we cannot use the same argument. In that case we need to take care of the constraints of $P_{Q_P\text{-join}}^\uparrow$ that correspond to s - t -cuts, where the x^* -load is strictly less than 3. Let us denote these cuts by $\mathcal{B}(x^*)$, that is,

$$(5) \quad \mathcal{B}(x^*) := \{C \subseteq V \mid s \in C, t \notin C, x^*(\delta(C)) < 3\} .$$

For a family $\mathcal{B} \subseteq \{C \subseteq V \mid s \in C, t \notin C\}$ of s - t -cuts, we say that a point $y \in P_{\text{HK}}^{\text{MV}}$ is \mathcal{B} -good, if for every $B \in \mathcal{B}$ we have

- (i) either $y(\delta(B)) \geq 3$,
- (ii) or $y(\delta(B)) = 1$, and y is integral on the edges $\delta(B)$.

Therefore, if $y \in P_{\text{HK}}^{\text{MV}}$ is $\mathcal{B}(x^*)$ -good, then $q = x^*/2 + y/2$ satisfies $q(\delta(C))/2 \geq 1$ for every Q_P -cut C . We will refer to a cut B satisfying condition (i) as a *type (i) cut*, and if it satisfies condition (ii) we will refer to it as a *type (ii) cut*. Note that condition (ii) translates to having a single edge $f \in \delta(B)$ with $y(f) = 1$ and $y(e) = 0$ for all other edges e from $\delta(B)$. The notion of \mathcal{B} -goodness was introduced by Zenklusen for the elements of the polytope P_{HK} in relation to metric PATH TSP.

Lemma 9. *The characteristic vector χ_U of any many-visits s - t path U is \mathcal{B} -good for any family \mathcal{B} of s - t -cuts.*

Proof. The lemma easily follows from the fact that a many-visits s - t path U crosses any s - t -cut an odd number of times. \square

We present our algorithm for the metric MANY-VISITS PATH TSP as [Algorithm 2](#).

Algorithm 2 A $3/2$ -approximation algorithm for the metric MANY-VISITS PATH TSP

Input: A complete undirected graph $G = (V, E)$, costs $c : E \rightarrow \mathbb{R}_{\geq 0}$ satisfying the triangle inequality, requests $r : V \rightarrow \mathbb{Z}_{\geq 1}$, distinct vertices $s, t \in V$.

Output: An s - t -path that visits each $v \in V$ exactly $r(v)$ times.

- 1: Calculate an optimal solution x^* to the Held-Karp relaxation of the MANY-VISITS PATH TSP, i.e. $x^* := \operatorname{argmin}\{c^\top x \mid x \in P_{\text{HK}}^{\text{MV}}\}$.
- 2: Determine a $\mathcal{B}(x^*)$ -good solution $y \in P_{\text{HK}}^{\text{MV}}$ minimizing $c^\top y$.
- 3: Let $B_1 \subset \dots \subset B_k$ denote the *type (ii)* cuts with respect to y .
- 4: Compute a connected multigraph P on $(V, \operatorname{supp}(y))$ such that
 - a: each vertex $v \in V - \{s, t\}$ has degree at least $2 \cdot r(v) - 1$,
 - b: each vertex $v \in \{s, t\}$ has degree at least $2 \cdot r(v) - 2$, and
 - c: P contains no parallel edges leaving B_i for $i = 1, \dots, k$.
- 5: Compute a minimum-cost matching M with respect to c on the vertices $\operatorname{odd}(P) \Delta \{s, t\}$.
- 6: Let P' denote the many-visits path $P + M$.
- 7: Do shortcuts in P' and obtain an s - t -walk P'' that visits each city v exactly $r(v)$ times.

return P'' .

In Step 4 of the algorithm, we use [Theorem 4](#) to obtain a multigraph with additional properties besides the degree requirements. In the single-visit counterpart of the problem, one can show that even though $x^*(\delta(B)) < 3$ and $y(\delta(B)) = 1$ for *type (ii)* cuts B , the corresponding point $q/2 = x^*/4 + y/4$ is still in $P_{Q_P\text{-join}}^\uparrow$. However, due to the possible parallel edges in P , the parity argument given by Zenklusen [58] does not hold, therefore we need to treat this case separately. For this reason we make the following distinction. Let E_y denote the set of edges that correspond to *type (ii)* cuts in y , that is

$$(6) \quad E_y := \{e \in E \mid \exists B \in \mathcal{B} : \operatorname{supp}(y) \cap \delta(B) = e\} .$$

We let $U^*(e) := 1$ for all $e \in E_y$, $U^*(e) := +\infty$ for the rest of the edges of $\text{supp}(y)$, and $U^*(e) := 0$ for edges $e \in E - \text{supp}(y)$. Finally, we set $L^*(e) := 0$ for every edge $e \in E$. According to the claim of [Theorem 4](#), we can compute a multigraph P satisfying the conditions in [Steps 4.a to 4.c](#), such that the cost of P is at most $\min\{c^\top x \mid x \in P^{\text{CG}}(\rho^*, L^*, U^*)\}$, where the polytope $P^{\text{CG}}(\rho^*, L^*, U^*)$ depends on the instance (G, c, r, s, t) and can be written in the following form:

$$(7) \quad P^{\text{CG}}(\rho^*, L^*, U^*) := \left\{ x \in \mathbb{R}_{\geq 0}^E \mid \begin{array}{l} \text{supp}(x) \text{ is connected} \\ x(E) = r(V) \\ x(\delta(v)) \geq 2 \cdot r(v) \quad \forall v \in V - \{s, t\} \\ x(\delta(v)) \geq 2 \cdot r(v) - 1 \quad \forall v \in \{s, t\} \\ 0 \leq x(e) \leq 1 \quad \forall e \in E_y \\ 0 \leq x(e) \leq +\infty \quad \forall e \in \text{supp}(y) - E_y \\ x(e) = 0 \quad \forall e \in E - \text{supp}(y) \end{array} \right\}$$

It is not difficult to see that $y \in P^{\text{CG}}(\rho, L, U)$, and thus

$$(8) \quad c(P) \leq \min \{c^\top x \mid x \in P^{\text{CG}}(\rho^*, L^*, U^*)\} \leq c^\top y ,$$

therefore $c(P) \leq c^\top y$ holds; this is one of the reasons behind restricting P to $\text{supp}(y)$. Moreover, according to [Lemma 9](#), the inequality $c^\top y \leq c(P^*)$ holds, hence the bound $c(P) \leq c(P^*)$ follows. Now we have all the ingredients to prove our main theorem.

Theorem 2. *There is a polynomial-time $3/2$ -approximation for the metric MANY-VISITS PATH TSP. The algorithm runs in time polynomial in n and $\log r$.*

Proof. Recall that $Q_P = \text{odd}(P) \Delta \{s, t\}$. First prove that $q = x^*/2 + y/2$ implies that $q/2$ is in $P_{Q_P\text{-join}}^\uparrow$. For that we need to show that we calculated the solution y in a way that q satisfies $q(\delta(C))/2 \geq 1$ for all cuts $C \subset V$ for which $|C \cap \text{odd}(P) \Delta \{s, t\}|$ is odd.

Clearly, $q \in P_{\text{HK}}^{\text{MV}}$, as q is the midpoint of two points from $P_{\text{HK}}^{\text{MV}}$. Therefore, for any Q_P -cut $C \subseteq V$ that is not an s - t -cut, we have $q(\delta(C))/2 \geq 1$ as needed. Moreover, by definition, for any Q_P -cut $C \subseteq V$ that is an s - t -cut and is not included in $\mathcal{B}(x^*)$, we have $x^*(\delta(C)) \geq 3$, and so

$$(9) \quad \frac{1}{2}q(\delta(C)) = \frac{1}{4}(x^*(\delta(C)) + y(\delta(C))) \geq 1,$$

as $y \in P_{\text{HK}}^{\text{MV}}$, and thus $y(\delta(C)) \geq 1$.

It remains to consider Q_P -cuts $C \subseteq V$ that are in $\mathcal{B}(x^*)$. Since y is $\mathcal{B}(x^*)$ -good by construction, either $y(\delta(C)) \geq 3$, or $y(\delta(C)) = 1$ with y being integral on the edges $\delta(C)$. If $y(\delta(C)) \geq 3$, then $q(\delta(C))/2 \geq 1$ follows from $x^*(\delta(C)) \geq 1$ and the definition of q . If $y(\delta(C)) = 1$ and y is integral on the edges $\delta(C)$, it holds that $y(e) = 0$ for all edges of $\delta(C)$ except for one $f \in \delta(C)$ where $y(f) = 1$. It is at this point where we exploit the restrictions imposed on P . Since $\text{supp}(P) \subseteq \text{supp}(y)$, and the load on an edge $e \in E_y$ is at most 1 in P , the only edge of P with a positive load on $\delta(C)$ is f , and that load is at most 1. Moreover, every cut has at least 1 load in P , which means

$|P \cap \delta(C)| = 1$. But an s - t -cut $C \subseteq V$ with $|\delta_P(C)|$ odd cannot be a Q_P -cut because of the following:

$$(10) \quad \begin{aligned} |C \cap \text{odd}(P)| &\equiv \sum_{v \in C} |\dot{\delta}_P(v)| \pmod{2} \\ &= 2 \cdot |\{uv \in P \mid u, v \in C\}| + |\delta_P(C)|. \end{aligned}$$

Equation (10) implies that $|C \cap \text{odd}(P)|$ is odd, and hence $|C \cap Q_P| = |C \cap (\text{odd}(P) \Delta \{s, t\})|$ is even because C is an s - t -cut. By the above, any cut of type (ii) partitions the vertices of $\text{odd}(P) \Delta \{s, t\}$ into two subsets of even cardinality.⁴ This means that no cut constraint of $P_{Q_P\text{-join}}^\uparrow$ requires a load of 1 for $q/2$ on C , and so $q/2 \in P_{Q_P\text{-join}}^\uparrow$ holds.

The cost of the matching M can therefore be bounded as follows:

$$(11) \quad c(M) \leq c\left(\frac{q}{2}\right) = \frac{1}{4}c^\top x^* + \frac{1}{4}c^\top y \leq \frac{1}{2}c(\mathbf{P}^*),$$

since $c^\top x^* \leq c(\mathbf{P}^*)$. Thus, the multigraph obtained from $P + M$ has cost at most $3/2 \cdot c(\mathbf{P}^*)$, as claimed.

Shortcuts and complexity. According to Theorem 4, every vertex $v \in V - \{s, t\}$ has degree at least $2 \cdot r(v) - 1$, while vertices s and t have degrees at least $2 \cdot r(s) - 2$ and $2 \cdot r(t) - 2$ respectively, in the multigraph P . The matching M provides 1 additional degree for vertices with the wrong parity, therefore P' will have an even degree at least $r(v)$ for all $v \in V - \{s, t\}$ and an odd degree at least $r(v) - 1$ for $v \in \{s, t\}$. This means that P' corresponds to a many-visits s - t -path that visits each vertex v at least $r(v)$ times, but possibly more. In Step 7 we proceed with taking shortcuts the way described in Algorithm 1, so that P'' is a feasible solution to the MANY-VISITS PATH TSP instance (G, c, r, s, t) .

Now we turn to the complexity analysis. The constraints of the Held-Karp relaxation (Equation (4)) of the MANY-VISITS PATH TSP can be tested in time polynomial in n and $\log r(V)$, hence calculating x^* takes a $\text{poly}(n, \log r(V))$ time as well [47, §58.5]. This means Step 1 takes time polynomial in n and $\log r(V)$. According to Lemma 14, Step 2 also has polynomial time complexity, and By Theorem 4, Step 4 takes polynomial time and calculating a matching in Step 5 can be done efficiently as well. Finally, since the number of edges in P is $r(V)$ and the matching M contributes at most $n/2$ edges, we remove at most $n/2$ edges from P' to obtain our solution P'' . This means that the number of operations performed in Step 7 can be bounded by $\mathcal{O}(n)$. The claimed time complexity follows. \square

Corollary 3. *There is a $3/2$ -approximation for the metric MANY-VISITS TSP that runs in time polynomial in n and $\log r$.*

⁴For a cut C with $y(\delta(C)) = 1$ and y being integral on $\delta(C)$, the term $|T \cap \delta(C)|$ in the proof of Theorem 2.1 of Zenklusen [58] corresponds to the term $|\delta_P(C)|$ in Equation (10). Since the spanning tree T computed on $\text{supp}(y)$ in the algorithm of [58] cannot contain parallel edges, $|T \cap \delta(C)|$ has a value of 1 without enforcing an upper bound on the edge $e \in \delta_T(C)$ for $y(e) = 1$.

