



Time- and Space-optimal Algorithm for the Many-visits TSP

ANDRÉ BERGER, Maastricht University, The Netherlands

LÁSZLÓ KOZMA, Freie Universität Berlin, Germany

MATTHIAS MNICH and ROLAND VINCZE, Technische Universität Hamburg, Germany

The many-visits traveling salesperson problem (MV-TSP) asks for an optimal tour of n cities that visits each city c a prescribed number k_c of times. Travel costs may be asymmetric, and visiting a city twice in a row may incur a non-zero cost. The MV-TSP problem finds applications in scheduling, geometric approximation, and Hamiltonicity of certain graph families.

The fastest known algorithm for MV-TSP is due to Cosmadakis and Papadimitriou (SICOMP, 1984). It runs in time $n^{O(n)} + O(n^3 \log \sum_c k_c)$ and requires $n^{\Theta(n)}$ space. An interesting feature of the Cosmadakis-Papadimitriou algorithm is its *logarithmic* dependence on the total length $\sum_c k_c$ of the tour, allowing the algorithm to handle instances with very long tours. The *superexponential* dependence on the number of cities in both the time and space complexity, however, renders the algorithm impractical for all but the narrowest range of this parameter.

In this article, we improve upon the Cosmadakis-Papadimitriou algorithm, giving an MV-TSP algorithm that runs in time $2^{O(n)}$, i.e., *single-exponential* in the number of cities, using *polynomial* space. The space requirement of our algorithm is (essentially) the size of the output, and assuming the Exponential-Time Hypothesis (ETH), the problem cannot be solved in time $2^{o(n)}$. Our algorithm is deterministic, and arguably both simpler and easier to analyze than the original approach of Cosmadakis and Papadimitriou. It involves an optimization over directed spanning trees and a recursive, centroid-based decomposition of trees.

CCS Concepts: • **Theory of computation** → **Graph algorithms analysis; Discrete optimization; Parameterized complexity and exact algorithms;**

Additional Key Words and Phrases: TSP, spanning trees, high-multiplicity scheduling

ACM Reference format:

André Berger, László Kozma, Matthias Mnich, and Roland Vincze. 2020. Time- and Space-optimal Algorithm for the Many-visits TSP. *ACM Trans. Algorithms* 16, 3, Article 35 (May 2020), 22 pages.

<https://doi.org/10.1145/3382038>

A preliminary version of this article appeared in the ACM-SIAM Symposium on Discrete Algorithms (SODA) 2019.

Research of L. K. supported by ERC Consolidator Grant No 617951 and DFG Grant KO 6140/1-1. Research of M. M. supported by DFG Grant MN 59/4-1.

Authors' addresses: A. Berger, Maastricht University, Department of Quantitative Economics, Maastricht, The Netherlands; email: a.berger@maastrichtuniversity.nl; L. Kozma, Freie Universität Berlin, Institute of Computer Science, Berlin, Germany; email: laszlo.kozma@fu-berlin.de; M. Mnich and R. Vincze, Technische Universität Hamburg, Institute for Algorithms and Complexity, Hamburg, Germany; emails: {mmnich, roland.vincze}@tuhh.de.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1549-6325/2020/05-ART35 \$15.00

<https://doi.org/10.1145/3382038>

1 INTRODUCTION

The traveling salesperson problem (TSP) is one of the cornerstones of combinatorial optimization, with origins going back (at least) to the 19th century work of Hamilton. For surveys on the rich history, variants, and current status of TSP, we refer to the dedicated books [3, 12, 22, 39]. In the standard TSP, given n cities and their pairwise distances, we seek a tour of minimum total distance that visits each city. If the distances obey the triangle inequality, then an optimal tour necessarily visits each city *exactly once* (apart from returning to the starting city in the end). In the general case of the TSP with arbitrary distances, the optimal tour may visit a city multiple times. Instances with non-metric distances arise from various applications that are modeled by the TSP, e.g., from scheduling problems.

To date, the fastest known exact algorithms for TSP (both in the metric and non-metric cases) are due to Bellman [6] and Held and Karp [25], running in time $O(2^n n^2)$ for n -city instances; both algorithms also require space $\Omega(2^n)$.

In this article, we study the general problem where each city has to be visited *exactly* a given number of times. More precisely, we are given a set V of n vertices, with pairwise distances (or costs) $d_{ij} \in \mathbb{N} \cup \{\infty\}$, for all $i, j \in V$. No further assumptions are made on the values d_{ij} ; in particular, they may be asymmetric, i.e., d_{ij} may not equal d_{ji} , and the cost d_{ii} of a self-loop may be non-zero. Also given are integers $k_i \geq 1$ for $i \in V$, which we refer to as *multiplicities*. A valid tour of length k is a sequence $(x_1, \dots, x_k) \in V^k$, where $k = \sum_{i \in V} k_i$, such that each $i \in V$ appears in the sequence exactly k_i times. The cost of the tour is $\sum_{i=1}^{k-1} d_{x_i, x_{i+1}} + d_{x_k, x_1}$. Our goal is to find a valid tour with minimum cost.

The problem is known as the *many-visits TSP* (MV-TSP). As an alternative name, *high-multiplicity TSP* also appears in the literature. When $k_i = 1$ for all $i \in V$, MV-TSP includes the standard metric TSP as a special case and thus, unless $P = NP$, it cannot be solved in polynomial time. Nonetheless, as discussed later, the problem can be solved efficiently when the number of cities (vertices) n is small, even if the length k of the tour is very large (possibly exponential in n).

As a natural TSP-generalization, MV-TSP is a fundamental problem of independent interest. In addition, MV-TSP has proved to be useful for modeling other problems, particularly in scheduling [9, 27, 47, 50]. Suppose there are k jobs of n different types to be executed on a single, universal machine. Processing a job, as well as switching to another type of job, come with certain costs, and the goal is to find the sequence of jobs with minimal total cost. Modeling this problem as a MV-TSP instance is straightforward, by letting d_{ij} denote the cost of processing a job of type i together with the cost of switching from type i to type j . (Note that d_{ii} is not necessarily smaller than d_{ij} for $i \neq j$; in some applications it is beneficial to switch between job types.) Emmons and Mathur [15] also describe an application of MV-TSP to the no-wait flow shop problem.

A different kind of application comes from geometric approximation. A standard technique to approximate geometric optimization problems is to reduce the size of the input by grouping certain input points together. Each group is then replaced by a single representative, and the reduced instance is solved exactly. For example, we may snap input points to nearby grid points, if doing so does not significantly affect the objective cost. Recently, this technique was used by Kozma and Mömke, to give an efficient polynomial-time approximation scheme (EPTAS) for the MAXIMUM SCATTER TSP in doubling metrics [37], addressing an open question of Arkin et al. [4]. In this case, the reduced problem is exactly the MV-TSP. Yet another application of MV-TSP is in settling the parameterized complexity of finding a Hamiltonian cycle in a graph class with restricted neighborhood structure [38].

To the best of our knowledge, MV-TSP was first considered in 1966 by Rothkopf [48]. In 1980, Psaraftis [47] gave a dynamic programming algorithm with run time $O(n^2 \cdot \prod_{i \in V} (k_i + 1))$. Observe that this quantity may be as high as $(k/n + 1)^n$, which is prohibitive even for moderately large

values of k . In 1984, Cosmadakis and Papadimitriou [13] observed that MV-TSP can be decomposed into a connectivity subproblem and an assignment subproblem. Taking advantage of this decomposition, they designed a family of algorithms, the best of which has run time $O^*(n^{2n}2^n + \log k)$.¹ This result can be seen as an early example of *fixed-parameter tractability*, where the rapid growth in complexity is restricted to a certain parameter.

The algorithm of Cosmadakis and Papadimitriou is, to date, the fastest solution to MV-TSP.² Its analysis is highly non-trivial, combining graph-theoretic insights and involved estimates of various combinatorial quantities. The analysis of the Cosmadakis-Papadimitriou algorithm is not known to be tight, but a lower bound of the form $2^{\Omega(n \log n)}$ is known to hold for its run time [13, p. 104]. Similarly, in the space requirement of the algorithm, a term of the form $2^{\Omega(n \log n)}$ appears hard to avoid. While it extends the tractability of TSP to a new range of parameters, the usefulness of the Cosmadakis-Papadimitriou algorithm is limited by its superexponential³ dependence on n in the run time, with the issue of superexponential *space* perhaps even more worrisome.

There have been further studies of the MV-TSP problem. Van der Veen and Zhang [50] discuss a problem equivalent to MV-TSP, called *K-group TSP*, and describe an algorithm with polylogarithmic dependence on the number k of visits, similarly to Cosmadakis and Papadimitriou. The value n , however, is assumed constant, and its effect on the run time is not explicitly computed (the dependence can be seen to be superexponential). Finally, Grigoriev and Van de Klundert [20] give an ILP formulation for MV-TSP with $O(n^2)$ variables. Applying Kannan's improvement [31] of Lenstra's algorithm [40] for solving fixed-dimensional ILPs to this formulation yields an algorithm with run time $n^{O(n^2)} \log k$. Further ILP formulations for MV-TSP are due to Sarin et al. [49] and Aguayo et al. [2], both of which again require superexponential time to be solved by standard algorithms. For further details about the history of the MV-TSP problem, we refer to the TSP textbook of Gutin and Punnen [22, § 11.10].

Our results. Our main result improves both the time and space complexity of the best known algorithm for MV-TSP, the first improvement in over 35 years. Specifically, we show that a *logarithmic* dependence on the number k of visits, a *single-exponential* dependence on the number n of cities, and a *polynomial* space complexity are simultaneously achievable. Moreover, while we build upon ideas from the previous best approach, our algorithm is arguably easier to describe, easier to implement, and easier to analyze than its predecessor. To introduce the techniques step-by-step, we describe *three* algorithms for solving MV-TSP. We refer to the three algorithms, after their core subroutines, as ENUM-MV, DP-MV, and DC-MV. The acronyms stand for enumeration, dynamic programming, and divide and conquer. We also describe an improved variant of the third algorithm, called DC-MV2. All our algorithms are deterministic. Their complexities are summarized in Theorem 1.1, proved in Section 2.

THEOREM 1.1.

- (i) ENUM-MV solves MV-TSP using space $O(n^2)$, in time $O^*(n^n + \log k)$.
- (ii) DP-MV solves MV-TSP using space $O(5^n)$, in time $O^*(5^n + \log k)$.
- (iii) DC-MV2 solves MV-TSP using space $O(n^2)$, in time $O^*(16^{n+o(n)} + \log k)$.

The Exponential-Time Hypothesis (ETH) [28, 29] implies that TSP cannot be solved in $2^{o(n)}$, i.e., sub-exponential, time. Under this hypothesis, the run time of our algorithm DC-MV2 is

¹Here, and in the following, the $O^*(\cdot)$ notation is used to suppress a factor bounded by a polynomial in n .

²It may seem that a linear dependence on the length k of the tour is necessary even to output the result. Observe, however, that a tour can be compactly represented by collapsing cycles and storing them together with their multiplicities.

³In this article, the term *superexponential* always refers to a quantity of the form $2^{\omega(n)}$.

asymptotically optimal for MV-TSP, up to the base of the exponential. Further, note that the space requirement of DC-MV2 is also (essentially) optimal, as a compact solution encodes for each of the $\Omega(n^2)$ edges the number t of times that this edge is traversed by an optimal tour. Note that throughout the article, we assume that each multiplicity can be stored in a constant number of machine words; if this is not the case, e.g., if k is exponential in n , a factor $O(\log k)$ should be applied to the given space bounds.

Our result leads to improvements in applications where MV-TSP is solved as a subroutine. For instance, as a corollary of Theorem 1.1(iii), the approximation scheme for Maximum Scatter TSP [37] can now be implemented in space *polynomial* in the error parameter ε .

It is interesting to contrast our results for MV-TSP with recent results for the *r-simple path* problem, where a path of length k is sought that visits each vertex *at most* r times. For that problem, the fastest known algorithms—due to Abasi et al. [1] and Gabizon et al. [18]—have run time *exponential* in $k/r \cdot \log r$, and such exponential dependence is necessary assuming ETH.

Overview of techniques. The Cosmadakis-Papadimitriou algorithm is based on the following high-level insight, common to most work on the TSP problem, whether exact or approximate. The task of finding a valid tour may be split into two separate tasks: (1) finding a structure that connects all vertices, and (2) augmenting the structure found in (1) to ensure that each vertex is visited the required number of times. Indeed, such an approach is also used in the well-known 3/2-approximation algorithm of Christofides for metric TSP [11]. There, the structure that guarantees connectivity is a minimum spanning tree, and “visitability” is enforced by the addition of a perfect matching that connects odd-degree vertices, ensuring that all vertices have even degree, and can thus be entered and exited, as required.

In the case of MV-TSP, Cosmadakis and Papadimitriou ensure connectivity (part (1)) by finding a *minimal connected Eulerian digraph* on the set V of input vertices. Indeed, such a digraph must be part of every solution, since a tour must balance every vertex (equal out-degree and in-degree), and all vertices must be mutually reachable. Minimality is meant here in the sense that no proper subgraph is Eulerian and is required only to reduce the search space. Assuming that a connected Eulerian subdigraph of the solution is found, it needs to be extended to the edge set of a valid tour (part (2)). If this is done with the cheapest possible set of edges, then the optimum must have been found. This second step amounts to solving a transportation problem, which takes polynomial time.

The first step, however, requires us to implicitly consider all possible minimal Eulerian digraphs. As it is NP-complete to test the non-minimality of a Eulerian digraph [46], the authors relax minimality and suggest the use of heuristics for pruning out non-minimal instances in practice. However, they obtain a saving in run time by observing that among all digraphs with the same *degree sequence* only one with smallest cost needs to be considered. Otherwise, in the final tour, the connected Eulerian subdigraph could be swapped with a cheaper one, while maintaining the validity of the tour.

Cosmadakis and Papadimitriou thus iterate over feasible degree sequences of connected Eulerian digraphs; for each such degree sequence they construct the cheapest realization (which may not be minimal) by dynamic programming; finally they construct, for each obtained Eulerian digraph, the cheapest extension to a valid tour by solving a transportation problem. The cheapest tour found over all iterations is returned as the solution. Iterating and optimizing over these structures is no easy task, and Cosmadakis and Papadimitriou invoke a number of graph-theoretic and combinatorial insights. For estimating the total cost of their procedure a sophisticated global counting argument is developed.

The key insight of our approach is that the machinery involving Eulerian digraphs is *not necessary* for solving MV-TSP. To ensure connectivity (i.e., task (1) above), a *directed spanning tree* is sufficient. This may seem surprising, as a directed tree fails to satisfy the main property of connected Eulerian digraphs: *strong connectivity*. Observe, however, that a collection of directed edges with the same out-degrees and in-degrees as a valid MV-TSP tour is itself a valid tour, as long as the digraph determined by the collection of edges is weakly connected. Requiring the solution to contain a tree is sufficient to avoid the case of disjoint cycles. The fact that the tree can be assumed to be *rooted*, i.e., all of its edges are directed away from some vertex, follows from the strong connectedness of the tour.⁴

Directed spanning trees are easier to enumerate and optimize over than Eulerian digraphs; this fact alone explains the reduced complexity of our approach. However, to obtain our main result, further ideas are needed. In particular, we find the cheapest directed spanning tree that is feasible for a given degree sequence—first by dynamic programming, then by a recursive partitioning of trees, based on centroid-decompositions.

2 IMPROVED ALGORITHMS FOR THE MANY-VISITS TSP

In this section, we describe and analyze our three algorithms. The first, ENUM-MV, is based on exact enumeration of trees (Section 2.2); the second, DP-MV, uses a dynamic programming approach to find an optimal tree (Section 2.3); and the third, DC-MV, is based on divide and conquer (Section 2.4 and Section 2.5). Before presenting the algorithms, we introduce some notation and structural observations that are subsequently used (Section 2.1).

2.1 Trees, Tours, and Degree Sequences

Let V be a set of vertices. We view a *directed multigraph* G with vertex set V as a *multiset* of edges (i.e., elements of $V \times V$). Accordingly, self-loops and multiple copies of the same edge are allowed. The *multiplicity* of an edge (i, j) in a directed multigraph G is denoted $m_G(i, j)$. The *out-degree* of a vertex $i \in V$ is $\delta_G^{\text{out}}(i) = \sum_{j \in V} m_G(i, j)$, and the *in-degree* of a vertex $i \in V$ is $\delta_G^{\text{in}}(i) = \sum_{j \in V} m_G(j, i)$. Given edge costs $d : V \times V \rightarrow \mathbb{N} \cup \{\infty\}$, the *cost* of G is simply the sum of its edge costs, i.e., $\text{cost}(G) = \sum_{i, j \in V} m_G(i, j) \cdot d(i, j)$.

For two directed multigraphs G and H over the same vertex set V , let $G + H$ denote the directed multigraph obtained by adding the corresponding edge multiplicities of G and H . Observe that as an effect, out-degrees and in-degrees are also added pointwise. Formally, $m_{G+H}(i, j) = m_G(i, j) + m_H(i, j)$, and $\delta_{G+H}^{\text{out}}(i) = \delta_G^{\text{out}}(i) + \delta_H^{\text{out}}(i)$, and $\delta_{G+H}^{\text{in}}(i) = \delta_G^{\text{in}}(i) + \delta_H^{\text{in}}(i)$, for all $i, j \in V$. The relation $\text{cost}(G + H) = \text{cost}(G) + \text{cost}(H)$ clearly holds.

In the following, for a directed multigraph G , we refer to its *underlying graph*, i.e., to the undirected graph consisting of those edges $\{i, j\}$ for which $m_G(i, j) + m_G(j, i) \geq 1$.

Consider a tour $C = (x_1, \dots, x_k) \in V^k$. We refer to the unique directed multigraph G consisting of the edges $(x_1, x_2), \dots, (x_{k-1}, x_k), (x_k, x_1)$ as the *edge set* of C . We state a simple but crucial observation.

LEMMA 2.1. *Let G be a directed multigraph over V with out-degrees $\delta_G^{\text{out}}(\cdot)$ and in-degrees $\delta_G^{\text{in}}(\cdot)$. Then G is the edge set of a tour that visits each vertex $i \in V$ exactly k_i times if and only if both of the following conditions hold:*

- (i) *the underlying graph of G is connected, and*
- (ii) *for all $i \in V$, we have $\delta_G^{\text{out}}(i) = \delta_G^{\text{in}}(i) = k_i$.*

⁴We thank Andreas Björklund for the latter observation that led to an improved run time and a simpler correctness argument.

PROOF. The fact that connectedness of G and $\delta_G^{\text{out}}(i) = \delta_G^{\text{in}}(i)$ is equivalent with the existence of a tour that uses each edge of G exactly once is the well-known ‘‘Euler’s theorem.’’ (See Reference [5, Theorem 1.6.3] for a short proof.) Clearly, visiting each vertex i exactly k_i times is equivalent with the condition that the tour contains k_i edges of the form (\cdot, i) and k_i edges of the form (i, \cdot) . \square

Moreover, given the edge set G of a tour C , a tour C' with edge set G can easily be recovered. This amounts to finding a Eulerian tour of G , which can be done in time linear in the length k of the tour. To avoid a linear dependence on k , we can apply the algorithm of Grigoriev and Van de Klundert [20] that constructs a compact representation of C' in time $O(n^4 \log k)$. As the edge sets of C and C' are equal, C' also visits each $i \in V$ exactly k_i times and $\text{cost}(C') = \text{cost}(C) = \text{cost}(G)$. Thus, in solving MV-TSP, we only focus on finding a minimum cost directed multigraph whose underlying undirected graph is connected and whose degrees match the multiplicities required by the problem.

A *directed spanning tree* of V is a tree with vertex set V whose edges are directed *away* from some vertex $r \in V$; in other words, the tree contains a directed path from r to every other vertex in V . (Directed spanning trees are alternatively called *branchings*, *arborescences*, or *out-trees*.) We refer to the vertex r as the *root* of the tree. We observe that every valid tour contains a directed spanning tree.

LEMMA 2.2. *Let G be the edge set of a tour of V with arbitrary non-zero multiplicities, and let $r \in V$ be an arbitrary vertex. Then there is a directed spanning tree T of G rooted at r and a directed multigraph X , such that $G = T + X$.*

PROOF. We choose T to be the single-source *shortest path tree* in G with source r . More precisely, let $d(v)$ denote the *distance* (i.e., number of edges) from r to v in G . Observe that in a valid tour all vertices are mutually reachable, so $d(v)$ is finite for all $v \in V$. Let us now build T by adding, for each vertex $v \in V \setminus \{r\}$, an edge (w, v) , where $w \in V$ is an arbitrary vertex such that $d(w) = d(v) - 1$. Such a w must exist, as the predecessor of v on a shortest path from r has this property. The fact that $d(\cdot)$ strictly increases along edges ensures that T is cycle-free, i.e., a tree. \square

We can thus split the MV-TSP problem into finding a directed spanning tree T with an arbitrary root r and an extension X , such that $T + X$ is a valid tour. We claim that in the decomposition $G = T + X$ of an optimal tour G , both T and X are optimal with respect to their degree sequences.