Proof. Let $G = (V, E)$ be a graph and (G, c, r) denote a metric MANY-VISITS TSP instance. Choose an arbitrary vertex $v \in V$, and construct a metric MANY-VISITS PATH TSP instance $(\hat{G}, \hat{c}, \hat{r}, s_v, t_v)$ as follows. Let \hat{G} be an undirected graph on the vertex set $\hat{V} := V - v \cup \{s_v, t_v\}$ and edge set $\hat{E} := \hat{V} \times \hat{V}$. We define $\hat{c}(s_v t_v)$ as the cost of a self loop at v , $c(vv)$, and $\hat{c}(s_v u) = \hat{c}(t_v u) := c(vu)$ for every vertex $u \in V - v$. Moreover, the self-loops at s_v and t_v have cost $c(vv)$ as well. It is easy to check that \hat{c} satisfies the triangle inequality. Finally, set $\hat{r}(s_v) := r(v)$ and $\hat{r}(t_v) := 1$.

Now we prove that the MANY-VISITS TSP instance (G, c, r) and the corresponding MANY-VISITS PATH TSP instance $(\hat{G}, \hat{c}, \hat{r}, s_v, t_v)$ can be reduced to each other. First let T be a solution to (G, c, r) . Choose an edge vu from T such that $v \neq u$, and let P denote the many-visits s_v - t_v -path obtained from T by deleting v , replacing each occurrence of any edge $vw \in T$ with a copy of $s_v w$ if $w \in V - \{u, v\}$ or with a copy of the loop on s_v if $w = v$, and replacing all but one occurrence of the edge $vu \in T$ with a copy of $s_v u$, while one copy of vu is substituted by $t_v u$. In other words, s_v ‘inherits’ all copies of all edges and self-loops incident to v , except one copy of uv , and t_v inherits one copy of uv . This means that $\deg_P(s_v) = 2 \cdot r(v) - 1$ and $\deg_P(t_v) = 1$. Note that each edge of T is replaced by an edge of the same cost, and every vertex $w \in V - v$ has the same degree in T and P , hence $\deg_P(w) = 2 \cdot r(w)$. Therefore, P is a feasible solution to $(\hat{G}, \hat{c}, \hat{r}, s_v, t_v)$ of the same cost as T .

Now consider a multigraph P that is a solution to $(\hat{G}, \hat{c}, \hat{r}, s_v, t_v)$. Identify the vertices s_v and t_v , denote the new vertex by v , and introduce an edge vv for every copy of the edge $s_v t_v$ in P . Let us denote the resulting multigraph by T . Since $c(uv) = \hat{c}(s_v u) = \hat{c}(t_v u)$ for all $u \in V - v$ and $c(vv) = \hat{c}(s_v s_v) = \hat{c}(t_v t_v) = \hat{c}(s_v t_v)$, replacing s_v and t_v by v the way described above does not change the cost of the multigraph. Moreover, the degree of v in T is $\deg_T(v) = \deg_P(s_v) + \deg_P(t_v) = 2 \cdot r(v) - 1 + 1 = 2 \cdot r(v)$. The degrees of vertices $w \in V - v$ remain unchanged, thus T is a feasible solution to (G, c, r) of the same cost as P .

We therefore showed that for every solution of (G, c, r) there exists a solution of $(\hat{G}, \hat{c}, \hat{r}, s_v, t_v)$ with the same cost, and vice versa. Let now (G, c, r) be a metric MANY-VISITS TSP instance. Pick an arbitrary vertex $v \in V$, and consider the corresponding metric MANY-VISITS PATH TSP instance $(\hat{G}, \hat{c}, \hat{r}, s_v, t_v)$, and obtain a $3/2$ -approximation P using [Algorithm 2](#). Identify s_v and t_v into v again, and substitute each copy of the edge $s_v t_v$ in P by a copy of the self-loop vv . By the above, the resulting multigraph T gives a $3/2$ -approximation to the instance (G, c, r) . \square

Remark 2. *Alternatively, one can directly obtain a $3/2$ -approximation for the metric MANY-VISITS TSP by performing [Step 4.a](#), [Step 5](#) and [Step 7](#) of [Algorithm 2](#). More precisely, calculate a connected multigraph T with degrees at least $2 \cdot r(v) - 1$ and cost at most the optimum using the result of [Theorem 4](#), then calculate a matching on the odd degree vertices and apply shortcuts. This procedure was described by a superset of the authors [\[4\]](#).*

Before we show how to calculate a $\mathcal{B}(x^*)$ -good point $y \in P_{\text{HK}}^{\text{MV}}$, let us show that the number of cuts in \mathcal{B} is polynomial in n , and that the set \mathcal{B} can be computed efficiently:

Lemma 10. *Let $q \in P_{\text{HK}}^{\text{MV}}$. Then the family $\mathcal{B}(q)$ of s - t -cuts of q -value strictly less than 3 satisfies $|\mathcal{B}(q)| \leq n^4$ and can be computed in $\mathcal{O}(mn^4)$ time, where $n := |V|$ and $m := \text{supp}(q)$.*

Proof. Let us define an auxiliary graph $H = (V, E')$ whose edge set E' consists of the edges in $\text{supp}(q)$ and an additional st edge. Let $q_H = q + \chi_{st}$. Clearly, for non- s - t -cuts we have $q_H(\delta_H(C)) = q(\delta(C))$, while for s - t -cuts we have $q_H(\delta_H(C)) = q(\delta(C)) + 1 \geq 2$ because of the newly added edge st . Therefore, the family $\mathcal{B}(q)$ can be written as

$$\mathcal{B}(q) = \{C \subset V \mid s \in C, t \notin C, q_H(\delta_H(C)) < 4\} .$$

The minimum cut has a load of at least 2, and due to Karger [32] the number of cuts with a load less than k times the minimum cut is at most $\mathcal{O}(n^{2k})$. Moreover, using an algorithm by Nagamochi et al. [42], we can enumerate the cuts of size at most k times the minimum cut in time $\mathcal{O}(m^2n + n^{2k}m)$. These results prove that the number of cuts in $|\mathcal{B}(q)|$ is $\mathcal{O}(n^4)$, and that they can be enumerated in time $\mathcal{O}(mn^4)$. \square

The dynamic program

Given a family \mathcal{B} of s - t -cuts, our goal is to determine a minimum cost \mathcal{B} -good point $y \in P_{\text{HK}}^{\text{MV}}$. We use the dynamic programming approach introduced by Traub and Vygen [54] and improved upon by Zenklusen [58]. More precisely, the goal of the dynamic program is to determine which cuts in \mathcal{B} are of [type \(i\)](#), and which ones are of [type \(ii\)](#). Our approach is constructive as the dynamic program also determines a point y that is $\mathcal{B}(x^*)$ -good.

Consider a $\mathcal{B}(x^*)$ -good point y . Let B_1, \dots, B_k denote the [type \(ii\)](#) s - t cuts in \mathcal{B} with respect to y , that is, $y(v_i u_i) = 1$ for exactly one edge $v_i u_i \in \delta(B_i)$ and $y(e) = 0$ for $e \in \delta(B_i) - v_i u_i$. It is not difficult to see that these cuts necessarily form a chain (see e.g. [58]), thus we set the indices such that $B_1 \subsetneq \dots \subsetneq B_k$. The endpoints of $v_i u_i$ are named such that $v_i \in B_i$, $u_i \notin B_i$. Furthermore, we define $B_0 := \emptyset$, $B_{k+1} := V$, $u_0 := s$ and $v_{k+1} := t$ for notational convenience. Note that u_i and v_{i+1} might coincide for some $i = 0, \dots, k + 1$.

The work of Zenklusen [58] argues that the ‘first’ and ‘last’ cuts are [type \(ii\)](#) cuts, that is, $B_1 = \{s\}$ and $B_k = V - \{t\}$, because the constraints of $P_{\text{HK}}^{\text{MV}}$ enforce a degree of 1 on vertices s and t . In the many-visits setting, however, this is not necessarily true, as the instance possibly requires more than one visit for s or t .

Assume for a moment that we knew the cuts B_1, \dots, B_k and the edges $v_i u_i$, and we are looking for a $\mathcal{B}(x^*)$ -good point $y \in P_{\text{HK}}^{\text{MV}}$ such that among all cuts in \mathcal{B} the cuts B_1, \dots, B_k are precisely those where (a) y is integral, and (b) $y(\delta(B_i)) = 1$ for all $i = 1, \dots, k$. Then the \mathcal{B} -good points $y \in P_{\text{HK}}^{\text{MV}}$ that satisfy these constraints (a) and (b) have the following properties for all $i = 1, \dots, k$:

- (P1) $y(v_i u_i) = 1$ and $y(e) = 0$ for all edges $e \in \delta(B_i) - v_i u_i$,
- (P2) the restriction of y to the vertex set $B_{i+1} - B_i$ is a solution to the Held-Karp relaxation for the MANY-VISITS PATH TSP with endpoints u_i and v_{i+1} , with the additional property that $y(\delta(B)) \geq 3$ for every cut $B \in \mathcal{B}$ such that $B_i \cup u_i \subseteq B \subseteq B_{i+1} - v_{i+1}$.

The dynamic program thus aims to find cuts B_1, \dots, B_k while exploiting the properties (P1) and (P2) above. Formally, it is defined to find a shortest path on an auxiliary directed graph. Let us define the auxiliary directed graph $H = (N, A)$ with node set N , arc set A , and length function $d : A \rightarrow \mathbb{R}_{\geq 0}$. The node set N is defined by $N = N^+ \cup N^-$, where

$$\begin{aligned} N^+ &= \{(B, u) \in \mathcal{B} \times V \mid u \notin B\} \cup \{(\emptyset, s)\}, \text{ and} \\ N^- &= \{(B, v) \in \mathcal{B} \times V \mid v \in B\} \cup \{(V, t)\}. \end{aligned}$$

The arc set A is given by $A = A_{\text{HK}} \cup A_E$, where

$$\begin{aligned} A_{\text{HK}} &= \{((B^+, u), (B^-, v)) \in N^+ \times N^- \mid B^+ \subseteq B^-, u, v \in B^- - B^+\}, \text{ and} \\ A_E &= \{((B^-, v), (B^+, u)) \in N^- \times N^+ \mid B^- = B^+\}. \end{aligned}$$

Finally, the lengths $d : A \rightarrow \mathbb{R}_{\geq 0}$ are defined as follows:

$$d(a) = \begin{cases} c(vu) & \text{if } a = ((B, v), (B, u)) \in A_E, \\ \text{OPT}(\text{LP}(a)) & \text{if } a \in A_{\text{HK}}, \end{cases}$$

where $\text{OPT}(\text{LP}(a))$ denotes the optimum value of

$$\begin{aligned} \min \quad & c^\top x \\ (\text{LP}(a)) \quad & \text{subject to} \quad x \in P_{\text{HK}}^{\text{MV}}(B^- - B^+, u, v) \\ & x(\delta(B)) \geq 3 \quad \text{for all } B \in \mathcal{B} \text{ s.t. } B^+ \subseteq B \subseteq B^-, \\ & u \in B, v \notin B, \end{aligned}$$

where $a = ((B^+, u), (B^-, v))$.

For y -values across the cuts $B \in \mathcal{B}$ so that $B \notin \{B_1, \dots, B_k\}$, we require that $y(\delta(B)) \geq 3$ holds. We ensure this by defining modified Held-Karp relaxations of the MANY-VISITS PATH TSP instances between cuts B_i and B_{i+1} for every $i = 0, \dots, k$. More precisely, such an instance is defined on the subgraph of G induced on the vertex set $B_{i+1} - B_i$ with distinguished vertices u_i and v_{i+1} , with the additional property that it has a y -load of at least 3 on each cut $B \in \mathcal{B}$ with $B_i \subset B \subset B_{i+1}$, as shown in (LP(a)). In case $u \neq v$, the polytope $P_{\text{HK}}^{\text{MV}}(W, u, v)$ is defined as follows:

$$(12) \quad P_{\text{HK}}^{\text{MV}}(W, u, v) := \left\{ x \in \mathbb{R}_{\geq 0}^E \left| \begin{array}{ll} x(\delta(C)) \geq 2 & \forall C \subset W, C \neq \emptyset, \\ & |C \cap \{u, v\}| \in \{0, 2\} \\ x(\delta(C)) \geq 1 & \forall C \subset W, |C \cap \{u, v\}| = 1 \\ x(\dot{\delta}(w)) = 2 \cdot r(w) & \forall w \in W - \{u, v\} \\ x(\dot{\delta}(u)) = 2 \cdot r(u) - 1 & \\ x(\dot{\delta}(v)) = 2 \cdot r(v) - 1 & \\ x(e) = 0 & \forall e \in E - E[W] \end{array} \right. \right\}$$

Note that unlike in the single-visit case [58], we allow u being equal to s or v being equal to t in Equation (12), and the corresponding polytopes $P_{\text{HK}}^{\text{MV}}(B_1, s, v_1)$ and $P_{\text{HK}}^{\text{MV}}(V - B_k, u_k, t)$ are feasible.

Let us now cover the case when for some index $i \in \{0, \dots, k\}$, vertices u_i and v_{i+1} coincide.⁵ In the single-visit PATH TSP, the solution is defined to be the all-zero vector if $u_i = v_{i+1}$ is the only vertex in $B_{i+1} - B_i$, and there exists no solution otherwise. However, since we allow for a vertex to be visited more than once (i.e. have a degree more than 2) in a solution to the Held-Karp relaxation for the MANY-VISITS PATH TSP, we use a different extension in our approach. We define the corresponding subproblem as the Held-Karp relaxation for the MANY-VISITS TSP. First assume that $u_i \neq s$ and $u_i \neq t$. Since $y(v_i u_i) = 1$ and $y(u_i u_{i+1}) = 1$ by construction, the degree requirement for u_i in the MANY-VISITS TSP subproblem is *two* less than that in $P_{\text{HK}}^{\text{MV}}$, namely $r(u_i) - 2$. If $u_i = s$ (or $u_i = t$), then due to $y(s u_1) = 1$ (or $y(v_k t) = 1$) the degree requirement for u_i in the subproblem is *one* less than that in $P_{\text{HK}}^{\text{MV}}$, which also equals to $r(u_i) - 2$. Note that if $u_i = v_{i+1}$ there is no cut $B \in \mathcal{B}$ with $u_i \in B$ and $v_{i+1} \notin B$, thus the linear program LP(a) has the form $\min\{c^\top x \mid x \in P_{\text{HK}}^{\text{MV}}(W, u, u)\}$, where:

$$(13) \quad P_{\text{HK}}^{\text{MV}}(W, u, u) := \left\{ x \in \mathbb{R}_{\geq 0}^E \mid \begin{array}{ll} x(\delta(C)) \geq 2 & \forall C \subset W, C \neq \emptyset, \\ x(\delta(w)) = 2 \cdot r(w) & \forall w \in W - u \\ x(\delta(u)) = 2 \cdot r(u) - 2 & \\ x(e) = 0 & \forall e \in E - E[W] \end{array} \right\}.$$

If the requirement for u is $r(u) = 1$ and $|W| > 1$, the polytope $P_{\text{HK}}^{\text{MV}}(W, u, u)$ is empty, and thus the linear program LP(a) has no solution. In this case the cost of the arc a is defined to be infinity. Note however that if $r(u) = 1$ and $W = \{u\}$, the corresponding linear program has a non-zero solution, namely a vector that has value $r(u) - 1$ in the coordinate of the self-loop uu , and 0 otherwise.

To find a \mathcal{B} -good point with minimum cost $c^\top y$, we compute a shortest (\emptyset, s) - (V, t) path with respect to d in H ; due to [Lemmas 9](#) and [11](#) this path has finite length. Let $(\emptyset, s), (B_1, v_1), (B_1, u_1), (B_2, v_2), \dots, (B_k, u_k), (V, t)$ be the nodes on this shortest path, and similarly as before, define $B_0 := \emptyset, u_0 := s$ and $B_{k+1} := V, v_{k+1} := t$. By construction of H , we have $B_0 \subset B_1 \subset \dots \subset B_{k+1}$. Let $x^i \in \mathbb{R}^E$ be an optimal solution to LP(a) for $a = ((B_i, u_i), (B_{i+1}, v_{i+1}))$. Set

$$(14) \quad y := \sum_{i=0}^k x^i + \sum_{i=1}^k \chi_{v_i u_i}.$$

By the definition of the lengths d in H , $c^\top y$ necessarily equals the length ℓ^* of a shortest (\emptyset, s) - (V, t) path in H with respect to d . We now show that y computed in [Equation \(14\)](#) is indeed a $\mathcal{B}(x^*)$ -good point of minimum cost.

Lemma 11. *The length ℓ^* of a shortest (\emptyset, s) - (V, t) path in H with respect to d satisfies $\ell^* \leq \min\{c^\top z \mid z \in P_{\text{HK}}^{\text{MV}}, z \text{ is } \mathcal{B}\text{-good}\}$.*

Proof. Let $\mathcal{B}_z \subseteq \mathcal{B}$ be the family of cuts $B \in \mathcal{B}$ such that $z(f) = 1$ for precisely one edge $f \in \delta(B)$, and $z(e) = 0$ for all other edges $e \in \delta(B) - f$. These are the sets in \mathcal{B}

⁵Note that the corresponding arc in H will have the form $((B^+, w), (B^-, w)) \in A_{\text{HK}}$.

that are **type (ii)** cuts with respect to z , and also \mathcal{B}_z forms a chain: $B_1 \subset \dots \subset B_k$ holds, where $B_i \in \mathcal{B}_z$ for $i = 1, \dots, k$. The cuts $\{B_1, \dots, B_k\}$ defines a partition of V into sets $B'_0 := B_1$, $B'_1 := B_2 - B_1$, \dots , $B'_{k-1} := B_k - B_{k-1}$, $B'_k := V - B_k$. For $i \in \{1, \dots, k\}$, let $v_i u_i$ be the unique edge in $\delta(B_i)$ where $z(v_i u_i) = 1$, so that $v_i \in B_i$ and $u_i \notin B_i$.

Consider the path along nodes $(B_0, u_0), (B_1, v_1), (B_1, u_1), \dots, (B_{k+1}, v_{k+1})$. It suffices to show that the length ℓ of the path is at most $c^\top z$. For each $i \in \{0, \dots, k\}$, the vector $z^i \in \mathbb{R}^E$ is defined to be the restriction of z to $E[B_{i+1} - B_i]$. Assume for a moment that z^i is a feasible solution of $\text{LP}(a)$ with $a = ((B_i, u_i), (B_{i+1}, v_{i+1}))$. Then the total length ℓ is equal to $\sum_{i=0}^k c^\top z^i + \sum_{i=1}^k c(v_i u_i)$ by definition, which is at most $\sum_{i=0}^k c^\top z^i + \sum_{i=1}^k c(v_i u_i) = c^\top z$. Since ℓ^* is minimum among all possible ℓ 's, we get $\ell^* \leq \ell \leq c^\top z$.

Since z is \mathcal{B} -good, and $z^i(\delta(B)) = z(\delta(B))$ for any cut B with $B_i \subsetneq B \subsetneq B_{i+1}$, that means $z^i(\delta(B)) = z(\delta(B)) \geq 3$. It remains to show that $z^i \in P_{\text{HK}}^{\text{MV}}(B'_i, u_i, v_{i+1})$ or $z^i \in P_{\text{HK}}^{\text{MV}}(B'_i, u_i, u_i)$ follows for $i = 0, \dots, k$.

Distinct endpoints. Let us start with the case when $u_i \neq v_{i+1}$. By definition, $z^i(e) = z(e)$ if both endpoints of e are in B'_i and $z^i(e) = 0$ otherwise. Hence, for vertices $w \in B'_i$ such that $w \notin \{u_i, v_{i+1}\}$, $z^i(\delta(w)) = z(\delta(w)) = 2 \cdot r(w)$. First assume that $u_i \neq s$ and $v_{i+1} \neq t$. Both of the endpoints u_i and v_{i+1} have a total z -value of 1 on the edge set $\delta(B'_i)$ due to $z(v_i u_i) = 1$ and $z(v_{i+1} u_{i+1}) = 1$, therefore the z^i -value on the edges incident to the endpoints is $z^i(\delta(u_i)) = z(\delta(u_i)) - 1 = 2 \cdot r(u_i) - 1$ and $z^i(\delta(v_{i+1})) = z(\delta(v_{i+1})) - 1 = 2 \cdot r(v_{i+1}) - 1$. If $u_i = s$ or $v_{i+1} = t$, their z^i -values equal to their z -values, that is $2 \cdot r(s) - 1$ or $2 \cdot r(t) - 1$, respectively. The degree constraints are therefore satisfied.

Finally, we have to show that for a cut $C \subseteq B'_i$, $z^i(\delta(C)) \geq 1$ holds if C is a u_i - v_{i+1} -cut, and $z^i(\delta(C)) \geq 2$ if C does not separate u_i and v_{i+1} . For the single-visit variant, the proof goes by showing that z^i is in the spanning tree polytope of $G[B'_i]$, using the fact that z is in the spanning tree polytope of G and the degree constraints of $v \in B'_i$ [58]. However, these terms do not immediately generalize to the many-visits setting, so we show that the connectivity of z^i follows from the properties of z .

Non- u_i - v_{i+1} -cuts: First let us consider the cuts that do not separate u_i and v_{i+1} , and prove that the total z^i -value across these cuts is at least 2. We may assume that $C \subset B'_i$ does not contain neither u_i nor v_{i+1} throughout this paragraph. In case $u_i, v_{i+1} \in C$, we can take $B'_i - C$ and we are done, as $z^i(\delta(C)) = z^i(\delta(B'_i - C)) = z(\delta(C)) \geq 2$, yielding $z^i(\delta(C)) \geq 2$. Assume first that $u_i \neq s$ and $v_{i+1} \neq t$, and let $C \subset B'_i$. Then $z^i(\delta(C)) \geq 2$ simply because C is a non- s - t -cut and thus $z(\delta(C)) \geq 2$. Now let u_i be equal to s , and let $C \subset B'_i$ be a cut that does not contain either s or v_1 . Likewise, $z^0(\delta(C)) \geq 2$ because C is a non- s - t -cut and thus $z(\delta(C)) \geq 2$. The argument for $v_{i+1} = t$ goes the same way.

u_i - v_{i+1} -cuts: Here we have to prove that if C is a u_i - v_{i+1} -cut, then $z^i(\delta(C)) \geq 1$.

Assume that $u_i \neq s$ and $v_{i+1} \neq t$. Without the loss of generality, we may assume that $u_i \in C$. Since C is a non- s - t -cut, $z(\delta(C)) \geq 2$. If we account for $z(v_i u_i) = 1$, then $z^i(\delta(C)) \geq 1$ follows. One can similarly prove the claim if $u_i \notin C$ and $v_{i+1} \in C$.

Assume now that $u_i = s$ and $u_i \in C$; then C is an s - t -cut, so $z(\delta(C)) \geq 1$. Moreover, $v_{i+1} \notin C$, so $v_{i+1}u_{i+1} \notin \delta(C)$, therefore $z^i(\delta(C)) = z(\delta(C)) \geq 1$. If $u_i = s$ and $v_{i+1} \in C$, then $z^i(\delta(C)) \geq z(\delta(C)) - z(v_{i+1}u_{i+1}) \geq 1$, since C is a non- s - t -cut and thus $z(\delta(C)) \geq 2$. The case $v_{i+1} = t$ can be proved similarly.

Same endpoints. Now we cover the case when $u_i = v_{i+1}$. We need to prove that $z^i \in P_{\text{HK}}^{\text{MV}}(B'_i, u_i, u_i)$, as defined in Equation (13). The argument about the degrees is analogous to the case above, the fact that $z^i(\dot{\delta}(w)) = 2 \cdot r(w)$ directly follows for vertices $w \in B'_i - u_i$. Moreover, if $u_i \notin \{s, t\}$, the endpoint u_i has a z -load of 2 on $\delta(B'_i)$ because of $z(v_i u_i) = 1$ and $z(v_{i+1} u_{i+1}) = z(u_i u_{i+1}) = 1$, hence $z^i(\dot{\delta}(u_i)) = z(\dot{\delta}(u_i)) - 2 = 2 \cdot r(u_i) - 2$, as desired. If $u_i \in \{s, t\}$, then $z(\dot{\delta}(u_i)) = 2 \cdot r(u_i) - 1$. Moreover, $z(s u_1) = 1$ or $z(v_k t) = 1$ if $u_i = s$ or $u_i = t$, respectively. In both cases this means $z^i(\dot{\delta}(u_i)) = z(\dot{\delta}(u_i)) - 1 = 2 \cdot r(u_i) - 2$.

Now turn to the cut constraints, and let $C \subset B'_i$ be a cut. We can assume that $u_i \notin C$, otherwise we take $B'_i - C$, and we are done. If $u_i \notin C$, then $z^i(\delta(C)) = z(\delta(C)) \geq 2$ because C is a non- s - t -cut. The proof is complete. \square

Lemma 12. $y \in P_{\text{HK}}^{\text{MV}}$.

Proof. To prove this claim, we use the properties of x^i . We again distinguish two cases, based on whether the endpoints of subproblem LP(a) are the same or different.

Degree constraints. First, consider the indices i with $u_i \neq v_{i+1}$. By definition, the vector x^i satisfies $x^i \in P_{\text{HK}}^{\text{MV}}(B_{i+1} - B_i, u_i, v_{i+1})$, meaning that it is a solution to the Held-Karp relaxation for MANY-VISITS PATH TSP in the induced subgraph $G[B_{i+1} - B_i]$ with endpoints u_i and v_{i+1} . Recall that $y = \sum_{i=0}^k x^i + \sum_{i=1}^k \chi_{v_i u_i}$ by definition. For a $v \in B_{i+1} - B_i$, the value $y(\dot{\delta}(v))$ is equal to

1. $x^i(\dot{\delta}(v))$ if $v = u_i = s$ or $v = v_{i+1} = t$,
2. $x^i(\dot{\delta}(v)) + 1 = 2 \cdot r(v)$ if $v \notin \{s, t\}$ and $v = u_i$ or $v = v_{i+1}$ for some i , due to the edge $v_i u_i$ or $v_{i+1} u_{i+1}$, respectively; and
3. $x^i(\dot{\delta}(v)) = 2 \cdot r(v)$ otherwise.