LEMMA 2.3. *Let G be the edge set of an optimal tour for MV-TSP, let T be a directed spanning tree, and let X be a directed multigraph such that $G = T + X$. Then, T has the smallest cost among all directed spanning trees with degrees $\delta_T^{\text{out}}(\cdot)$ and $\delta_T^{\text{in}}(\cdot)$, and X has the smallest cost among all directed multigraphs with degrees $\delta_X^{\text{out}}(\cdot)$ and $\delta_X^{\text{in}}(\cdot)$.*

PROOF. Suppose there is a directed spanning tree T' such that $\text{cost}(T') < \text{cost}(T)$, and $\delta_{T'}^{\text{out}}(i) = \delta_T^{\text{out}}(i)$, and $\delta_{T'}^{\text{in}}(i) = \delta_T^{\text{in}}(i)$ for all $i \in V$. But then $T' + X$ is connected, has the same degree sequence as G , while $\text{cost}(T' + X) < \text{cost}(G)$, contradicting the optimality of G .

Similarly, suppose there is a directed multigraph X' such that $\text{cost}(X') < \text{cost}(X)$, and $\delta_{X'}^{\text{out}}(i) = \delta_X^{\text{out}}(i)$, and $\delta_{X'}^{\text{in}}(i) = \delta_X^{\text{in}}(i)$ for all $i \in V$. But then $T + X'$ is connected, has the same degree sequence as G , while $\text{cost}(T + X') < \text{cost}(G)$, contradicting the optimality of G . \square

Next, we characterize the feasible degree sequences of directed spanning trees.⁵

LEMMA 2.4. *Let $V = \{x_1, \dots, x_n\}$ be a set of vertices, where $n \geq 2$. There is a directed spanning tree of V with root x_1 whose out-degrees and in-degrees are, respectively, $\delta^{\text{out}}(\cdot)$ and $\delta^{\text{in}}(\cdot)$, if and only if*

⁵The feasibility of degree sequences for various graph classes is a well-studied subject, see, e.g., References [8, 16, 17, 23, 24, 33, 35]. The simple condition we state is similar to the condition for *undirected* trees, given by Berge [7, page 117].

- (i) $\delta^{\text{in}}(x_1) = 0$,
- (ii) $\delta^{\text{in}}(x_i) = 1$ for $i = 2, \dots, n$,
- (iii) $\delta^{\text{out}}(x_1) > 0$, and
- (iv) $\sum_i \delta^{\text{out}}(x_i) = n - 1$.

PROOF. In the forward direction, in a directed spanning tree all non-root vertices have exactly one parent, proving (i) and (ii). The root must have at least one child (iii), and the total number of edges is $n - 1$, proving (iv).

In the backward direction, we argue by induction on n . In the case $n = 2$, we have $\delta^{\text{in}}(x_1) = \delta^{\text{out}}(x_2) = 0$, and $\delta^{\text{in}}(x_2) = \delta^{\text{out}}(x_1) = 1$, hence an edge (x_1, x_2) satisfies the degree requirements.

Consider now the case of $n > 2$ vertices. From (ii)–(iv) it follows that for some $k \in \{2, \dots, n\}$, we have $\delta^{\text{in}}(x_k) = 1$ and $\delta^{\text{out}}(x_k) = 0$, i.e., x_k is a leaf.

Let x_j be a vertex for some $j \in \{1, \dots, n\} \setminus \{k\}$ such that $\delta^{\text{out}}(x_j) \geq 1$, and $\delta^{\text{in}}(x_j) + \delta^{\text{out}}(x_j) \geq 2$; by (ii)–(iv) there must be such a vertex. We decrease $\delta^{\text{out}}(x_j)$ by one. Conditions (i)–(iv) clearly hold for $V \setminus \{x_k\}$. By induction, we can build a tree on $V \setminus \{x_k\}$ and attach x_k to this tree as a leaf with the edge (x_j, x_k) . \square

Let $\text{DS}(n)$ denote the number of different pairs of sequences $(\delta'_1, \dots, \delta'_n), (\delta''_1, \dots, \delta''_n)$ that are feasible for a directed spanning tree, i.e., for vertex set $V = \{x_1, \dots, x_n\}$, for some directed spanning tree T with root x_1 , we have $\delta_T^{\text{out}}(x_i) = \delta'_i$, and $\delta_T^{\text{in}}(x_i) = \delta''_i$, for all $i \in \{1, \dots, n\}$. By Lemma 2.4, $\text{DS}(n)$ equals the number of ways to distribute $n - 1$ out-degrees to n vertices, such that a designated vertex has non-zero out-degree. This task is the same as distributing $n - 2$ balls arbitrarily into n bins, of which there are $\binom{2n-3}{n-1} = O(4^n)$ ways.

2.2 ENUM-MV: Polynomial Space and Superexponential Time

Given the vertex set $V = \{1, \dots, n\}$, multiplicities $k_i \in \mathbb{N}$, and cost function d , we wish to find a tour of minimum cost that visits each $i \in V$ exactly k_i times. From Lemma 2.2, our first algorithm presents itself. It simply iterates over all directed spanning trees T with vertex set V and extends each of them optimally to a valid tour C_T . Among all valid tours constructed, one with smallest cost is returned (Algorithm 1).

This simple algorithm already improves on the previous best run time (although it is still superexponential) and reduces the space requirement from superexponential to polynomial.

ALGORITHM 1: ENUM-MV for solving MV-TSP using enumeration

Input: Vertex set V , cost function d , multiplicities k_i .

Output: A tour of minimum cost that visits each $i \in V$ exactly k_i times.

- 1: **for** each directed spanning tree T with root $r \in V$ **do**
 - 2: Find minimum cost directed multigraph X such that for all $i \in V$:

$$\delta_X^{\text{out}}(i) = k_i - \delta_T^{\text{out}}(i),$$

$$\delta_X^{\text{in}}(i) = k_i - \delta_T^{\text{in}}(i).$$
 Let $C_T \leftarrow T + X$.
 - 3: **return** C_T with smallest cost.
-

The correctness of the algorithm is immediate: From Lemma 2.2, it follows that all C_T 's considered are valid (connected, and with degrees matching the required multiplicities), and by Lemma 2.3, the optimal tour C must be considered during the execution.

The iteration of Line 1 requires us to enumerate all labeled trees with vertex set V . There are n^{n-2} such trees [10], and standard techniques can be used to enumerate them with a constant overhead

per item (see, e.g., Kapoor and Ramesh [32]). For each considered tree, we orient the edges in a unique way, away from r .

Let T be the current tree. In Line 2, we find a minimum cost directed multigraph X , with given out-degree and in-degree sequence, such as to extend T into a valid tour. If, for some vertex i , it holds that $\delta_T^{\text{in}}(i) > k_i$ or $\delta_T^{\text{out}}(i) > k_i$, we proceed to the next spanning tree, since the current tree cannot be extended to a valid tour. Observe that this may happen only if $k_i < n - 1$, for some $i \in V$. Otherwise, we find the optimal X by solving a transportation problem in polynomial time. Throughout the execution, we remember the best tour found so far, which we output in the end. We describe next the transportation subroutine, which is common to all our algorithms, and is essentially the same as in the Cosmadakis-Papadimitriou algorithm.

The transportation problem. The subproblem we need to solve is finding a minimum cost directed multigraph X over vertex set V , with given out-degree and in-degree requirements. We can map this problem to an instance of the Hitchcock transportation problem [26], where the goal is to transport a given amount of goods from N warehouses to M outlets with given pairwise shipping costs. Note that this is a special case of the more general min-cost max-flow problem [14].

More precisely, let us define a digraph with vertices $\{s, t\} \cup \{s_i, t_i \mid i \in V\}$. Edges are $\{(s, s_i), (t_i, t) \mid i \in V\}$, and $\{(s_i, t_j) \mid i, j \in V\}$. We set cost 0 to edges (s, s_i) and (t_i, t) and cost d_{ij} (i.e., the costs given in the MV-TSP instance) to (s_i, t_j) . We set capacity ∞ to edges (s_i, t_j) and capacities $k_i - \delta_T^{\text{out}}(i)$ to (s, s_i) , and $k_i - \delta_T^{\text{in}}(i)$ to (t_i, t) . The construction is identical to the one used by Cosmadakis and Papadimitriou, apart from the fact that in our case the capacity of (s, s_i) may be different from the capacity of (t_i, t) . Observe that the sum of capacities of (s, s_i) -edges equals the sum of capacities of (t_i, t) -edges over all $i \in V$. Thus, a maximal $s - t$ flow saturates all these edges. The amount of flow transmitted on the edge (s_i, t_j) gives the multiplicity of edge (i, j) in the sought-after multigraph, for all $i, j \in V$. A minimum cost maximum flow clearly maps to a minimum cost edge set with the given degree constraints.

In the Cosmadakis-Papadimitriou algorithm, the transportation subproblems are solved via the scaling method of Edmonds and Karp [14]. This algorithm proceeds by solving $O(\log k)$ approximate versions of the problem, where the costs are the same as in the original problem, but the capacities are scaled (i.e., divided and rounded down) by a factor 2^p for $p = \lceil \log k \rceil, \dots, 0$. Each approximate problem is solved in $O(n^3)$ time by performing flow augmentations on the optimal flow found in the previous approximation, multiplied by two. The overall run time for solving the described transportation problem is therefore $O(n^3 \log k)$.

Cosmadakis and Papadimitriou describe an improvement that also applies for our case. Namely, they show that the total run time for solving several instances with the same costs can be reduced if the capacities on corresponding edges in two different instances may differ by at most n . The strategy is to solve all but the last $\log n$ approximate problems only once, as these are (essentially) the same for all instances. For different instances, we only need to solve the last $\log n$ approximate problems, i.e., at the finest levels of approximation. This gives a run time of $O(n^3 \log k)$ for solving the “master problem,” and $O(n^3 \log n)$ for solving each individual instance. We refer to Cosmadakis and Papadimitriou [13] as well as to Edmonds and Karp [14] for details.

In our case, the different instances of the transportation problem are for finding the directed multigraphs X for different trees T . Each of these instances agree in the underlying graph and cost function and may differ only in the capacities. As the maximum degree of each tree T is at most $n - 1$, the differences are bounded, as required.