By the above, $y(\dot{\delta}(s)) = 2 \cdot r(s) - 1$, $y(\dot{\delta}(t)) = 2 \cdot r(t) - 1$, and $y(\dot{\delta}(v)) = 2 \cdot r(v)$ for $v \notin \{s, t\}$, therefore the degree constraints are satisfied for all $v \in V$.

If $u_i = v_{i+1}$, the value $y(\dot{\delta}(v))$ is equal to

1. $x^i(\dot{\delta}(v)) + 1$ if $v = s$ or $v = t$, because of the edge $v u_1$ or $v_k v$, respectively,
2. $x^i(\dot{\delta}(v)) + 2 = 2 \cdot r(v)$ if $v \notin \{s, t\}$ and $v = u_i = v_{i+1}$ for some i , due to the edge $v_i u_i$ and $v_{i+1} u_{i+1}$; and
3. $x^i(\dot{\delta}(v)) = 2 \cdot r(v)$ otherwise.

Again, we get $y(\dot{\delta}(s)) = 2 \cdot r(s) - 1$, $y(\dot{\delta}(t)) = 2 \cdot r(t) - 1$, and $y(\dot{\delta}(v)) = 2 \cdot r(v)$ for $v \notin \{s, t\}$, therefore the degree constraints are satisfied for all $v \in V$.

Cut constraints. It remains to show that y satisfies the cut constraints. As in the proof of Lemma 11, instead of building on a spanning subgraph polytope, we directly prove that the cut constraints hold. As before, the cuts $\{B_1, \dots, B_k\}$ define a partition of V into sets $B'_0 := B_1$, $B'_1 := B_2 - B_1$, \dots , $B'_{k-1} := B_k - B_{k-1}$, $B'_k := V - B_k$.

Let us first consider the value of y on s - t -cuts. For $B_i \in \{B_1, \dots, B_k\}$ the y -load on $\delta(B_i)$ equals to 1 due to the edge $v_i u_i$, therefore it satisfies the constraint $y(\delta(B_i)) \geq 1$. If C is a s - t -cut such that $C \notin \{B_1, \dots, B_k\}$, then there is at least one index $i \in \{0, \dots, k\}$, such that both $B'_i \cap C$ and $B'_i - C$ are not empty. In other words, there is at least one vertex from B'_i on both sides of the cut C .

If $u_i \neq v_{i+1}$, x^i satisfies the constraints of LP for $a = ((B_i, u_i), (B_{i+1}, v_{i+1}))$, we have $x^i \in P_{\text{HK}}^{\text{MV}}(B'_i, u_i, v_{i+1})$. That means x^i has a load of at least 1 on edges leaving every proper subset of B'_i , including $B'_i \cap C$, and $x^i(\delta(B'_i \cap C)) \geq 1$ implies $y(\delta(B'_i \cap C)) \geq 1$, which yields $y(\delta(C)) \geq 1$.

If $u_i = v_{i+1}$, then $x^i \in P_{\text{HK}}^{\text{MV}}(B'_i, u_i, u_i)$, which means that $y(\delta(B'_i - C, B'_i \cap C)) \geq 2$, so $y(\delta(C)) \geq 1$ follows.

If C is a non- s - t -cut, we distinguish the following three cases. Note that in neither of the cases is $B'_0 \subseteq C$ or $B'_k \subseteq C$ a possibility, as that would make C an s - t -cut.

If $C \subsetneq B'_i$ for some i , and C is a u_i - v_{i+1} -cut so that $u_i \in C$ (or $v_{i+1} \in C$), then $y(\delta(C))$ has at least 1 load from the fact that $x^i \in P_{\text{HK}}^{\text{MV}}(B'_i, u_i, v_{i+1})$, and 1 load from the edge $v_{i-1}u_i$ (or $v_i u_{i+1}$). If C is *not* a u_i - v_{i+1} -cut, and u_i, v_{i+1} are in C , then $y(\delta(C)) \geq 2$ because of the edges $v_i u_i$ and $v_{i+1} u_{i+1}$; while if u_i, v_{i+1} are not in C then $y(\delta(C)) \geq 2$ follows from $x^i \in P_{\text{HK}}^{\text{MV}}(B'_i, u_i, v_{i+1})$. Note that if $u_i = v_{i+1}$, C can only be a non- u_i - v_{i+1} -cut. In that case $x^i(\delta(C)) \geq 2$ because $x^i \in P_{\text{HK}}^{\text{MV}}(B'_i, u_i, u_i)$, and thus $y(\delta(C)) \geq 2$ follows.

If $C = \bigcup_{i \in \mathcal{I}} B'_i$ for some $\mathcal{I} \subset \{0, \dots, k\}$, let us define $i_{\min} := \min\{i \mid B'_i \subset C\}$ and $i_{\max} := \max\{i \mid B'_i \subset C\}$. Then $y(\delta(C)) \geq y(v_{i_{\min}-1} u_{i_{\min}}) + y(v_{i_{\max}} u_{i_{\max}+1}) = 2$.

Else there exists a set B'_i such that $C \cap B'_i \neq \emptyset$ and $C \not\subseteq B'_i$ and $B'_i \not\subseteq C$ hold, then let us define i_{\min} and i_{\max} as follows:

$$\begin{aligned} i_{\min} &:= \min\{i \mid C \cap B'_i \neq \emptyset, C \not\subseteq B'_i, B'_i \not\subseteq C\} , \\ i_{\max} &:= \max\{i \mid C \cap B'_i \neq \emptyset, C \not\subseteq B'_i, B'_i \not\subseteq C\} . \end{aligned}$$

Suppose that $i_{\min} \neq i_{\max}$. Then, $\delta(C)$ has at least 1 y -load on $\delta(B'_{i_{\min}} \cap C, B'_{i_{\min}} - C)$, as well as at least 1 y -load between on $\delta(B'_{i_{\max}} \cap C, B'_{i_{\max}} - C)$, thus $y(\delta(C)) \geq 2$. In case of $i_{\min} = i_{\max}$, there must exist another index $i \neq i_{\min}$ such that $B'_i \subset C$ (as otherwise we are back in one of the previous two cases), in which case there is at least one edge e in $\delta(C \cap B'_i)$ such that $y(e) = 1$ (either $v_i u_i$ or $v_{i+1} u_{i+1}$ or both); in total $y(\delta(C)) \geq 2$ holds.

This concludes the proof of Lemma 12. \square

Lemma 13. y is \mathcal{B} -good.

Proof. The proof follows the lines of the corresponding proof of Lemma 3.3 of Zenklusen [58]: there the claim can be deduced from cut constraints of P_{HK} , while in our case it follows from those of the polytope $P_{\text{HK}}^{\text{MV}}$. Nevertheless, we include the full proof here for the sake of completeness.

For $i \in \{1, \dots, k\}$, we have by construction of y that $y(v_i u_i) = 1$ and $y(e) = 0$ for other edges $e \in \delta(B_i)$. This means that all cuts B_i satisfy (ii) of the definition of \mathcal{B} -goodness, i.e. the y -value is 1 and y is integral. Let us show that for any other cut $B \in \mathcal{B} - \{B_1, \dots, B_k\}$, the y -load satisfies (i) of the definition.

First suppose that $\{B_1, \dots, B_k\} \cup B$ is not a chain, in this case there is some index $j \in \{0, \dots, k\}$, such that neither $B \subseteq B_j$ nor $B_j \subseteq B$ is true. Hence

$$\begin{aligned} y(\delta(B)) + 1 &= y(\delta(B)) + y(\delta(B_j)) \\ &\geq y(\delta(B - B_j)) + y(\delta(B_j - B)) \\ &\geq 4 . \end{aligned}$$

The first line follows from $y(\delta(B_j)) = 1$, this was shown at the beginning of the proof. The first inequality holds by the cut functions $C \rightarrow y(\delta(C))$ being symmetric and submodular. Since B and B_j are s - t -cuts, $B - B_j$ and $B_j - B$ are non- s - t -cuts, and the y -load of both of these cuts is at least 2, hence the second inequality follows.

Suppose that $\{B_1, \dots, B_k\} \cup B$ is a chain, then there is an index j such that $B_j \subsetneq B \subsetneq B_{j+1}$. If $u_j \in B$ and $v_{j+1} \notin B$, then $x^j(\delta(B)) \geq 3$ because of the constraints of the corresponding linear program $\text{LP}(a)$, where $a = ((B_j, u_j), (B_{j+1}, v_{j+1}))$. Since $y \geq x^j$ holds for all j component-wise, $y(\delta(B)) \geq 3$ follows. If $u_j \notin B$ and $v_{j+1} \in B$, then both the edges $v_j u_j$ and $v_{j+1} u_{j+1}$ are in $\delta(B)$; moreover $x^j(\delta(B)) \geq 1$ since B is a u_j - v_{j+1} -cut, therefore

$$y(\delta(B)) \geq x^j(\delta(B)) + y(v_j u_j) + y(v_{j+1} u_{j+1}) \geq 3 .$$

Finally, if B is not an u_j - v_{j+1} -cut, $x^j(\delta(B)) \geq 2$ since $x^j \in P_{\text{HK}}^{\text{MV}}(B_{j+1} - B_j, u_j, v_{j+1})$. Moreover, either $u_i v_i$ or $u_{i+1} v_{i+1}$ is an edge in $\delta(B)$, depending on whether u_i and v_{j+1} are in B or not; both of the possibilities imply $y(\delta(B)) \geq 3$. \square

Lemma 14. *Let $\mathcal{B} \subseteq \{C \subseteq V \mid s \in C, t \notin C\}$. One can determine in time polynomial in $|\mathcal{B}|$ and the input size of (G, s, t, c) a \mathcal{B} -good point $y \in P_{\text{HK}}^{\text{MV}}$ of minimum cost.*

Proof. The number of nodes and arcs in H are polynomial in $|\mathcal{B}|$. Calculating a shortest path on H takes time polynomial in $|H|$. The feasibility of the linear programs $(\text{LP}(a))$ can be checked in time $\text{poly}(n, \log r, |\mathcal{B}|)$, therefore an optimal solution can also be found, using the ellipsoid method, in time polynomial in n , $\log r$ and $|\mathcal{B}|$ (see the discussion in §58.5 of Schrijver's book [47]). \square

Remark 3. *It is worth considering how Algorithm 2 proceeds when applied to the single-visit TSP, that is, when $r(v) = 1$ for each $v \in V$. The output of Algorithm 3 in Step 4 is then a connected multigraph with $r(V) - 1 = n - 1$ edges. This means that each vertex v has degree at least $2 \cdot r(v) - 1 = 1$, which boils down to a connected graph with $n - 1$ edges, therefore a spanning tree on G with the additional properties (R1)-(R3). Thus Algorithm 2 performs the same operations, as the algorithm of Zenklusen [58] for the PATH TSP.*

5 Approximation Algorithm for the Bounded Degree g -Polymatroid Element with Multiplicities Problem

5.1 Polyhedral background

In what follows, we make use of some basic notions and theorems of the theory of generalized polymatroids. For background, see for example the paper of Frank and Tardos [18] or Chapter 14 in the book by Frank [17].

Given a ground set S , a set function $b : 2^S \rightarrow \mathbb{Z}$ is *submodular* if

$$b(X) + b(Y) \geq b(X \cap Y) + b(X \cup Y)$$

holds for every pair of subsets $X, Y \subseteq S$. A set function $p : 2^S \rightarrow \mathbb{Z}$ is *supermodular* if $-p$ is submodular. As a generalization of matroid rank functions, Edmonds introduced the notion of polymatroids [12]. A set function b is a *polymatroid function* if $b(\emptyset) = 0$, b is non-decreasing, and b is submodular.

We define

$$P(b) := \{x \in \mathbb{R}_{\geq 0}^S \mid x(Y) \leq b(Y) \text{ for every } Y \subseteq S\} .$$

The set of integral elements of $P(b)$ is called a *polymatroidal set*. Similarly, the *base polymatroid* $B(b)$ is defined by

$$B(b) := \{x \in \mathbb{R}^S \mid x(Y) \leq b(Y) \text{ for every } Y \subseteq S, x(S) = b(S)\} .$$

Note that a base polymatroid is just a facet of the polymatroid $P(b)$. In both cases, b is called the *border function* of the polyhedron. Although non-negativity of x is not assumed in the definition of $B(b)$, this follows by the monotonicity of b and the definition of $B(b)$: $x(s) = x(S) - x(S - s) \geq b(S) - b(S - s) \geq 0$ holds for every $s \in S$. The set of integral elements of $B(b)$ is called a *base polymatroidal set*. Edmonds [12] showed that the vertices of a polymatroid or a base polymatroid are integral, thus $P(b)$ is the convex hull of the corresponding polymatroidal set, while $B(b)$ is the convex hull of the corresponding base polymatroidal set. For this reason, we will call the sets of integral elements of $P(b)$ and $B(b)$ simply a polymatroid and a base polymatroid.