As an alternative to the Edmonds-Karp algorithm, we may solve the arising transportation problems with a *strongly polynomial* algorithm, e.g., the one by Orlin [45] or its extension due to Kleinschmidt and Schannath [34]. Note that these algorithms were not yet available

when Cosmadakis and Papadimitriou obtained their result. The run time for the transportation subproblem then becomes $O(n^3 \log n)$, i.e., independent of k . Such an improvement is likely of theoretical interest only; furthermore, it assumes that operations on the multiplicities k_i take constant time. If this assumption is unrealistic, e.g., if k is exponential in n , then we may fall back to the Edmonds-Karp algorithm, with the term $O(n^3 \log k)$ added to the run time.

Analysis of Algorithm 1. We iterate over all $O(n^{n-2})$ directed spanning trees and solve a transportation problem for each, with run time $O(n^3 \log n)$. The total run time $O(n^{n+1} \log n)$ follows. The space requirement of the algorithm is dominated by that of solving a (single) transportation problem and of storing the edge set of a single tour (apart from minor bookkeeping).

2.2.1 Improved Enumeration Algorithm. We can slightly improve the run time of ENUM-MV by observing that the solution of the transportation problem depends only on the degree sequence of the current tree T and not the actual edges of T . Therefore, different trees with the same degree sequence can be extended in the same way. Observe that several trees may have the same degree sequence; in an extreme case, all $(n-2)!$ simple Hamiltonian paths with the same endpoints have the same degree sequence.

ALGORITHM 2: ENUM-MV for solving MV-TSP using enumeration (improved)

Input: Vertex set V , cost function d , multiplicities k_i .

Output: A tour of minimum cost that visits each $i \in V$ exactly k_i times.

- 1: **for** each feasible degree sequence $\delta^{\text{out}}(\cdot), \delta^{\text{in}}(\cdot)$ of a directed spanning tree with root r **do**
 - 2: Find minimum cost directed multigraph X such that for all $i \in V$:

$$\delta_X^{\text{out}}(i) = k_i - \delta^{\text{out}}(i),$$

$$\delta_X^{\text{in}}(i) = k_i - \delta^{\text{in}}(i).$$
 - 3: **for** each directed spanning tree T with $\delta_T^{\text{out}} = \delta^{\text{out}}$ and $\delta_T^{\text{in}} = \delta^{\text{in}}$ **do**
 Let $C_T \leftarrow T + X$.
 - 4: **return** C_T with smallest cost.
-

Algorithm 2 implements this idea, iterating over all trees, grouped by their degree sequences. It solves the transportation problem only once for each degree sequence (there are $DS(n)$ of them).

The correctness is immediate, as all directed spanning trees T are still considered as before. Assume that we can iterate over all feasible degree sequences and all corresponding trees with $O(n)$ overhead per item. By Lemma 2.4, the first task only requires us to consider all ways of distributing $n-2$ out-degrees among n vertices. For completeness, we describe a procedure for this task in Appendix A. We give the subroutine for the second task in Section 2.2.2. We thus obtain the run time $O(n^{n-1} + DS(n) \cdot n^4 \log n) = O(n^{n-1})$. The space requirement is asymptotically unchanged.

2.2.2 Generating Trees by Degree Sequence. In the proof of Lemma 2.4, we generate *one* directed tree from its degree sequence. In this section, we show how to generate *all* trees for a given degree sequence (Algorithm 3).

The initial call to the BUILD TREE() procedure is with a feasible input degree sequence $(\delta^{\text{out}}, \delta^{\text{in}})$ and an empty “stub” ($\mathbf{x}^{\text{stub}} = \{\}$) as arguments. The algorithm finds the first *unattached* vertex that is either a leaf of the final tree or whose subtree is already complete (that is, $\delta^{\text{out}}(x_i) = 0$), then it finds all possibilities for attaching x_i to the rest of the tree. Observe that the fact that x_i has no more capacity for outgoing edges prevents cycles. The procedure is then called recursively with modified arguments: Edge (x_j, x_i) is added to the stub and the out-degree of x_j and the in-degree of x_i are decremented.

ALGORITHM 3: BUILDTREE for generating all directed trees for a given degree sequence

Input: Degree sequence $(\delta^{\text{out}}, \delta^{\text{in}})$ and forest \mathbf{x}^{stub} .

Output: A generator of all trees with degree sequence $(\delta^{\text{out}}, \delta^{\text{in}})$, attached to \mathbf{x}^{stub} .

```

1: procedure BUILDTREE( $\delta^{\text{out}}, \delta^{\text{in}}, \mathbf{x}^{\text{stub}}$ )
2:   if  $\sum \delta^{\text{in}}(\cdot) = 1$  then
3:     Let  $i$  be the index where  $\delta^{\text{in}}(x_i) = 1$ .
4:     yield  $\mathbf{x}^{\text{stub}} \cup \{(x_1, x_i)\}$ 
5:   else
6:     Let  $i$  be the first index where  $\delta^{\text{out}}(x_i) = 0$  and  $\delta^{\text{in}}(x_i) = 1$ .
7:     for every index  $j \neq i$ , such that  $\delta^{\text{out}}(x_j) + \delta^{\text{in}}(x_j) \geq 2$  do      ▷ attach  $x_i$  to parent  $x_j$ 
8:        $(\delta^{\text{out}'}, \delta^{\text{in}'}) \leftarrow (\delta^{\text{out}}, \delta^{\text{in}})$ 
9:        $\delta^{\text{out}'}(x_j) \leftarrow \delta^{\text{out}'}(x_j) - 1$ 
10:       $\delta^{\text{in}'}(x_i) \leftarrow 0$ 
11:      BUILDTREE( $\delta^{\text{out}'}, \delta^{\text{in}'}, \mathbf{x}^{\text{stub}} \cup \{(x_j, x_i)\}$ )

```

At each recursive level there are as many new calls as possible candidates for the next edge, with both degree demands decreased by one, and with the stub gaining one additional edge. During the intermediate calls the stub may be disconnected, i.e., it represents a forest. At the $(n - 1)$ -th level exactly two degree-one vertices remain—say, x_i with in-degree 1 and the root x_1 with out-degree 1. Adding the edge (x_1, x_i) finishes the construction.

Observe that there are no “dead ends” during this process, i.e., every call of BUILDTREE eventually results in a valid directed tree in the last level of the recursion, and there are no discarded graphs during the process. The argument mirrors the one in the proof of Lemma 2.4. In every call to BUILDTREE, the degree sequence $(\delta^{\text{out}}, \delta^{\text{in}})$ encodes a valid directed tree on the subset of vertices with non-zero in-degree, together with the root x_1 . Let h denote the number of vertices with non-zero in-degree. We argue by induction on h that the construction succeeds. For $h = 1$ (Lines 2–4) this is clear. For $h \geq 2$ (Lines 6–11), we attach x_i to a vertex x_j of a valid tree, constructed by the recursive call with $h - 1$ vertices with non-zero in-degree.

For every directed tree with the required degree-sequence, repeatedly peeling off the lowest-index leaf gives a valid edge-insertion order that is found by the algorithm. (The order in which vertices are considered is essentially the *Prüfer sequence* [42] of the tree.) Finally, as the recursive calls correspond to attaching x_i to *different* parents, no tree is constructed more than once.

Note that Algorithm 3 does not return the list of trees realizing a degree sequence $(\delta^{\text{out}}, \delta^{\text{in}})$ at once; it is instead a *generator* that serves as the head of the loop in Step 3 of Algorithm 2, outputting a different tree in every iteration. We used the keyword **yield** instead of **return** to indicate that the algorithm does not terminate after “returning” the first result.

2.3 DP-MV: Exponential Space and Single-exponential Time

Next, we improve the run time to single-exponential, by making use of Lemma 2.3. Specifically, we observe that for every feasible degree sequence only the tree with minimum cost needs to be considered. The enumeration in Algorithm 3 can easily be modified to return, instead of all trees, just one with smallest cost; this, however, would not improve the asymptotic run time. Instead, in this section, we describe a dynamic programming algorithm resembling the algorithms by Bellman [6] and Held and Karp [25] for directly computing the best directed tree for a given degree sequence.

The outline of the algorithm is shown in Algorithm 4, and it is identical for DP-MV and DC-MV, described in Section 2.4 and Section 2.5. The algorithms DP-MV and DC-MV differ in the way they find the minimum cost directed tree, i.e., Line 2 of the generic Algorithm 4.

ALGORITHM 4: Solving MV-TSP by optimizing over trees (common for DP-MV, DC-MV, and DC-MV2).

Input: Vertex set V , cost function d , multiplicities k_i .

Output: A tour of minimum cost that visits each $i \in V$ exactly k_i times.

- 1: **for** each feasible degree sequence $\delta^{\text{out}}(\cdot), \delta^{\text{in}}(\cdot)$ of a directed spanning tree **do**
 - 2: Find a minimum-cost directed tree T with $\delta_T^{\text{in}} = \delta^{\text{in}}$ and $\delta_T^{\text{out}} = \delta^{\text{out}}$.
 - 3: Find a minimum-cost directed multigraph X such that for all $i \in V$:

$$\delta_X^{\text{in}}(i) = k_i - \delta_T^{\text{in}}(i),$$

$$\delta_X^{\text{out}}(i) = k_i - \delta_T^{\text{out}}(i).$$
 Let $C_T \leftarrow T + X$.
 - 4: **return** C_T with smallest cost.
-

The dynamic programming approach (DP-MV) resembles Algorithm 3 for iterating over all directed trees with a given degree sequence. Specifically, let $(\delta^{\text{out}}, \delta^{\text{in}})$ be the degree sequence for which we wish to find a minimum-cost tree. We build a dynamic programming table \mathcal{T} that holds an optimal tree and its cost for *every* feasible degree sequence. The solution can thus be read from $\mathcal{T}[\delta^{\text{out}}, \delta^{\text{in}}]$. Observe that specifying a degree sequence allows us to restrict the problem to arbitrary subsets of V by simply setting the degrees of non-participating vertices to zero.

To compute $\mathcal{T}[\delta^{\text{out}}, \delta^{\text{in}}]$, we find the leaf with smallest index x_i and all non-leaves x_j that may be connected to x_i by an edge in the optimal tree (similarly to Algorithm 3). For each choice of x_j , we compute the optimal tree by adding the connecting edge (x_j, x_i) to the optimal tree over $V \setminus \{x_i\}$, with the degree sequence suitably updated.

The correctness of the dynamic programming algorithm follows from an observation similar to Lemma 2.3; every subtree of the optimal tree must be optimal for its degree sequence, as otherwise it could be swapped for a cheaper subtree. The details are shown in Algorithm 5.