Hassin [24] introduced polyhedra bounded simultaneously by a non-negative, monotone non-decreasing submodular function b over a ground set S from above and by a non-negative, monotone non-decreasing supermodular function p over S from below, satisfying the so-called *cross-inequality* linking the two functions:

$$b(X) - p(Y) \geq b(X - Y) - p(Y - X) \quad \text{for every pair of subsets } X, Y \subseteq S .$$

We say that a pair (p, b) of set functions over the same ground set S is a *paramodular pair* if $p(\emptyset) = b(\emptyset) = 0$, p is supermodular, b is submodular, and they satisfy the cross-inequality. The slightly more general concept of generalized polymatroids was introduced by Frank [16]. A *generalized polymatroid*, or *g -polymatroid* is a polyhedron of the form

$$Q(p, b) := \{x \in \mathbb{R}^S \mid p(Y) \leq x(Y) \leq b(Y) \text{ for every } Y \subseteq S\} ,$$

where (p, b) is a paramodular pair. Here, (p, b) is called the *border pair* of the polyhedron. It is known (see e.g. [17]) that a g-polymatroid defined by an integral paramodular pair is a non-empty integral polyhedron.

A special g-polymatroid is a box $\beta(L, U) = \{x \in \mathbb{R}^S \mid L \leq x \leq U\}$ where $L : S \rightarrow \mathbb{Z} \cup \{-\infty\}$, $U : S \rightarrow \mathbb{Z} \cup \{\infty\}$ with $L \leq U$. Another illustrious example is base polymatroids. Indeed, given a polymatroid function b with finite $b(S)$, its *complementary set function* p is defined for $X \subseteq S$ by $p(X) := b(S) - b(S - X)$. It is not difficult to check that (p, b) is a paramodular pair and that $B(b) = Q(p, b)$.

Theorem 14.3.9 (Frank [17]). *The intersection Q' of a g-polymatroid $Q = Q(p, b)$ and a box $\beta = \beta(L, U)$ is a g-polymatroid. If $L(Y) \leq b(Y)$ and $p(Y) \leq U(Y)$ hold for every $Y \subseteq S$, then Q' is non-empty, and its unique border pair (p', b') is given by*

$$(15) \quad \begin{aligned} p'(Z) &= \max\{p(Z') - U(Z' - Z) + L(Z - Z') \mid Z' \subseteq S\} , \\ b'(Z) &= \min\{b(Z') - L(Z' - Z) + U(Z - Z') \mid Z' \subseteq S\} . \end{aligned}$$

Given a g-polymatroid $Q(p, b)$ and $Z \subset S$, by *deleting* $Z \subseteq S$ from $Q(p, b)$ we obtain a g-polymatroid $Q(p, b) \setminus Z$ defined on set $S - Z$ by the restrictions of p and b to $S - Z$, that is,

$$Q(p, b) \setminus Z := \{x \in \mathbb{R}^{S-Z} \mid p(Y) \leq x(Y) \leq b(Y) \text{ for every } Y \subseteq S - Z\} .$$

In other words, $Q(p, b) \setminus Z$ is the projection of $Q(p, b)$ to the coordinates in $S - Z$.

Extending the notion of contraction from matroids to g-polymatroids is not immediate. A set can be naturally identified with its characteristic vector, that is, in the case of matroids contraction is basically an operation defined on 0–1 vectors. In our proof, we will need a generalization of this to the integral elements of a g-polymatroid. However, such an element might have coordinates larger than one as well, hence finding the right definition is not straightforward. In the case of matroids, the most important property of contraction is the following: I is an independent of M/Z if and only if $F \cup I$ is independent in M for any maximal independent set F of Z .

With this property in mind, we define the g-polymatroid obtained by the contraction of an integral vector $z \in Q(p, b)$ to be the polymatroid $Q(p', b') := Q(p, b)/z$ on the same ground set S with the border functions

$$\begin{aligned} p'(X) &:= p(X) - z(X) \\ b'(X) &:= b(X) - z(X) . \end{aligned}$$

Observe that p' is obtained as the difference of a supermodular and a modular function, implying that it is supermodular. Similarly, b' is submodular. Moreover, $p'(\emptyset) = b'(\emptyset) = 0$, and

$$\begin{aligned} b'(X) - p'(Y) &= b(X) - z(X) - p(Y) + z(Y) \\ &\geq b(X - Y) + p(Y - X) - z(X - Y) + z(Y - X) \\ &= b'(X - Y) - p'(Y - X), \end{aligned}$$

hence (p', b') is indeed a paramodular pair. The main reason for defining the contraction of an element $z \in Q(p, b)$ is shown by the following lemma.

Lemma 15. *Let $Q(p', b')$ be the polymatroid obtained by contracting $z \in Q(p, b)$. Then $x + z \in Q(p, b)$ for every $x \in Q(p', b')$.*

Proof. Let $x \in Q(p', b')$. By definition, this implies $p'(Y) \leq x(Y) \leq b'(Y)$ for $Y \subseteq S$. Thus $p(Y) = p'(Y) + z(Y) \leq x(Y) + z(Y) \leq b'(Y) + z(Y) = b(Y)$, concluding the proof. \square

5.2 The algorithm

The aim of this section is to prove [Theorem 5](#) and [Theorem 6](#). [Theorem 5](#) extends the result of Király et al. [[33](#)] from matroids to g-polymatroids. However, adapting their algorithm is not immediate due to the following major differences. A crucial step of their approach is to relax the problem by deleting a constraint corresponding to a hyperedge ε with small $g(\varepsilon)$ value. This step is feasible when the solution is a 0–1 vector, but it is not applicable for g-polymatroids (or even for polymatroids) where an integral element might have coordinates larger than 1. This difficulty is compounded by the presence of multiplicity vectors, that makes both the computations and the tracking of changes after hyperedge deletions more complicated. Finally, in contrast to matroids that are defined by a submodular function (the rank function), g-polymatroids are determined by a pair of supermodular and submodular functions. Thus the structure of the family of tight sets is more complex, which affects the proof of one of the key claims ([Claim 16](#)).

We start by formulating a linear programming relaxation for the BOUNDED DEGREE G-POLYMATROID ELEMENT PROBLEM:

$$\begin{array}{ll}
 \text{minimize} & \sum_{s \in S} c(s) x(s) \\
 \text{(LP)} \quad \text{subject to} & p(Z) \leq x(Z) \leq b(Z) \quad \forall Z \subseteq S \\
 & f(\varepsilon) \leq \sum_{s \in \varepsilon} m_\varepsilon(s) x(s) \leq g(\varepsilon) \quad \forall \varepsilon \in \mathcal{E}
 \end{array}$$

Although the program has an exponential number of constraints, it can be separated in polynomial time using submodular minimization [[28,41,46](#)]. [Algorithm 3](#) generalizes the approach by Király et al. [[33](#)]. We iteratively solve the linear program, delete elements which get a zero value in the solution, update the solution values and perform a contraction on the polymatroid, or remove constraints arising from the hypergraph. In the first round, the bounds on the coordinates solely depend on p and b , while in the subsequent rounds the whole problem is restricted to the unit cube.

Theorem 5. *There is an algorithm for the BOUNDED DEGREE G-POLYMATROID ELEMENT WITH MULTIPLICITIES problem which returns an integral element x of $Q(p, b)$ of cost at most the optimum value such that $f(\varepsilon) - 2\Delta + 1 \leq \sum_{s \in \varepsilon} m_\varepsilon(s)x(s) \leq g(\varepsilon) + 2\Delta - 1$ for each $\varepsilon \in \mathcal{E}$, where $\Delta = \max_{s \in S} \{\sum_{\varepsilon \in \mathcal{E}: s \in \varepsilon} m_\varepsilon(s)\}$. The run time of the algorithm is polynomial in n and $\log \sum_\varepsilon (f(\varepsilon) + g(\varepsilon))$.*

Proof. Our algorithm is presented as [Algorithm 1](#).

Algorithm 3 Approximation algorithm for the BOUNDED DEGREE G-POLYMATROID ELEMENT WITH MULTIPLICITIES problem.

Input: A g-polymatroid $Q(p, b)$ on ground set S , cost function $c : S \rightarrow \mathbb{R}$, a hypergraph $H = (S, \mathcal{E})$, lower and upper bounds $f, g : \mathcal{E} \rightarrow \mathbb{Z}_{\geq 0}$, multiplicities $m_\varepsilon : S \rightarrow \mathbb{Z}_{\geq 0}$ for $\varepsilon \in \mathcal{E}$ satisfying $m_\varepsilon(s) = 0$ for $s \in S - \varepsilon$.

Output: $z \in Q(p, b)$ of cost at most OPT_{LP} , violating the hyperedge constraints by at most $2\Delta - 1$.

- 1: Initialize $z(s) \leftarrow 0$ for every $s \in S$.
 - 2: **while** $S \neq \emptyset$ **do**
 - 3: Compute a basic optimal solution x for (LP).
 (Note: starting from the second iteration, $0 \leq x \leq 1$.)
 a: Delete any element s with $x(s) = 0$. Update each hyperedge $\varepsilon \leftarrow \varepsilon - s$ and $m_\varepsilon(s) \leftarrow 0$. Update the g-polymatroid $Q(p, b) \leftarrow Q(p, b) \setminus s$ by deletion.
 b: For all $s \in S$ update $z(s) \leftarrow z(s) + \lfloor x \rfloor(s)$.
 Apply polymatroid contraction $Q(p, b) \leftarrow Q(p, b) / \lfloor x \rfloor$, that is, redefine $p(Y) := p(Y) - \lfloor x \rfloor(Y)$ and $b(Y) := b(Y) - \lfloor x \rfloor(Y)$ for every $Y \subseteq S$.
 Update $f(\varepsilon) \leftarrow f(\varepsilon) - \sum_{s \in \varepsilon} m_\varepsilon(s) \lfloor x \rfloor(s)$ and $g(\varepsilon) \leftarrow g(\varepsilon) - \sum_{s \in \varepsilon} m_\varepsilon(s) \lfloor x \rfloor(s)$ for each $\varepsilon \in \mathcal{E}$.
 c: If $m_\varepsilon(\varepsilon) \leq 2\Delta - 1$, let $\mathcal{E} \leftarrow \mathcal{E} - \varepsilon$.
 d: **if** it is the first iteration **then**
 Take the intersection of $Q(p, b)$ and the unit cube $[0, 1]^S$, that is, $p(Y) := \max\{p(Y') - |Y' - Y| \mid Y' \subseteq S\}$ and $b(Y) := \min\{b(Y') + |Y - Y'| \mid Y' \subseteq S\}$ for every $Y \subseteq S$.
 - 4: **return** z
-

Correctness. First we show that if the algorithm terminates then the returned solution z satisfies the requirements of the theorem. In a single iteration, the g-polymatroid $Q(p, b)$ is updated to $(Q(p, b) \setminus D) / \lfloor x \rfloor$, where $D = \{s : x(s) = 0\}$ is the set of deleted elements. In the first iteration, the g-polymatroid thus obtained is further intersected with the unit cube. By Lemma 15, the vector $x - \lfloor x \rfloor$ restricted to $S - D$ remains a feasible solution for the modified linear program in the next iteration. Note that this vector is contained in the unit cube as its coordinates are between 0 and 1. This remains true when a lower degree constraint is removed in Step 3.c as well, therefore the cost of z plus the cost of an optimal LP solution does not increase throughout the procedure. Hence the cost of the output z is at most the cost of the initial LP solution, which is at most the optimum.

By Lemma 15, the vector $x - \lfloor x \rfloor + z$ is contained in the original g-polymatroid, although it might violate some of the lower and upper bounds on the hyperedges. We only remove the constraints corresponding to the lower and upper bounds for a hyperedge ε when $m_\varepsilon(\varepsilon) \leq 2\Delta - 1$. As the g-polymatroid is restricted to the unit cube after the first iteration, these constraints are violated by at most $2\Delta - 1$, as the total value of $\sum_{s \in \varepsilon} m_\varepsilon(s)z(s)$ can change by a value between 0 and $2\Delta - 1$ in the remaining iterations.

It remains to show that the algorithm terminates successfully. The proof is based on similar arguments as in Király et al. [33, proof of Theorem 2].