Analysis of Algorithm 5. Observe that the number of possible entries $\mathcal{T}[\cdot, \cdot]$, i.e., the number of feasible degree sequences for trees on subsets of V , is at most $\sum_{k=1}^n \binom{n}{k} \text{DS}(k)$. Using our previous estimate $\text{DS}(n) = O(4^n)$, this sum evaluates, by the binomial theorem, to $O(5^n)$. The overall run time of Algorithm 4, including the transportation subproblem, with Algorithm 5 as a subroutine, is $O^*(5^n)$, since the entry for a given degree sequence is computed at most once over all iterations, and the values are stored through the entire execution. In practice, storing an entire tree in each $\mathcal{T}[\cdot, \cdot]$ is wasteful; for the optimum tree to be constructible from the table, it is sufficient to store the node to which the lowest-index leaf is connected. The claimed time and space complexities follow.

2.4 DC-MV: Polynomial Space and Single-exponential Time

In this section, we show how to reduce the space complexity to polynomial, while maintaining a single-exponential run time.

The outer loop (Algorithm 4) remains the same, but we replace the subroutine for finding an optimal directed spanning tree (Algorithm 5) with an approach based on divide and conquer (Algorithm 6). The approach is inspired by the algorithm of Gurevich and Shelah [21] for finding an optimal TSP tour.

ALGORITHM 5: DP-MV for generating the optimal directed tree for a given degree sequence.

Input: Degree sequence $(\delta^{\text{out}}, \delta^{\text{in}})$.

Output: (T, c) , where T is optimum tree with root x_1 , degrees $\delta^{\text{out}}, \delta^{\text{in}}$, and cost c .

```

1: procedure  $\mathcal{T}$  [ $\delta^{\text{out}}, \delta^{\text{in}}$ ]
2:   if  $\sum \delta^{\text{in}}(\cdot) = 1$  then
3:     Let  $i$  be the index where  $\delta^{\text{in}}(x_i) = 1$ .
4:     return  $(\{(x_1, x_i)\}, d_{x_1, x_i})$ 
5:   else  $(\sum \delta^{\text{in}}(\cdot) > 1)$ 
6:     Let  $i$  be the first index where  $\delta^{\text{out}}(x_i) = 0$  and  $\delta^{\text{in}}(x_i) = 1$ .
7:     for every index  $j \neq i$ , such that  $\delta^{\text{out}}(x_j) + \delta^{\text{in}}(x_j) \geq 2$  do      ▷ attach  $x_i$  to parent  $x_j$ 
8:        $(\delta^{\text{out}'}, \delta^{\text{in}'}) \leftarrow (\delta^{\text{out}}, \delta^{\text{in}})$ 
9:        $\delta^{\text{out}'}(x_j) \leftarrow \delta^{\text{out}'}(x_j) - 1$ 
10:       $\delta^{\text{in}'}(x_i) \leftarrow 0$ 
11:       $(T_j, c_j) \leftarrow \mathcal{T}[\delta^{\text{out}'}, \delta^{\text{in}'}]$ 
12:       $T_j \leftarrow T_j \cup \{(x_j, x_i)\}$ 
13:       $c_j \leftarrow c_j + d_{x_j, x_i}$ 
14:     return  $(T_j, c_j)$  with minimum  $c_j$ 

```

The algorithm relies on the following observation about tree-separators: Let (V_1, V_2) be a partition of V , i.e., $V_1 \cup V_2 = V$, and $V_1 \cap V_2 = \emptyset$, and let $|V| = n$. We say that (V_1, V_2) is *balanced* if $\max\{|V_1|, |V_2|\} \leq \lceil 2n/3 \rceil$. The following result is folklore (see, e.g., Reference [44, § 9.1]); we give a short proof for completeness:

LEMMA 2.5. *Every tree T with vertex set V admits a balanced partition (V_1, V_2) of V such that all edges of T between V_1 and V_2 are incident to a vertex $v \in V_1$.*

PROOF. A very old result of Jordan [30] states that every tree has a *centroid*, which is a vertex whose removal splits the tree into subtrees not larger than half of the original tree. To find such a centroid, move from an arbitrary vertex, one edge at a time, towards the largest subtree.

Let T_1, \dots, T_m be the vertex sets of the trees, in decreasing order of size, resulting from deleting the centroid of T . Let e_1, \dots, e_m denote the edges that connect the respective trees to the centroid in T . If $|T_1| \geq \lceil n/3 \rceil$, then we have the balanced partition $(V \setminus T_1, T_1)$, with a single crossing edge.

Otherwise, let m' be the smallest index such that $\lceil n/3 \rceil \leq |T_1| + \dots + |T_{m'}| \leq \lceil 2n/3 \rceil$. Observe that as $|T_i| < \lceil n/3 \rceil$ for all i , such an index m' must exist. Now the balanced partition is $(V \setminus (T_1 \cup \dots \cup T_{m'}), T_1 \cup \dots \cup T_{m'})$. The partition fulfills the stated condition, as all crossing edges are incident to the centroid, which is on the left side. \square

At a high-level, DC-MV works as follows: It *guesses* the partition (V_1, V_2) of the vertex set V according to a balanced separator of the (unknown) optimal rooted tree T , satisfying the conditions of Lemma 2.5. It also *guesses* the distinguished vertex $v \in V_1$ to which all edges that cross the partition are incident. (Throughout the description of the algorithm, guessing should be understood as iterating through all possible choices.)

The balanced separator splits T into a tree with vertex set V_1 and a forest with vertex set V_2 . There are two cases to consider, depending on whether the root $r = x_1$ of T falls in V_1 or V_2 . In the first case, V_1 induces a directed subtree of T rooted at r ; in the second case, V_1 induces a directed subtree of T rooted at v .

Observe that the out-degrees and in-degrees of vertices in V_1 are feasible for a directed tree, except for vertex v , which has additional degrees due to the edges crossing the partition. We refer to these out-degrees and in-degrees as the *excess* of v . The excess of v can be computed using Lemma 2.4. This excess determines the number and orientation of edges across the cut, even if the endpoints, other than v , of the edges are unknown.

By the same argument as in Lemma 2.3, the subtree of T induced by V_1 is optimal for its corresponding degree sequence after subtracting the excess of v ; therefore, we can find it by a recursive call to the procedure.

On the other side of the partition, we have a *collection* of trees. To obtain an instance of our original problem, we add a virtual vertex w to V_2 to play the role of v . For all vertices u , we set the distance between u and the virtual vertex w to be the same as the distance between u and v . We set the out-degree and in-degree of w to the *excess* of v to allow it to connect to all vertices of V_2 that v connects to in T . Now, we can find the optimal tree on this side, too, by a recursive call to the procedure.

ALGORITHM 6: DC-MV for generating the optimal directed tree for a given degree sequence.

Input: Vertex set V , degree sequence $(\delta^{\text{out}}, \delta^{\text{in}})$, and $n = |V|$.

Output: (T, c) , where T is optimum tree for $\delta^{\text{out}}, \delta^{\text{in}}$ and c is its cost.

```

1: procedure  $\mathcal{T}$   $[V, \delta^{\text{out}}, \delta^{\text{in}}]$ 
2:   if  $|V| \leq 5$  then
3:     Directly find optimum tree  $T$  with cost  $c$ .
4:     return  $(T, c)$ 
5:   else
6:     for each partition  $(V_1, V_2)$  of  $V$ , such that  $|V_1|, |V_2| \leq \lceil 2n/3 \rceil$  do
7:       for each  $v \in V_1$  do
8:          $(\text{excess}_v^{\text{out}}, \text{excess}_v^{\text{in}}) \leftarrow \left( \sum_{u \in V_1} \delta^{\text{out}}(u) - |V_1| + 1, \sum_{u \in V_1} \delta^{\text{in}}(u) - |V_1| + 1 \right)$ 
9:          $(\delta^{\text{out}'}, \delta^{\text{in}'}) \leftarrow (\delta^{\text{out}}, \delta^{\text{in}})$ 
10:         $\delta^{\text{out}'}(v) \leftarrow \delta^{\text{out}'}(v) - \text{excess}_v^{\text{out}}, \quad \delta^{\text{in}'}(v) \leftarrow \delta^{\text{in}'}(v) - \text{excess}_v^{\text{in}}$ 
11:         $V_2' \leftarrow V_2 \cup \{w\}$ 
12:         $d_{wu} \leftarrow d_{vu}, d_{uw} \leftarrow d_{uv}$ , for all  $u \in V_2$ 
13:         $(\delta^{\text{out}'}(w), \delta^{\text{in}'}(w)) \leftarrow (\text{excess}_v^{\text{out}}, \text{excess}_v^{\text{in}})$ 
14:        if  $\delta^{\text{out}'}(\cdot), \delta^{\text{in}'}(\cdot) \geq 0$  and  $\delta^{\text{out}'}(v) + \delta^{\text{in}'}(v) \geq 1$  and  $\delta^{\text{out}'}(w) + \delta^{\text{in}'}(w) \geq 1$ 
15:          then
16:             $c_1 \leftarrow \mathcal{T} [V_1, \delta^{\text{out}'}|_{V_1}, \delta^{\text{in}'}|_{V_1}]$ 
17:             $c_2 \leftarrow \mathcal{T} [V_2', \delta^{\text{out}'}|_{V_2'}, \delta^{\text{in}'}|_{V_2'}]$ 
18:             $c \leftarrow c_1 + c_2$ 
19:             $T \leftarrow$  merge  $T_1$  and  $T_2$  by identifying  $v$  and  $w$ 
19:   return  $(T, c)$  with minimum  $c$ 

```

On both sides of the partition, the roots of the directed trees are uniquely determined by the remaining degrees. Observe that if the original root coincides with the centroid v , then v and w will be the roots of the trees in the recursive calls. After obtaining the optimal trees on the two sides (assuming the guesses were correct), we reconstruct T by gluing the two trees together, identifying the vertices v and w . As this operation adds all degrees, we obtain a valid tree for the




Fig. 1. Illustration of DC-MV. (left) Optimal tree and balanced centroid-partitioning, the centroid shown as a hollow circle; (right) optimal trees on the two sides of the partitioning, centroid v and its virtual pair w shown as hollow circles; (above) root is in the left side of the partition, shown as double circle; (below) root is in the right side of the partition, shown as double circle.

original degree sequence; furthermore, the tree must be optimal. We illustrate the two cases of this process in Figure 1 with pseudocode given in Algorithm 6.

The threshold value 5 in Line 2 of Algorithm 6 is to ensure progress; the subproblem size $\lceil 2n/3 \rceil + 1$ (the second term accounts for the virtual vertex w) needs to be less than n , which holds for $n > 5$. For $n \leq 5$, we proceed by a brute-force approach, constructing all feasible trees and choosing one with minimal cost. In Line 8, the excess out-degree and in-degree of v are calculated. Both types of degrees need to sum to $|V_1| - 1$, from which the expression follows. The condition in Line 14 ensures that we only solve feasible problems. The notation $\delta|_X$ indicates a restriction of a degree sequence to a subset X of the vertices.