Termination. Suppose, for sake of contradiction, that the algorithm does not terminate. Then there is some iteration after which none of the simplifications in Steps 3.a to 3.c can be performed. This implies that for the current basic LP solution x it holds $0 < x(s) < 1$ for each $s \in S$ and $m_\varepsilon(\varepsilon) \geq 2\Delta$ for each $\varepsilon \in \mathcal{E}$. We say that a set Y is *p-tight* (or *b-tight*) if $x(Y) = p(Y)$ (or $x(Y) = b(Y)$), and let $\mathcal{T}^p = \{Y \subseteq S : x(Y) = p(Y)\}$ and $\mathcal{T}^b = \{Y \subseteq S : x(Y) = b(Y)\}$ denote the collections of *p-tight* and *b-tight* sets with respect to solution x .

Let \mathcal{L} be a maximal independent laminar system in $\mathcal{T}^p \cup \mathcal{T}^b$.

Claim 16. $\text{span}(\{\chi_Z \mid Z \in \mathcal{L}\}) = \text{span}(\{\chi_Z \mid Z \in \mathcal{T}^p \cup \mathcal{T}^b\})$

Proof of Claim 16. The proof uses an uncrossing argument. Let us suppose indirectly that there is a set R from $\mathcal{T}^p \cup \mathcal{T}^b$ for which $\chi_R \notin \text{span}(\{\chi_Z \mid Z \in \mathcal{L}\})$. Choose this set R so that it is incomparable to as few sets of \mathcal{L} as possible. Without loss of generality, we may assume that $R \in \mathcal{T}^p$. Now choose a set $T \in \mathcal{L}$ that is incomparable to R . Note that such a set necessarily exists as the laminar system is maximal. We distinguish two cases.

Case 1. $T \in \mathcal{T}^p$. Because of the supermodularity of p , we have

$$\begin{aligned} x(R) + x(T) &= p(R) + p(T) \leq p(R \cup T) + p(R \cap T) \leq x(R \cup T) + x(R \cap T) \\ &= x(R) + x(T), \end{aligned}$$

hence equality holds throughout. That is, $R \cup T$ and $R \cap T$ are in \mathcal{T}^p as well. In addition, since $\chi_R + \chi_T = \chi_{R \cup T} + \chi_{R \cap T}$ and χ_R is not in $\text{span}(\{\chi_Z \mid Z \in \mathcal{L}\})$, either $\chi_{R \cup T}$ or $\chi_{R \cap T}$ is not contained in $\text{span}(\{\chi_Z \mid Z \in \mathcal{L}\})$. However, both $R \cup T$ and $R \cap T$ are incomparable with fewer sets of \mathcal{L} than R , which is a contradiction.

Case 2. $T \in \mathcal{T}^b$. Because of the cross-inequality, we have

$$\begin{aligned} x(T) - x(R) &= b(T) - p(R) \geq b(T \setminus R) - p(R \setminus T) \geq x(T \setminus R) - x(R \setminus T) \\ &= x(T) - x(R), \end{aligned}$$

implying $T \setminus R \in \mathcal{T}^b$ and $R \setminus T \in \mathcal{T}^p$. Since $\chi_R + \chi_T = \chi_{R \setminus T} + \chi_{R \cap T} + 2\chi_{R \cup T}$ and χ_R is not in $\text{span}(\{\chi_Z \mid Z \in \mathcal{L}\})$, one of the vectors $\chi_{R \setminus T}$, $\chi_{R \cap T}$ and $\chi_{R \cup T}$ is not contained in $\text{span}(\{\chi_Z \mid Z \in \mathcal{L}\})$. However, any of these three sets is incomparable with fewer sets of \mathcal{L} than R , which is a contradiction.

The case when $R \in \mathcal{T}^b$ is analogous to the above. This completes the proof of the Claim. \diamond

We say that a hyperedge $\varepsilon \in \mathcal{E}$ is *tight* if $f(\varepsilon) = \sum_{s \in \varepsilon} m_\varepsilon(s)x(s)$ or $g(\varepsilon) = \sum_{s \in \varepsilon} m_\varepsilon(s)x(s)$. As x is a basic solution, there is a set $\mathcal{E}' \subseteq \mathcal{E}$ of tight hyperedges such that $\{m_\varepsilon \mid \varepsilon \in \mathcal{E}'\} \cup \{\chi_Z \mid Z \in \mathcal{L}\}$ are linearly independent vectors with $|\mathcal{E}'| + |\mathcal{L}| = |S|$.

We derive a contradiction using a token-counting argument. We assign 2Δ tokens to each element $s \in S$, accounting for a total of $2\Delta|S|$ tokens. The tokens are then redistributed in such a way that each hyperedge in \mathcal{E}' and each set in \mathcal{L} collects at least 2Δ tokens, while at least one extra token remains. This implies that $2\Delta|S| > 2\Delta|\mathcal{E}'| + 2\Delta|\mathcal{L}|$, leading to a contradiction.

We redistribute the tokens as follows. Each element s gives Δ tokens to the smallest member in \mathcal{L} it is contained in, and $m_\varepsilon(s)$ tokens to each hyperedge $\varepsilon \in \mathcal{E}'$ it is contained in. As $\sum_{\varepsilon \in \mathcal{E}': s \in \varepsilon} m_\varepsilon(s) \leq \Delta$ holds for every element $s \in S$, thus we redistribute at most 2Δ tokens per element and so the redistribution step is valid. Now consider any set $U \in \mathcal{L}$. Recall that $\mathcal{L}^{\max}(U)$ consists of the maximal members of \mathcal{L} lying inside U . Then $U - \bigcup_{W \in \mathcal{L}^{\max}(U)} W \neq \emptyset$, as otherwise $\chi_U = \sum_{W \in \mathcal{L}^{\max}(U)} \chi_W$, contradicting the independence of \mathcal{L} . For every set Z in \mathcal{L} , $x(Z)$ is an integer, meaning that $x(U - \bigcup_{W \in \mathcal{L}^{\max}(U)} W)$ is an integer. But also $0 < x(s) < 1$ for every $s \in S$, which means that $U - \bigcup_{W \in \mathcal{L}^{\max}(U)} W$ contains at least 2 elements. Therefore, each set U in \mathcal{L} receives at least 2Δ tokens, as required. By assumption, $m_\varepsilon(\varepsilon) \geq 2\Delta$ for every hyperedge $\varepsilon \in \mathcal{E}'$, which means that each hyperedge in \mathcal{E}' receives at least 2Δ tokens, as required.

If $\sum_{\varepsilon \in \mathcal{E}': s \in \varepsilon} m_\varepsilon(s) \leq \Delta$ holds for any $s \in S$ or $\mathcal{L}^{\max}(S)$ is not a partition of S , then an extra token exists. Otherwise, $\sum_{\varepsilon \in \mathcal{E}'} m_\varepsilon = \Delta \cdot \chi_S = \Delta \cdot \sum_{W \in \mathcal{L}^{\max}(S)} \chi_W$, contradicting the independence of $\{m_\varepsilon \mid \varepsilon \in \mathcal{E}'\} \cup \{\chi_Z \mid Z \in \mathcal{L}\}$.

Time complexity. Solving an LP, as well as removing a hyperedge in [Step 3.a](#) or removing an element from a hyperedge in [Step 3.c](#) can be done in polynomial time. In [Steps 3.b](#) and [3.d](#), we calculate the value of the current functions p and b for a set Y only when it is needed during the ellipsoid method. We keep track of the vectors $[x]$ that arise during contraction steps (there is only a polynomial number of them), and every time a query for p or b happens, it takes into account every contraction and removal that occurred until that point.

[Step 3.a](#) can be repeated at most $|S|$ times, while [Step 3.c](#) can be repeated at most $|\mathcal{E}|$ times. Starting from the second iteration, we are working in the unit cube. That is, when [Step 3.b](#) adds the integer part of a variable $x(s)$ to $z(s)$ and reduces the problem, then the given variable will be 0 in the next iteration and so element s is deleted. This means that the total number of iterations of [Step 3.b](#) is at most $\mathcal{O}(|S|)$. \square

Now we consider case when only lower or only upper bounds are given.

Theorem 6. *There is an algorithm for LOWER BOUNDED DEGREE G-POLYMATROID ELEMENT WITH MULTIPLICITIES which returns an integral element x of $Q(p, b)$ of cost at most the optimum value such that $f(\varepsilon) - \Delta + 1 \leq \sum_{s \in \varepsilon} m_\varepsilon(s)x(s)$ for each $\varepsilon \in \mathcal{E}$. An analogous result holds for UPPER BOUNDED DEGREE G-POLYMATROID ELEMENT, where $\sum_{s \in \varepsilon} m_\varepsilon(s)x(s) \leq g(\varepsilon) + \Delta - 1$. The run time of these algorithms is polynomial in n and $\log \sum_\varepsilon f(\varepsilon)$ or $\log \sum_\varepsilon g(\varepsilon)$, respectively.*

Proof. The proof is similar to the proof of [Theorem 5](#), the main difference appears in the counting argument. When only lower bounds are present, the condition in [Step 3.c](#) changes: we delete a hyperedge ε if $f(\varepsilon) \leq \Delta - 1$. Suppose, for the sake of contradiction, that the algorithm does not terminate. Then there is an iteration after which none of the simplifications in [Steps 3.a](#) to [3.c](#) can be performed. This implies that in the current basic solution $0 < x(s) < 1$ holds for each $s \in S$ and $f(\varepsilon) \geq \Delta$ for

each $\varepsilon \in \mathcal{E}$. We choose a subset $\mathcal{E}' \subseteq \mathcal{E}$ and a maximal independent laminar system \mathcal{L} of tight sets the same way as in the proof of [Theorem 5](#). Recall that $|\mathcal{E}'| + |\mathcal{L}| = |S|$.

Let Z_1, \dots, Z_k denote the members of the laminar system \mathcal{L} . As \mathcal{L} is an independent system, $Z_i - \bigcup_{W \in \mathcal{L}^{\max}(Z_i)} W \neq \emptyset$. Since $x(s) < 1$ for all $s \in S$,

$$x\left(Z_i - \bigcup_{W \in \mathcal{L}^{\max}(Z_i)} W\right) < \left|Z_i - \bigcup_{W \in \mathcal{L}^{\max}(Z_i)} W\right| .$$

As we have integers on both sides of this inequality, we get

$$\left|Z_i - \bigcup_{W \in \mathcal{L}^{\max}(Z_i)} W\right| - x\left(Z_i - \bigcup_{W \in \mathcal{L}^{\max}(Z_i)} W\right) \geq 1 \quad \text{for all } i = 1, \dots, k .$$

Moreover, $\sum_{s \in \varepsilon} m_\varepsilon(s)x(s) \geq f(\varepsilon) \geq \Delta$ for all hyperedges; therefore,

$$\begin{aligned} |\mathcal{E}'| + |\mathcal{L}| &\leq \sum_{\varepsilon \in \mathcal{E}'} \frac{\sum_{s \in \varepsilon} m_\varepsilon(s)x(s)}{\Delta} + \sum_{i=1}^k \left[\left|Z_i - \bigcup_{W \in \mathcal{L}^{\max}(Z_i)} W\right| - x\left(Z_i - \bigcup_{W \in \mathcal{L}^{\max}(Z_i)} W\right) \right] \\ &= \sum_{s \in S} \frac{x(s)}{\Delta} \sum_{\substack{\varepsilon \in \mathcal{E}' \\ s \in \varepsilon}} m_\varepsilon(s) + \sum_{W \in \mathcal{L}^{\max}(S)} |W| - \sum_{W \in \mathcal{L}^{\max}(S)} x(W) \leq |S| . \end{aligned}$$

In the last line, the first term is at most $x(S)$ since $\sum_{\varepsilon \in \mathcal{E}; s \in \varepsilon} m_\varepsilon(s) \leq \Delta$ holds for each element $s \in S$. From $x(S) - \sum_{W \in \mathcal{L}^{\max}(S)} x(W) \leq |S| - \sum_{W \in \mathcal{L}^{\max}(S)} |W|$ the upper bound of $|S|$ follows. As $|S| = |\mathcal{L}| + |\mathcal{E}'|$, we have equality throughout. This implies that

$$\sum_{\varepsilon \in \mathcal{E}'} m_\varepsilon = \Delta \cdot \chi_S = \Delta \cdot \sum_{W \in \mathcal{L}^{\max}(S)} \chi_W ,$$

contradicting linear independence.

If only upper bounds are present, we remove a hyperedge ε in [Step 3.c](#) when $g(\varepsilon) + \Delta - 1 \geq m_\varepsilon(\varepsilon)$. Suppose, for the sake of contradiction, that the algorithm does not terminate. Then there is an iteration after which none of the simplifications in [Steps 3.a](#) to [3.c](#) can be performed. This implies that in the current basic solution $0 < x(s) < 1$ holds for each $s \in S$ and $m_\varepsilon(\varepsilon) - g(\varepsilon) \geq \Delta$ for each $\varepsilon \in \mathcal{E}$. Again, we choose a subset $\mathcal{E}' \subseteq \mathcal{E}$ and a maximal independent laminar system \mathcal{L} of tight sets the same way as in the proof of [Theorem 5](#).