Analysis of Algorithm 6. For a set V of size n , we consider at most 2^n partitions (Line 6) and at most $\lceil 2n/3 \rceil$ choices of v (Line 7), and we recur on subsets of size at most $\lceil 2n/3 \rceil + 1$. All remaining operations take $O(n)$ time. The run time $t(n)$ of DC-MV thus satisfies:

$$\begin{aligned}
 t(n) &\leq 2^n \cdot n \cdot 2t(\lceil 2n/3 \rceil + 1) \\
 &= n^{O(\log n)} \cdot 2^{n(\sum_k (2/3)^k) + O(\log n)} \\
 &= n^{O(\log n)} \cdot O(2^{3n}) .
 \end{aligned}$$

Including the transportation problem, the overall run time of Algorithm 4 using DC-MV is thus $DS(n) \cdot (8^n n^{O(\log n)} + n^3 \log n) + O(n^3 \log k) = O^*(32^{n+o(n)} + \log k)$. For the space complexity, observe that we do not precompute $\mathcal{T}[\cdot, \cdot, \cdot]$ and, apart from minor bookkeeping, only a single tour is stored at each time, spread over $O(\log n)$ recursive levels. The claimed bounds follow.

2.5 dc-MV2: Polynomial Space and Improved Single-exponential Time

Finally, we modify the divide and conquer algorithm, improving its run time for a given degree sequence, from $8^{n+o(n)}$ to $4^{n+o(n)}$, while maintaining polynomial space.

The result is based on a strengthening of the structural observation about tree-separators (Lemma 2.5). Specifically, we observe that the vertex set of a tree can be partitioned in a “perfectly balanced” way if we allow the edges across the partition to be incident to a *logarithmic* number of vertices on one side (instead of a single vertex as in Lemma 2.5).

Again, let (V_1, V_2) be a partition of $V(T)$, and let $|V(T)| = n$. We say that (V_1, V_2) is *perfectly balanced* if $\max\{|V_1|, |V_2|\} \leq \lceil n/2 \rceil$.

LEMMA 2.6. *Every tree T admits a perfectly balanced partition (V_1, V_2) such that all edges of T between V_1 and V_2 are incident to a vertex in $\{v_1, \dots, v_k\} \subseteq V_1$, where $k \leq \lfloor \log_2 n \rfloor$.*

PROOF. We proceed by finding the special vertices v_j one-by-one for $j = 1, \dots, \lfloor \log_2 n \rfloor$ and updating the partition (V_1, V_2) in each iteration, maintaining the property that all edges between V_1 and V_2 are incident to a vertex in $\{v_1, \dots, v_j\}$. We additionally require the following conditions in every iteration: (1) v_1, \dots, v_j are, in this order, on a simple path entirely within V_1 , and (2) $|V_1| \leq |V_2|$. Furthermore, if $|V_1| < |V_2|$, then we require that there is a subtree T' of T attached to v_j such that (3) T' is entirely contained in V_2 , (4) $|T'| \leq n/2^j$, and (5) $|V_1| + |T'| > n/2$. In words, moving T' to the left side of the partition would make the left side larger than the right side. The claim of the Lemma follows from conditions (1)–(5), when $j = \lfloor \log_2 n \rfloor$.

Let us prove that the conditions hold by induction on j . For the base case, $j = 1$, we reuse our construction from Lemma 2.5. Let v_1 be the centroid of T , and iteratively move the subtrees of T attached to v_1 to the initially empty V_2 from the initially full V_1 , until $|V_2| \geq |V_1|$. Denote the last subtree that is moved as T' . Conditions (1)–(5) are easily verified.

For the inductive step, assume that v_1, \dots, v_j have been found, satisfying conditions (1)–(5), and let T' be the subtree attached to v_j that fulfills conditions (3)–(5). Let v^* be the centroid of T' , and let q_1, \dots, q_m be the internal vertices (in this order) on the path from v_j to v^* . Let $q_{m+1} = v^*$. Iteratively, for $\ell = 1, \dots, m + 1$, move the nodes q_ℓ from V_2 to V_1 , and after moving each q_ℓ , move one-by-one the subtrees of T' attached to q_ℓ , other than the one containing the centroid v^* . By the definition of the centroid, and by (4), all subtrees that we move have size at most $n/2^{j+1}$. We stop when $|V_1| \geq |V_2|$. Observe that this must happen eventually, due to condition (5).

If the last move is a vertex q_ℓ , then $|V_2| \leq |V_1| \leq |V_2| + 1$, so a perfectly balanced partition has been found; we label q_ℓ as v_{j+1} and stop the entire process. In this case, verifying conditions (1)–(5) is no longer needed. If the last move is a subtree, then we move it back to V_2 , label it T' , and label the last moved vertex q_ℓ as v_{j+1} . By construction, all edges across the partition are incident to one of the vertices v_1, \dots, v_{j+1} . This is because intermediate vertices $q_1, \dots, q_{\ell-1}$ have been moved across the partition together with all their pendent trees; they cannot thus be incident to any edges across the partition. (Recall that v_1, \dots, v_{j+1} form a simple path, but are not necessarily consecutive vertices on this path.) By construction, conditions (1)–(5) are satisfied with v_{j+1} , the new T' , and the current V_1, V_2 . (See Figure 2 for illustration.) \square

Our improved algorithm DC-MV2 follows the same outline as DC-MV, described in Section 2.4. We highlight the main differences.




Fig. 2. Illustration of the proof of Lemma 2.6.

The algorithm *guesses* the partition (V_1, V_2) of the vertex set V , according to a perfectly balanced separator of the (unknown) optimal rooted tree T , satisfying the conditions of Lemma 2.6. It also guesses the distinguished vertices $\{v_1, \dots, v_k\} \subseteq V_1$ to which all edges that cross the partition are incident. Again, the separator splits T into a tree with vertex set V_1 and a forest with vertex set V_2 . Again, we consider separately the cases when the root $r = x_1$ of T falls in V_1 or V_2 . In the first case, V_1 induces a directed subtree of T rooted at r ; in the second case, V_1 induces a directed subtree of T rooted at v_i for some $1 \leq i \leq k$. (See Figure 3.)

Whereas in DC-MV all excess degrees belong to a single vertex, they are now distributed among v_1, \dots, v_k . Therefore, we compute the *total excess* of the vertices v_1, \dots, v_k and we *guess* their distribution. Distributing the excess *in-degree* is easy. Assuming that our guesses were correct, at most one of the vertices v_1, \dots, v_k has an incoming edge across the partition. We just need to guess, therefore, which vertex v_i (if any) has an excess in-degree of 1, while all others have excess in-degree 0 (see Figure 3). For the excess *out-degrees*, we need to distribute a total value of at most n among $k = O(\log n)$ vertices, amounting to $n^{O(\log n)}$ possible choices.

We update the degrees of v_1, \dots, v_k in V_1 , removing their excesses. Assuming our guesses to be correct, the subtree of T induced by V_1 is optimal for its degree sequence, by Lemma 2.2.

On the side of V_2 , we have a collection of trees. To obtain an instance of our problem (i.e., a single tree), we add virtual vertices w_1, \dots, w_k to V_2 , to play the role of v_1, \dots, v_k , and a virtual root w_0 forcibly connected to the virtual vertices.

To achieve the correct connections, we first set the out-degrees and in-degrees of w_1, \dots, w_k to match the excesses of v_1, \dots, v_k . This allows them to connect to all vertices of V_2 that v_1, \dots, v_k connect to in T . Observe that the location of the root of T is determined by the in-degrees of the virtual vertices. If all w_i have in-degree 0, then the root is on the left side of the partition (i.e., in V_1). If w_ℓ , for some $\ell \in \{1, \dots, k\}$ has in-degree 1, then the root is on the right side of the partition (i.e., in V_2) in a subtree attached to v_ℓ .

In the first case, we make sure that all edges (w_0, w_i) are directed away from w_0 . In the second case, edge (w_ℓ, w_0) is directed towards w_0 , and all edges (w_0, w_i) for $i = 1, \dots, k$, and $i \neq \ell$ are directed away from w_0 (see Figure 3). To enforce these connections, we appropriately adjust the in-degrees and out-degrees of w_0, \dots, w_k . We also set distances involving virtual vertices w_0, \dots, w_k , such as to forbid unwanted connections. (See Algorithm 8.)

After obtaining the optimal trees on the two sides (assuming the guesses were correct), we reconstruct T by gluing the two trees together, removing w_0 with all its incident edges and identifying v_i with w_i for all $i = 1, \dots, k$. The optimality of the obtained tree follows by observing that a

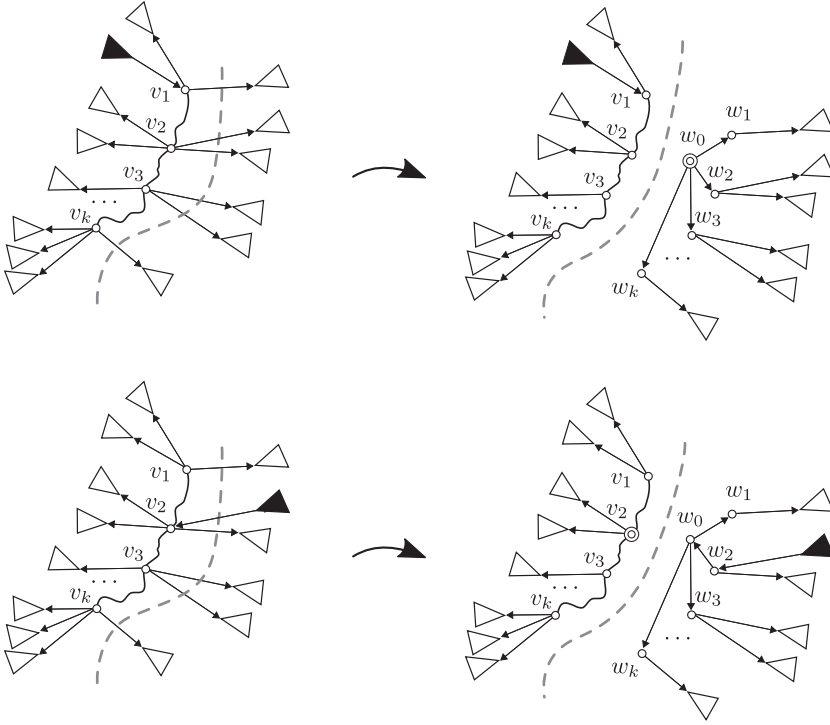


Fig. 3. Illustration of dc-MV2. (left) Optimal tree and perfectly balanced partitioning; triangles indicate subtrees, shaded triangle contains root. (right) Optimal trees on the two sides of the partitioning, virtual vertices w_0, w_1, \dots, w_k shown; (above) root falls to the left side of the partition; (below) root falls to the right side of the partition; new roots shown as double circles.

tree with smaller cost would also induce an improved solution on at least one of the subproblems, contradicting their optimality. Again, we remark that “guessing” means iterating through all possible choices. We refer to the detailed pseudocode described in Algorithm 7 and the illustration in Figure 3.

Some further remarks about Algorithm 7 are in order. The threshold 9 in Line 2 is to ensure progress; we need the recursive subproblem size $\lceil n/2 \rceil + \lceil \log_2 n \rceil + 1$ (the additive logarithmic term is due to the virtual vertices) to be less than n , which holds for $n > 9$. In practice, even some values $n \leq 9$ may be too large to be solved by brute force; in such cases, we can switch back to Algorithm 6 (dc-MV) or Algorithm 5 (dp-MV), without affecting the asymptotic guarantees.

Lines 6, 7, and 9 contain the “guesses,” specifically, for the partition of V , for the choice of v_1, \dots, v_k , and for distributing the excess degrees among v_1, \dots, v_k . Line 12 requires a test similar to the corresponding test in Algorithm 6, that the guessed values correctly describe a tree on both sides of the partition. Line 11 contains the update of degrees and distances on the two sides of the partition, separated into Algorithm 8 for readability.

Analysis of Algorithm 7. For a set V of size n , we consider at most 2^n partitions (Line 6), at most $n^{O(\log n)}$ choices for the v_i (Line 7), and at most $n^{O(\log n)}$ choices for their excess in-degrees and out-degrees (Line 9). We recur on subsets of size at most $\lceil n/2 \rceil + \lceil \log_2 n \rceil + 1$.

ALGORITHM 7: DC-MV2 for generating the optimal directed tree for a given degree sequence.

Input: Vertex set V , degree sequence $(\delta^{\text{out}}, \delta^{\text{in}})$, and $n = |V|$.
Output: (T, c) , where T is optimum tree for $\delta^{\text{out}}, \delta^{\text{in}}$ and c is its cost.

- 1: **procedure** $\mathcal{T} [V, \delta^{\text{out}}, \delta^{\text{in}}]$
- 2: **if** $|V| \leq 9$ **then**
- 3: Directly find optimum tree T with cost c .
- 4: **return** (T, c)
- 5: **else**
- 6: **for** each partition (V_1, V_2) of V , such that $|V_1|, |V_2| \leq \lceil n/2 \rceil$ **do**
- 7: **for** each $\{v_1, \dots, v_k\} \subseteq V_1$, where $k \leq \lfloor \log_2 n \rfloor$ **do**
- 8: $(\text{excess}^{\text{out}}, \text{excess}^{\text{in}}) \leftarrow \left(\sum_{u \in V_1} \delta^{\text{out}}(u) - |V_1| + 1, \sum_{u \in V_1} \delta^{\text{in}}(u) - |V_1| + 1 \right)$
- 9: **for** each *excess-partition* $\left\{ \text{excess}_{v_i}^{\text{out}} \right\}_{1 \leq i \leq k}, \left\{ \text{excess}_{v_i}^{\text{in}} \right\}_{1 \leq i \leq k}$ **do**
- 10: $(\delta^{\text{out}'}, \delta^{\text{in}'}) \leftarrow (\delta^{\text{out}}, \delta^{\text{in}})$
- 11: Update degrees and distances appropriately. ▷ Algorithm 8
- 12: **if** $\delta^{\text{out}'}, \delta^{\text{in}'}$ are *valid* for directed trees **then**
- 13: $c_1 \leftarrow \mathcal{T} [V_1, \delta^{\text{out}'|_{V_1}}, \delta^{\text{in}'|_{V_1}}]$
- 14: $c_2 \leftarrow \mathcal{T} [V_2', \delta^{\text{out}'|_{V_2'}}, \delta^{\text{in}'|_{V_2'}}]$
- 15: $c \leftarrow c_1 + c_2$
- 16: $T \leftarrow$ merge T_1 and T_2 by identifying v_i and w_i and removing w_0
- 17: **return** (T, c) with minimum c

All remaining operations take $O(n)$ time. The run time $t(n)$ can thus be bounded as:

$$\begin{aligned}
 t(n) &\leq 2^n \cdot n^{O(\log n)} \cdot 2t(\lceil n/2 \rceil + \lfloor \log_2 n \rfloor + 1) \\
 &= n^{O(\log^2 n)} \cdot 2^{n(\sum_k (1/2)^k) + O(\log^2 n)} \\
 &= n^{O(\log^2 n)} \cdot O(2^{2n}) .
 \end{aligned}$$

Including the transportation problem, the overall run time of Algorithm 4 using DC-MV2 is therefore $DS(n) \cdot (4^n n^{O(\log n)} + n^3 \log n) + O(n^3 \log k) = O^*(16^{n+o(n)} + \log k)$, with essentially the same space complexity as before.

We remark that in instances of MV-TSP where some of the multiplicities are smaller than $n - 1$, not all degree sequences of trees are realizable. Pruning out unrealizable degree sequences can lead to improvements in efficiency. As a special case, suppose that all multiplicities are equal to 1 (i.e., the standard TSP). Then, the only possible tree is a *path*, with a unique degree sequence, up to the choice of the leaf. In this case the run time of our approach reduces to $4^{n+o(n)}$, matching the Gurevich-Shelah algorithm [21].

3 DISCUSSION

We described three new algorithms, along with variations, for the many-visits TSP problem. In particular, we showed how the problem can be solved in time single-exponential in the number n of cities, while using space only polynomial in n . This yields the first improvement over the Cosmadakis-Papadimitriou algorithm in more than 35 years.

It remains an interesting open question to improve the bases of the exponentials in our run times, even in the special case when all edge costs are equal to 1 or 2. Such instances of MV-TSP

arise, e.g., in the MAXIMUM SCATTER TSP application by Kozma and Mömke [37]. Recent algebraic techniques [19, 36, 41] may be of help.

In practice, one may reduce the search space of our algorithms via heuristics, for instance, by forcing certain (directed) edges to be part of the solution. An edge (i, j) is part of the solution if the optimal tour visits it *at least once*. This may be reasonable if the edges in question are very cheap.

ALGORITHM 8: Updating degrees and distances in DC-MV2: Line 12 of Algorithm 7.

```

1:  $\delta^{\text{out}'}(v_i) \leftarrow \delta^{\text{out}'}(v_i) - \text{excess}_{v_i}^{\text{out}}, \quad \delta^{\text{in}'}(v_i) \leftarrow \delta^{\text{in}'}(v_i) - \text{excess}_{v_i}^{\text{in}}, \quad \text{for } 1 \leq i \leq k$ 
2:  $V'_2 \leftarrow V_2 \cup \{w_0, w_1, \dots, w_k\}$  ▷ virtual vertices
3:  $\delta^{\text{out}'}(w_i) \leftarrow \text{excess}_{v_i}^{\text{out}}, \quad \delta^{\text{in}'}(w_i) \leftarrow \text{excess}_{v_i}^{\text{in}} + 1, \quad \text{for all } 1 \leq i \leq k$ 
4:  $d_{w_0, w_i} \leftarrow 0, \quad d_{w_i, w_0} \leftarrow 0, \quad \text{for all } 1 \leq i \leq k$ 
5:  $d_{w_0, u} \leftarrow \infty, \quad d_{u, w_0} \leftarrow \infty, \quad \text{for all } u \in V_2 \setminus \{w_1, \dots, w_k\}$ 
6:  $d_{w_i, w_j} \leftarrow \infty, \quad \text{for all } 1 \leq i, j \leq k$ 
7:  $d_{w_i, u} \leftarrow d_{v_i, u}, \quad \text{and } d_{u, w_i} \leftarrow d_{u, v_i} \quad \text{for all } 1 \leq i \leq k, \quad \text{and } u \in V_2 \setminus \{w_0, \dots, w_k\}$ 
8:  $\delta^{\text{out}'}(w_0) \leftarrow k, \quad \delta^{\text{in}'}(w_0) \leftarrow 0$ 
9: if  $\text{excess}_{v_i}^{\text{in}} > 0$  for some  $i \in \{1, \dots, k\}$  then ▷ root is in  $V_2$ , in subtree attached to  $v_i$ 
10:    $\delta^{\text{out}'}(w_i) \leftarrow \delta^{\text{out}'}(w_i) + 1$ 
11:    $\delta^{\text{in}'}(w_i) \leftarrow \delta^{\text{in}'}(w_i) - 1$ 
12:    $\delta^{\text{out}'}(w_0) \leftarrow \delta^{\text{out}'}(w_0) - 1$ 
13:    $\delta^{\text{in}'}(w_0) \leftarrow \delta^{\text{in}'}(w_0) + 1$ 

```

APPENDIX

A DEFERRED SUBROUTINES

ALGORITHM 9: Generating all possible r -subsets of $\{1, \dots, n\}$.

```

1: Input: Positive integers  $n$  and  $r$ .
2: Output: A generator of all possible  $r$ -subsets of  $\{1, \dots, n\}$ .
3: procedure COMBINATIONS( $n, r$ )
4:    $\text{comb} \leftarrow [1, \dots, r]$ 
5:   yield  $\text{comb}$ 
6:   while true do
7:      $i \leftarrow$  last index such that  $\text{comb}_i \neq n - r + i$ , if no such index exists, break
8:      $\text{comb}_i \leftarrow \text{comb}_i + 1$ 
9:     for every index  $j \leftarrow i+1, \dots, r$  do
10:        $\text{comb}_j \leftarrow \text{comb}_{j-1} + 1$ 
11:   yield  $\text{comb}$ 

```

The algorithm COMBINATIONS implements algorithm NEXKSB by Nijenhuis and Wilf [43, page 26]. It takes two integers as input, n and r , and generates all (ordered) subsets of size r , of the base set $\{1, \dots, n\}$. It starts with the set $[1, 2, \dots, r]$ and generates all r -subsets in lexicographical order, up to $[n - r + 1, n - r + 2, \dots, n - r + r] = [n - r + 1, n - r + 2, \dots, n]$. In every iteration, it increments the rightmost entry comb_i not equal to $n - r + i$, and makes comb_j equal to $\text{comb}_{j-1} + 1$ for all j indices between $[i + 1, r]$. The algorithm stops when there are no such indices i , which happens when reaching $[n - r + 1, \dots, n]$.