Let Z_1, \dots, Z_k denote the members of the laminar system \mathcal{L} . As \mathcal{L} is an independent system, $Z_i - \bigcup_{W \in \mathcal{L}^{\max}(Z_i)} W \neq \emptyset$ and so

$$x\left(Z_i - \bigcup_{W \in \mathcal{L}^{\max}(Z_i)} W\right) \geq 1 .$$

By $\sum_{s \in \varepsilon} m_\varepsilon(s)x(s) \leq g(\varepsilon)$, we get $\sum_{s \in \varepsilon} m_\varepsilon(s) - \sum_{s \in \varepsilon} m_\varepsilon(s)x(s) \geq m_\varepsilon(\varepsilon) - g(\varepsilon) \geq \Delta$. Thus,

$$\begin{aligned} |\mathcal{E}'| + |\mathcal{L}| &\leq \sum_{\varepsilon \in \mathcal{E}'} \frac{\sum_{s \in \varepsilon} m_\varepsilon(s) - \sum_{s \in \varepsilon} m_\varepsilon(s)x(s)}{\Delta} + \sum_{i=1}^k x(Z_i - \bigcup_{W \in \mathcal{L}^{\max}(Z_i)} W) \\ &= \sum_{s \in S} \frac{1 - x(s)}{\Delta} \sum_{\substack{\varepsilon \in \mathcal{E}' \\ s \in \varepsilon}} m_\varepsilon(s) + \sum_{W \in \mathcal{L}^{\max}(S)} x(W) \\ &\leq \sum_{s \in S} \frac{1 - x(s)}{\Delta} \sum_{\substack{\varepsilon \in \mathcal{E}' \\ s \in \varepsilon}} m_\varepsilon(s) + x(S) \leq |S|. \end{aligned}$$

In the last line, the first term is at most $|S| - x(S)$ since $\sum_{\varepsilon \in \mathcal{E}': s \in \varepsilon} m_\varepsilon(s) \leq \Delta$ holds for every element $s \in S$. Therefore, the upper bound of $|S|$ follows. As $|S| = |\mathcal{L}| + |\mathcal{E}'|$, we have equality throughout. This implies that $\sum_{\varepsilon \in \mathcal{E}'} m_\varepsilon = \Delta \cdot \chi_S = \Delta \cdot \sum_{W \in \mathcal{L}^{\max}(S)} \chi_W$, contradicting linear independence. \square

Remark 4. Note that [Theorems 5 and 6](#) only provide a solution if there exists a (fractional) solution to the underlying linear program in [\(LP\)](#). Consequently, [Theorem 4](#) only provides a solution if the polytope P^{CG} in [Equation \(1\)](#) is not empty.

We have seen in [Section 5.1](#) that base polymatroids are special cases of g-poq-lymatroids. This implies that the results of [Theorem 6](#) immediately apply to polymatroids. In the LOWER BOUNDED DEGREE POLYMATROID BASIS WITH MULTIPLICITIES problem, we are given a base polymatroid $B(b) = (S, b)$ with a cost function $c : S \rightarrow \mathbb{R}$, and a hypergraph $H = (S, \mathcal{E})$ on the same ground set. The input contains lower bounds $f : \mathcal{E} \rightarrow \mathbb{Z}_{\geq 0}$ and multiplicity vectors $m_\varepsilon : \varepsilon \rightarrow \mathbb{Z}_{\geq 1}$ for every hyperedge $\varepsilon \in \mathcal{E}$. The objective is to find a minimum-cost element $x \in B(b)$ such that $f(\varepsilon) \leq \sum_{s \in \varepsilon} m_\varepsilon(s)x(s)$ holds for each $\varepsilon \in \mathcal{E}$.

Corollary 17. *There is a polynomial-time algorithm for the LOWER BOUNDED DEGREE POLYMATROID BASIS WITH MULTIPLICITIES problem which returns an integral element x of $B(b)$ of cost at most the optimum value such that $f(\varepsilon) - \Delta + 1 \leq \sum_{s \in \varepsilon} m_\varepsilon(s)x(s)$ for each $\varepsilon \in \mathcal{E}$.*

5.3 Proof of Theorem 4

In this section we show that [Algorithm 3](#) can be applied in order to obtain an approximation to the MINIMUM BOUNDED DEGREE CONNECTED MULTIGRAPH WITH EDGE BOUNDS problem, as described in [Theorem 4](#).

Theorem 4. *There is an algorithm for the MINIMUM BOUNDED DEGREE CONNECTED MULTIGRAPH WITH EDGE BOUNDS problem that, in time polynomial in n and $\log \sum_v \rho(v)$, returns a connected multigraph T with $\rho(V)/2$ edges, where each vertex v has degree at least $\rho(v) - 1$ and the cost of T is at most the cost of $\min\{c^\top x \mid x \in$*

$P^{\text{CG}}(\rho, L, U)\}$, where

$$(1) \quad P^{\text{CG}}(\rho, L, U) := \left\{ x \in \mathbb{R}_{\geq 0}^E \left| \begin{array}{ll} \text{supp}(x) \text{ is connected} & \\ x(E) = \rho(V)/2 & \\ x(\delta(v)) \geq \rho(v) & \forall v \in V \\ L(vw) \leq x(vw) \leq U(vw) & \forall v, w \in V \end{array} \right. \right\}.$$

Let us take a MINIMUM BOUNDED DEGREE CONNECTED MULTIGRAPH WITH EDGE BOUNDS problem instance (G, c, ρ, L, U) on a graph $G(V, E)$, where c, ρ, L, U are non-negative and $\rho(V) = \sum_{v \in V} \rho(v)$ is even.⁶ Note that we do not require c to satisfy the triangle inequality. We start with defining the specific input variables passed over Algorithm 3. Then, we show that given the specified input, the algorithm yields an approximate solution to the MINIMUM BOUNDED DEGREE CONNECTED MULTIGRAPH WITH EDGE BOUNDS problem. From now on we use $\hat{\rho} = \rho(V)/2 - |V| + 1$.

We first set the base set S as the edge set E of our original graph G . In the hypergraph $H = (S, \mathcal{E})$, the elements of S thus correspond to the edges of G . Moreover, there is a hyperedge ε for every vertex in V , defined the following way: $\mathcal{E} := \{\delta(v) \mid v \in V\}$. The multiplicity of an element s in a hyperedge ε is 1, that is, $m_\varepsilon(s) := 1$ if s corresponds to a regular edge $e \in E$, and $m_\varepsilon(s) := 2$ if s corresponds to a self-loop. We set the lower bound f for a hyperedge ε according to the degree requirement of the corresponding vertex v , that is $f(\varepsilon) := \rho(v)$.

We now define the second input of Algorithm 3, a g-polymatroid $Q(S, p, b)$. This is done in two steps, by first defining an auxiliary polymatroid $Q'(S, p', b')$, then taking the intersection of the g-polymatroid Q' with a box. We define the border function p' as the zero vector on S , and $b'(Z)$ as follows:

Lemma 18. *Let b' denote the following function defined on sets $Z \subseteq S$:*

$$(16) \quad b'(Z) = \begin{cases} |V(Z)| - \text{comp}(Z) + \hat{\rho}, & \text{if } Z \neq \emptyset, \\ 0, & \text{otherwise} \end{cases}.$$

Then b' is a polymatroid function.

Proof. By definition, $b'(\emptyset) = 0$ and b' is monotone increasing. It remains to show that b' is submodular. Let $X, Y \subseteq S$. The submodular inequality clearly holds if one of X and Y is empty. If none of X and Y is empty then the submodular inequality follows from the fact that $|V(Z)| - \text{comp}(Z)$ is the rank function of the graphical matroid. \square

Consider the g-polymatroid $B(p', b')$ determined by the border functions defined in Equation (16). Let us define the set

$$B = \{x \in \mathbb{Z}_{\geq 0}^E : x(E) = \rho(V)/2, \text{ supp}(x) \text{ is connected}\}.$$

Lemma 19. $B = B(p', b') \cap \mathbb{Z}_{\geq 0}^E$.

⁶Due to the handshaking lemma, the sum of degrees in a graph is even, therefore $\rho(V)$ being even is necessary.

Proof. Take an integral element $x \in B(p', b')$ and let $C \subseteq E$ be an arbitrary cut between V_1 and V_2 for some partition $V_1 \uplus V_2$ of V . Then

$$\begin{aligned} x(C) &= x(E) - x(E(V_1) \cup E(V_2)) \\ &\geq |V| - 1 + \hat{\rho} - (|V_1| + |V_2| - \text{comp}(E(V_1) \cup E(V_2)) + \hat{\rho}) \\ &\geq 1, \end{aligned}$$

thus $\text{supp}(x)$ is connected. As $x(E) = |V| - 1 + \hat{\rho} = \rho(V)/2$, we obtain $x \in B$, showing that $B(p', b') \subseteq B$.

To see the other direction, take an element $x \in B$. As $\text{supp}(x)$ is connected, $x(E - F) \geq \text{comp}(F) + |V| - |V(F)| - 1$ for every $F \subseteq E$. That is,

$$\begin{aligned} x(F) &= x(E) - x(E - F) \\ &\leq r(V) - (|V - V(F)| + \text{comp}(F) - 1) \\ &= |V(F)| - \text{comp}(F) + \hat{\rho}, \end{aligned}$$

thus $x(F) \leq b'(F)$. As $x(E) = r(V) = |V| - 1 + \hat{\rho}$, we obtain $x \in B(p', b')$, showing $B \subseteq B(p', b')$. \square

So far we proved that the integral points of $Q'(S, p', b')$ correspond to a connected multigraph on V that has $\rho(V)/2$ edges. Let $\beta(L, U)$ be the box defined by

$$\beta(L, U) := \{x \in \mathbb{R}_{\geq 0}^S \mid L(s) \leq x(s) \leq U(s) \ \forall s \in S\} .$$

Let us define the polymatroid $Q = (S, p, b)$ as the intersection of the polymatroid Q' and the box β , where the border functions p, b are defined as in Equation (15). We now prove that taking $H = (S, \mathcal{E})$ and $Q(S, p, b)$ as input, the output of Algorithm 3 corresponds to a multigraph with the properties stated in Theorem 4.

Proof of Theorem 4. Consider the linear program (LP) that is defined in the iterative rounding method for the G-POLYMATROID ELEMENT WITH MULTIPLICITIES problem. The constraints regarding the bounds on the hyperedges imply $\rho(v) \leq x(\delta(v))$ for every $v \in V$: note that $m_\varepsilon(s) = 2$ for self loops and 1 for simple edges, and this equals to the contribution of an edge uv to the value $x(\delta(v))$. This, together with Lemma 19 and the fact that x is contained in the box $\beta(L, U)$, implies that Algorithm 3 returns an integral solution z such that the cost of z is at most the minimum cost element of P^{CG} .

According to Theorem 6, the integral solution z violates the bounds f on the hyperedges by at most $\Delta - 1$, where $\Delta := \max_{s \in S} \{\sum_{\varepsilon \in \mathcal{E}: s \in \varepsilon} m_\varepsilon(s)\}$. But we defined $m_\varepsilon(s)$ to be equal to 2 if s corresponds to a self-loop in G and 1 otherwise, meaning that the solution z violates the bounds on the hyperedges f and thus the bounds on the vertices ρ by at most 1. The solution z is also connected, with a total number of edges $\rho(V)/2$ and z satisfies the edge bounds L, U ; due to Theorem 6. Therefore, the solution z corresponds to a multigraph, that admits the properties in the claim of Theorem 4. \square

5.4 Strongly polynomial time implementation

The transportation problem in [Algorithm 1](#) can be solved in strongly polynomial time [[34, 43](#)]. Computing compact path-cycle representations uses the algorithm of Grigoriev and van de Klundert [[23](#)]; which, along with computing the Eulerian trail and making shortcuts, can be done in strongly polynomial time. Moreover, the algorithm uses the $3/2$ -approximation for the PATH TSP by Zenklusen [[58](#)] as a black-box, which can also be implemented in strongly polynomial time. Making shortcuts in [Algorithm 1](#) can be performed in strongly polynomial time as well, thus we can find a $5/2$ -approximation for the metric MANY-VISITS PATH TSP in strongly polynomial time.

[Algorithm 2](#) involves solving three types of LPs. According to §58.5 in Schrijver’s book [[47](#)], if the feasibility of a linear program for a vector x can be tested in polynomial time, then the ellipsoid method can find a solution in strongly polynomial time.

In [Step 1](#), we calculate an optimal solution to P_{HK}^{MV} as defined in [Equation \(2\)](#), while a number of linear programs of form $LP(a)$ arise throughout the dynamic program in [Step 2](#). The feasibility of the cut constraints can be checked in strongly polynomial time, by solving a minimum cut problem. The number of degree constraints in both types of these LPs and the number of constraints $x(\delta(B)) \geq 3$ in $(LP(a))$ is polynomial in n . Finally, in (LP) , the number of constraints involving hyperedges is polynomial in n , and one can check the feasibility of the constraints involving the border functions using submodular minimization [[29](#)]. This means all of the linear programs arising in [Algorithm 2](#) can be solved in strongly polynomial time.

According to [Lemma 10](#), the number of cuts in \mathcal{B} is polynomial in n , hence [Steps 1, 2](#) and [4](#) can be performed in strongly polynomial time. This is true for computing a matching in [Step 5](#), as well as all the remaining graph operations, using the same arguments as in case of [Algorithm 1](#). Therefore we provide a $3/2$ -approximation for the metric MANY-VISITS PATH TSP in strongly polynomial time.

6 Discussion

In this paper we gave an approximation algorithm for a far-reaching generalization of the metric PATH TSP, the metric MANY-VISITS PATH TSP where each city v has a (potentially exponentially large) requirement $r(v) \geq 1$. Our algorithm yields a $3/2$ -approximation for the metric MANY-VISITS PATH TSP in time polynomial in the number n of cities and the logarithm of the $r(v)$ ’s. It therefore generalizes the recent fundamental result by Zenklusen [[58](#)], who obtained a $3/2$ -approximation for the metric PATH TSP, finishing a long history of research.

At the heart of our algorithm is the first polynomial-time approximation algorithm for the minimum-cost degree bounded g -polymatroid element with multiplicities problem. That algorithm yields a solution of cost at most the optimum, which violates the lower bounds only by a constant factor depending on the weighted maximum element frequency Δ .

Finally, we show a simple approach, that gives a $5/2$ -approximation for the metric MANY-VISITS TSP in strongly polynomial time, and an $\mathcal{O}(1)$ -approximation for the metric MANY-VISITS ATSP in polynomial time.

Acknowledgements. The authors are grateful to Rico Zenklusen for discussions on techniques to obtain a $3/2$ -approximation for the metric version of the MANY-VISITS TSP, and to Tamás Király and Gyula Pap for their suggestions. Kristóf Bérczi was supported by the János Bolyai Research Fellowship of the Hungarian Academy of Sciences and by the ÚNKP-19-4 New National Excellence Program of the Ministry for Innovation and Technology. Projects no. NKFI-128673 and no. ED_18-1-2019-0030 (Application-specific highly reliable IT solutions) have been implemented with the support provided from the National Research, Development and Innovation Fund of Hungary, financed under the FK_18 and the Thematic Excellence Programme funding schemes, respectively.

References

- [1] A. Allahverdi, C. T. Ng, T. C. E. Cheng, and M. Y. Kovalyov. A survey of scheduling problems with setup times or costs. *European J. Oper. Res.*, 187(3):985–1032, 2008.
- [2] H.-C. An, R. Kleinberg, and D. B. Shmoys. Improving Christofides’ algorithm for the s - t path TSP. *J. ACM*, 62(5):Art. 34, 28, 2015.
- [3] J. E. Beasley, M. Krishnamoorthy, Y. M. Sharaiha, and D. Abramson. Scheduling aircraft landings—the static case. *Transp. Sci.*, 34(2):180–197, May 2000.
- [4] K. Bérczi, A. Berger, M. Mních, and R. Vincze. Degree-bounded generalized polymatroids and approximating the metric many-visits TSP. Technical report, 2019. <https://arxiv.org/abs/1911.09890>.
- [5] A. Berger, L. Kozma, M. Mních, and R. Vincze. Time- and space-optimal algorithm for the many-visits TSP. *ACM Trans. Algorithms*, 2020.
- [6] L. Bianco, P. Dell’Olmo, and S. Giordani. Minimizing total completion time subject to release dates and sequence-dependent processing times. *Ann. Oper. Res.*, 86:393–415, 1999.
- [7] K. Chaudhuri, S. Rao, S. Riesenfeld, and K. Talwar. A push-relabel approximation algorithm for approximating the minimum-degree MST problem and its generalization to matroids. *Theoret. Comput. Sci.*, 410(44):4489–4503, 2009.
- [8] K. Chaudhuri, S. Rao, S. Riesenfeld, and K. Talwar. What would Edmonds do? Augmenting paths and witnesses for degree-bounded MSTs. *Algorithmica*, 55(1):157–189, 2009.
- [9] N. Christofides. Worst-case analysis of a new heuristic for the travelling salesman problem. Technical Report 388, Carnegie Mellon University, 1976.

-
- [10] S. S. Cosmadakis and C. H. Papadimitriou. The traveling salesman problem with many visits to few cities. *SIAM J. Comput.*, 13(1):99–108, 1984.
- [11] M. A. Deppert and K. Jansen. Near-linear approximation algorithms for scheduling problems with batch setup times. In *Proc. SPAA 2019*, 2019.
- [12] J. Edmonds. Submodular functions, matroids, and certain polyhedra. In *Combinatorial Structures and their Applications (Proc. Calgary Internat. Conf., Calgary, Alta., 1969)*, pages 69–87. Gordon and Breach, New York, 1970.
- [13] J. Edmonds and R. M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. In *Combinatorial Structures and their Applications (Proc. Calgary Internat. Conf., Calgary, Alta., 1969)*, pages 93–96. Gordon and Breach, New York, 1970.
- [14] U. Feige and M. Singh. Improved approximation ratios for traveling salesperson tours and paths in directed graphs. In *Proc. APPROX-RANDOM 2007*, pages 104–118. 2007.
- [15] H. Fleischner. *Eulerian graphs and related topics. Part 1. Vol. 2*, volume 50 of *Annals of Discrete Mathematics*. North-Holland Publishing Co., Amsterdam, 1991.
- [16] A. Frank. Generalized polymatroids. In *Finite and infinite sets*, pages 285–294. 1984.
- [17] A. Frank. *Connections in combinatorial optimization*, volume 38 of *Oxford Lecture Series in Mathematics and its Applications*. Oxford University Press, Oxford, 2011.
- [18] A. Frank and É. Tardos. Generalized polymatroids and submodular flows. *Math. Prog.*, 42(1-3):489–563, 1988.
- [19] M. Fürer and B. Raghavachari. Approximating the minimum-degree Steiner tree to within one of optimal. *J. Algorithms*, 17(3):409–423, 1994.
- [20] M. R. Garey and D. S. Johnson. *Computers and intractability*. W. H. Freeman and Co., San Francisco, Calif., 1979.
- [21] M. X. Goemans. Minimum bounded degree spanning trees. In *Proc. FOCS 2006*, pages 273–282, 2006.
- [22] C. Gottschalk and J. Vygen. Better s - t -tours by Gao trees. *Math. Prog.*, 172(1-2, Ser. B):191–207, 2018.
- [23] A. Grigoriev and J. van de Klundert. On the high multiplicity traveling salesman problem. *Discrete Optim.*, 3(1):50–62, 2006.
- [24] R. Hassin. Minimum cost flow with set-constraints. *Networks*, 12(1):1–21, 1982.
-

-
- [25] C. Hierholzer and C. Wiener. Ueber die Möglichkeit, einen Linienzug ohne Wiederholung und ohne Unterbrechung zu umfahren. *Math. Ann.*, 6(1):30–32, 1873.
- [26] F. L. Hitchcock. The distribution of a product from several sources to numerous localities. *J. Math. Phys. Mass. Inst. Tech.*, 20:224–230, 1941.
- [27] J. A. Hoogeveen. Analysis of Christofides’ heuristic: some paths are more difficult than cycles. *Oper. Res. Lett.*, 10(5):291–295, 1991.
- [28] S. Iwata, L. Fleischer, and S. Fujishige. A combinatorial strongly polynomial algorithm for minimizing submodular functions. *J. ACM*, 48(4):761–777, 2001.
- [29] S. Iwata, L. Fleischer, and S. Fujishige. A combinatorial strongly polynomial algorithm for minimizing submodular functions. *J. ACM*, 48(4):761–777, 2001.
- [30] K. Jain. A factor 2 approximation algorithm for the generalized Steiner network problem. *Combinatorica*, 21(1):39–60, 2001.
- [31] K. Jansen, K.-M. Klein, M. Maack, and M. Rau. Empowering the configuration-IP—new PTAS results for scheduling with setups times. In *Proc. ITCS 2019*, volume 124 of *Leibniz Int. Proc. Inform.*, pages Art. No. 44, 19. 2019.
- [32] D. R. Karger. Global min-cuts in RNC, and other ramifications of a simple min-cut algorithm. In *Proc. SODA 1993*, pages 21–30, 1993.
- [33] T. Király, L. C. Lau, and M. Singh. Degree bounded matroids and submodular flows. *Combinatorica*, 32(6):703–720, 2012.
- [34] P. Kleinschmidt and H. Schannath. A strongly polynomial algorithm for the transportation problem. *Math. Program.*, 68:1–13, 1995.
- [35] J. Könemann and R. Ravi. A matter of degree: Improved approximation algorithms for degree-bounded minimum spanning trees. *SIAM J. Comput.*, 31(6):1783–1793, 2002.
- [36] J. Könemann and R. Ravi. Primal-dual meets local search: approximating MST’s with nonuniform degree bounds. In *Proc. STOC 2003*, pages 389–395, 2003.
- [37] Ł. Kowalik, S. Li, W. Nadara, M. Smulewicz, and M. Wahlström. Many visits TSP revisited. In *Proc. ESA 2020*.
- [38] L. Kozma and T. Mömke. Maximum scatter TSP in doubling metrics. In *Proc. SODA 2017*, pages 143–153, 2017.
- [39] M. Lampis. Algorithmic meta-theorems for restrictions of treewidth. *Algorithmica*, 64(1):19–37, 2012.

-
- [40] A. Lieder, D. Briskorn, and R. Stolletz. A dynamic programming approach for the aircraft landing problem with aircraft classes. *EJOR*, 243(1):61–69, May 2015.
- [41] S. T. McCormick. Submodular function minimization. *Handbooks in operations research and management science*, 12:321–391, 2005.
- [42] H. Nagamochi, K. Nishimura, and T. Ibaraki. Computing all small cuts in an undirected network. *SIAM J. Discrete Math.*, 10(3):469–481, 1997.
- [43] J. B. Orlin. A faster strongly polynomial minimum cost flow algorithm. *Oper. Res.*, 41(2):338–350, 1993.
- [44] H. N. Psaraftis. A dynamic programming approach for sequencing groups of identical jobs. *Oper. Res.*, 28(6):1347–1359, 1980.
- [45] M. Rothkopf. Letter to the editor—the traveling salesman problem: On the reduction of certain large problems to smaller ones. *Oper. Res.*, 14(3):532–533, 1966.
- [46] A. Schrijver. A combinatorial algorithm minimizing submodular functions in strongly polynomial time. *J. Combinatorial Theory, Ser. B*, 80(2):346–355, 2000.
- [47] A. Schrijver. *Combinatorial optimization. Polyhedra and efficiency. Vol. B*, volume 24 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2003. Matroids, trees, stable sets, Chapters 39–69.
- [48] A. Sebő. Eight-fifth approximation for the path TSP. In *Proc. IPCO 2013*, volume 7801 of *Lecture Notes Comput. Sci.*, pages 362–374. 2013.
- [49] A. Sebő and A. van Zuylen. The salesman’s improved paths: a $3/2 + 1/34$ approximation. In *Proc. FOCS 2016*, pages 118–127. 2016.
- [50] A. Serdyukov. On some extremal walks in graphs. *Upravlyaemye systemy*, 17:76–79, 1978.
- [51] M. Singh and L. C. Lau. Approximating minimum bounded degree spanning trees to within one of optimal. *J. ACM*, 62(1):Art. 1, 19, 2015.
- [52] O. Svensson, J. Tarnawski, and L. A. Végh. A constant-factor approximation algorithm for the asymmetric traveling salesman problem. In *Proc. STOC 2018*, pages 204–213. 2018.
- [53] V. Traub and J. Vygen. Approaching $\frac{3}{2}$ for the s - t -path TSP. *J. ACM*, 66(2):Art. 14, 17, 2019.
- [54] V. Traub and J. Vygen. An improved approximation algorithm for ATSP. In *Proc. STOC 2020*, 2020.
-

-
- [55] V. Traub, J. Vygen, and R. Zenklusen. Reducing path TSP to TSP. In *Proc. STOC 2020*, 2020.
- [56] J. A. A. van der Veen, G. J. Woeginger, and S. Zhang. Sequencing jobs that require common resources on a single machine: a solvable case of the TSP. *Math. Prog.*, 82:235–254, 1998.
- [57] J. Vygen. Reassembling trees for the traveling salesman. *SIAM J. Discrete Math.*, 30(2):875–894, 2016.
- [58] R. Zenklusen. A 1.5-Approximation for Path TSP. In *Proc. SODA 2019*, pages 1539–1549, 2019.