Fig. 4. Sequence $[2, 4, 7, 8]$ representing the degree sequence $[1, 1, 2, 0, 0]$.

An example is shown in Figure 4. The corresponding integer sequence would be $[1, 1, 2, 0, 0]$; however, COMBINATIONS returned the sequence $[2, 4, 7, 8]$ —that is, a sequence of the positions (separating bars) of the integer sequence above. To obtain the actual degree sequence, we use a short script COMBINATIONTOSEQUENCE in Algorithm 10 that converts the sequence with the positions of the bars to the degree sequence.

ALGORITHM 10: Converting $[a_1, \dots, a_m]$ into $[a_1, a_2 - a_1 - 1, \dots, r + m - a_m]$.

```

1: Input: List of positions  $a = [a_1, \dots, a_m]$ , integer  $r$ .
2: Output: A sequence of  $m + 1$  integers that sum up to  $r$ .
3: procedure COMBINATIONTOSEQUENCE( $a, r$ )
4:    $\text{seq}_1 \leftarrow a_1 - 1$ 
5:   for every index  $i \leftarrow 2, \dots, m$  do
6:      $\text{seq}_i \leftarrow a_i - a_{i-1} - 1$ 
7:    $\text{seq}_{m+1} \leftarrow r + m - a_m$ 
8:   return seq

```

Finally, the procedure DISTRIBUTE(n, k) calls COMBINATIONTOSEQUENCE(a, n) for each output a of COMBINATIONS($n + k, n$).

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their many insightful remarks.

REFERENCES

- [1] Hasan Abasi, Nader H. Bshouty, Ariel Gabizon, and Elad Haramaty. 2014. On r -simple k -path. In *Proceedings of the MFCS*. 1–12.
- [2] Maichel M. Aguayo, Subhash C. Sarin, and Hanif D. Sherali. 2018. Single-commodity flow-based formulations and accelerated Benders algorithms for the high-multiplicity asymmetric traveling salesman problem and its extensions. *J. Oper. Res. Soc.* 69, 5 (2018), 734–746.
- [3] David L. Applegate, Robert E. Bixby, Vašek Chvátal, and William J. Cook. 2006. *The Traveling Salesman Problem: A Computational Study*. Princeton University Press.
- [4] Esther M. Arkin, Yi-Jen Chiang, Joseph S. B. Mitchell, Steven Skiena, and Tae-Cheon Yang. 1999. On the maximum scatter traveling salesperson problem. *SIAM J. Comput.* 29, 2 (1999), 515–544.
- [5] Jørgen Bang-Jensen and Gregory Z. Gutin. 2002. *Digraphs—Theory, Algorithms, and Applications*. Springer.
- [6] Richard Bellman. 1962. Dynamic programming treatment of the travelling salesman problem. *J. Assoc. Comput. Mach.* 9 (1962), 61–63.
- [7] Claude Berge. 1973. *Graphs and Hypergraphs*. North-Holland Publishing Co., Amsterdam-London; American Elsevier Publishing Co., Inc., New York. xiv+528 pages. Translated from the French by Edward Minieka, North-Holland Mathematical Library, Vol. 6.
- [8] Annabell Berger. 2014. A note on the characterization of digraphic sequences. *Discrete Math.* 314 (2014), 38–41.
- [9] Nadia Brauner, Yves Crama, Alexander Grigoriev, and Joris van de Klundert. 2005. A framework for the complexity of high-multiplicity scheduling problems. *J. Combin. Optim.* 9, 3 (2005), 313–323.
- [10] Arthur Cayley. 1889. A theorem on trees. *Quart. J. Pure Appl. Math.* 23 (1889), 376–378.
- [11] Nicos Christofides. 1976. *Worst-case Analysis of a New Heuristic for the Travelling Salesman Problem*. Technical Report 388. Carnegie Mellon University.

- [12] W. J. Cook. 2011. *In Pursuit of the Traveling Salesman: Mathematics at the Limits of Computation*. Princeton University Press.
- [13] Stavros S. Cosmadakis and Christos H. Papadimitriou. 1984. The traveling salesman problem with many visits to few cities. *SIAM J. Comput.* 13, 1 (1984), 99–108.
- [14] Jack Edmonds and Richard M. Karp. 1970. Theoretical improvements in algorithmic efficiency for network flow problems. In *Proceedings of the International Conference on Combinatorial Structures and their Applications*. Gordon and Breach, New York, 93–96.
- [15] Hamilton Emmons and Kamlesh Mathur. 1995. Lot sizing in a no-wait flow shop. *Oper. Res. Lett.* 17, 4 (1995), 159–164.
- [16] Pál Erdős and Tibor Gallai. 1960. Graphs with prescribed degrees of vertices (Hungarian). *Mat. Lapok* 11 (1960), 264–274.
- [17] Delbert R. Fulkerson. 1960. Zero-one matrices with zero trace. *Pacific J. Math.* 10 (1960), 831–836.
- [18] Ariel Gabizon, Daniel Lokshtanov, and Michał Pilipczuk. 2015. Fast algorithms for parameterized problems with relaxed disjointness constraints. In *Proceedings of the ESA*. 545–556.
- [19] Alexander Golovnev. 2014. Approximating asymmetric TSP in exponential time. *Int. J. Found. Comput. Sci.* 25, 1 (2014), 89–100.
- [20] Alexander Grigoriev and Joris van de Klundert. 2006. On the high multiplicity traveling salesman problem. *Discrete Optim.* 3, 1 (2006), 50–62.
- [21] Yuri Gurevich and Saharon Shelah. 1987. Expected computation time for Hamiltonian path problem. *SIAM J. Comput.* 16, 3 (1987), 486–502.
- [22] Gregory Gutin and Abraham Punnen. 2002. *The Traveling Salesman Problem and Its Variations*. Springer.
- [23] S. Louis Hakimi. 1962. On realizability of a set of integers as degrees of the vertices of a linear graph. *J. Soc. Indust. Appl. Math.* 10 (1962), 496–506.
- [24] Václav Havel. 1955. Eine Bemerkung über die Existenz der endlichen Graphen. *Časopis Pěst. Mat.* 80 (1955), 477–480.
- [25] Michael Held and Richard M. Karp. 1962. A dynamic programming approach to sequencing problems. *J. Soc. Indust. Appl. Math.* 10 (1962), 196–210.
- [26] Frank L. Hitchcock. 1941. The distribution of a product from several sources to numerous localities. *J. Math. Phys. Mass. Inst. Tech.* 20 (1941), 224–230.
- [27] Dorit S. Hochbaum and Ron Shamir. 1991. Strongly polynomial algorithms for the high multiplicity scheduling problem. *Oper. Res.* 39, 4 (1991), 648–653.
- [28] Russell Impagliazzo and Ramamohan Paturi. 2001. On the complexity of k -SAT. *J. Comput. Syst. Sci.* 62, 2 (2001), 367–375.
- [29] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. 2001. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.* 63, 4 (2001), 512–530.
- [30] Camille Jordan. 1869. Sur les assemblages de lignes. *J. reine angew. Math.* 70 (1869), 185–190.
- [31] Ravi Kannan. 1983. Improved algorithms for integer programming and related lattice problems. In *Proceedings of the STOC*. 193–206.
- [32] Sanjiv Kapoor and H. Ramesh. 1995. Algorithms for enumerating all spanning trees of undirected and weighted graphs. *SIAM J. Comput.* 24, 2 (1995), 247–265.
- [33] Hyunju Kim, Zoltán Toroczkai, Péter L. Erdős, István Miklós, and László A. Székely. 2009. Degree-based graph construction. *J. Phys. A* 42, 39 (2009), 392001, 10.
- [34] Peter Kleinschmidt and Heinz Schannath. 1995. A strongly polynomial algorithm for the transportation problem. *Math. Program.* 68 (1995), 1–13.
- [35] Daniel J. Kleitman and D. L. Wang. 1973. Algorithms for constructing graphs and digraphs with given valences and factors. *Discrete Math.* 6 (1973), 79–88.
- [36] Mikko Koivisto and Pekka Parviainen. 2010. A space-time tradeoff for permutation problems. In *Proceedings of the SODA*. 484–492.
- [37] László Kozma and Tobias Mömke. 2017. Maximum scatter TSP in doubling metrics. In *Proceedings of the SODA*. 143–153.
- [38] Michael Lampis. 2012. Algorithmic meta-theorems for restrictions of treewidth. *Algorithmica* 64, 1 (2012), 19–37.
- [39] E. L. Lawler, D. B. Shmoys, A. H. G. R. Kan, and J. K. Lenstra. 1985. *The Traveling Salesman Problem*. John Wiley & Sons.
- [40] H. W. Lenstra, Jr. 1983. Integer programming with a fixed number of variables. *Math. Oper. Res.* 8, 4 (1983), 538–548.
- [41] Daniel Lokshtanov and Jesper Nederlof. 2010. Saving space by algebraization. In *Proceedings of the STOC*. 321–330.
- [42] J. W. Moon. 1970. *Counting Labelled Trees*. Canadian Mathematical Congress.
- [43] Albert Nijenhuis and Herbert S. Wilf. 1978. *Combinatorial Algorithms* (2nd ed.). Academic Press, Inc. [Harcourt Brace Jovanovich, Publishers], New York-London. xv+302 pages.
- [44] T. Nishizeki and N. Chiba. 1988. *Planar Graphs: Theory and Algorithms*. Elsevier Science.

- [45] James B. Orlin. 1993. A faster strongly polynomial minimum cost flow algorithm. *Oper. Res.* 41, 2 (1993), 338–350.
- [46] Christos H. Papadimitriou and Mihalis Yannakakis. 1981. On minimal Eulerian graphs. *Inf. Proc. Lett.* 12, 4 (1981), 203–205.
- [47] Harilaos N. Psaraftis. 1980. A dynamic programming approach for sequencing groups of identical jobs. *Oper. Res.* 28, 6 (1980), 1347–1359.
- [48] Michael Rothkopf. 1966. Letter to the editor—The traveling salesman problem: On the reduction of certain large problems to smaller ones. *Oper. Res.* 14, 3 (1966), 532–533.
- [49] Subhash C. Sarin, Hanif D. Sherali, and Liming Yao. 2011. New formulation for the high multiplicity asymmetric traveling salesman problem with application to the Chesapeake problem. *Optim. Lett.* 5, 2 (2011), 259–272.
- [50] Jack A. A. van der Veen and Shuzhong Zhang. 1996. Low-complexity algorithms for sequencing jobs with a fixed number of job-classes. *Comput. Oper. Res.* 23, 11 (1996), 1059–1067.

Received July 2019; revised January 2020; accepted February 2020