

VU Research Portal

A Two-Step Warm Start Method Used for Solving Large-Scale Stochastic Mixed-Integer Problems

Markhorst, Berend; Leitner, Markus; Berkhout, Joost; Zocca, Alessandro; Mei, Rob van der

2024

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Markhorst, B., Leitner, M., Berkhout, J., Zocca, A., & Mei, R. V. D. (2024). *A Two-Step Warm Start Method Used for Solving Large-Scale Stochastic Mixed-Integer Problems*.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

A Two-Step Warm Start Method Used for Solving Large-Scale Stochastic Mixed-Integer Problems

Berend Markhorst^{a,b,*}, Markus Leitner^c, Joost Berkhout^b, Alessandro Zocca^b, Rob van der Mei^{a,b}

^a*CWI stochastics department, Science Park 123, 1098 XG, Amsterdam, The Netherlands*

^b*VU mathematics department, De Boelelaan 1105, Amsterdam, 1081 HV, The Netherlands*

^c*Department of Operations Analytics, Vrije Universiteit Amsterdam, De Boelelaan 1105, Amsterdam, 1081 HV, The Netherlands*

Abstract

Two-stage stochastic programs become computationally challenging when the number of scenarios representing parameter uncertainties grows. Motivated by this, we propose the TULIP-algorithm (“*Two-step warm start method Used for solving Large-scale stochastic mixed-Integer Problems*”), a two-step approach for solving two-stage stochastic (mixed) integer linear programs with an exponential number of constraints. In this approach, we first generate a reduced set of representative scenarios and solve the root node of the corresponding integer linear program using a cutting-plane method. The generated constraints are then used to accelerate solving the original problem with the full scenario set in the second phase. We demonstrate the generic effectiveness of TULIP on two benchmark problems: the Stochastic Capacitated Vehicle Routing Problem and the Two-Stage Stochastic Steiner Forest Problem. The results of our extensive numerical experiments show that TULIP yields significant computational gains compared to solving the problem directly with branch-and-cut.

Keywords: Scenario Reduction, Branch-and-cut, Stochastic Programming, Mathematical Optimization, Two-Stage Stochastic Steiner Forest, Stochastic Capacitated Vehicle Routing Problem

1. Introduction

Decision-making under uncertainty is used in many real-world problems, where the objective is to make optimal choices despite unpredictable future events. This uncertainty is prevalent in various fields within mathematical optimization, such as network design [1], supply chain planning [2], energy markets [3], airport operations and scheduling [4], and inventory management [5], where decisions must be made today while considering potential future outcomes, which are often uncertain. To address such complexities, mathematical frameworks such as Stochastic Programming (SP) have been widely used to incorporate uncertainties directly into the decision-making process, thus enabling more robust and informed decisions, see [6, 7] for an overview.

In SP problems, the parameters subject to uncertainty have underlying distributions. In case these distributions are continuous, they cannot be exactly embedded within (mixed-)integer linear programs ((M)ILP). To deal with this issue explicitly, these uncertain outcomes can be discretized into a finite set, for example with Sample Average Approximation [8], and represented by a scenario tree, see [9, Section 1.2.2]. Each branch in this tree corresponds to an outcome of the uncertain parameters. One of the

main challenges in solving SP problems is the need to consider a large number of scenarios to accurately capture the underlying distributions. These problems become computationally expensive and often intractable to solve as the number of scenarios increases, especially when the uncertainty spans a high-dimensional space.

A typical approach in the literature for this problem is the use of decomposition methods, such as the integer L-shaped method [10] based on Benders’ decomposition, see [11] for an extensive overview. Another suitable approach is Lagrangian relaxation, see [12] for an introduction and [13] for an overview and a related method for network design problems is dual ascent [14]. With distribution- [15] and problem-based [16] scenario generation, one represents the scenario tree with a representable, but strongly reduced subset of scenarios, which can be used to approximate the solution of the original problem.

Among the various stochastic programming models, two-stage SP models are often used, see [17] for an illustrative introduction. In this framework, decisions are made in two stages: a set of initial decisions is made before the uncertainty is revealed, followed by corrective actions once the uncertain parameters are known. This structure allows decision-makers to balance the trade-off between the costs of initial decisions and the expected costs of future corrective actions, providing an effective way to optimize in the presence of uncertainty.

To improve the tractability of exact solution meth-

*Corresponding author

Email address: berend.markhorst@cwi.nl (Berend Markhorst)

ods for two-stage integer SP problems, in this work, we propose a generic and effective method called “*Two-step warm start method Used for solving Large-scale stochastic mixed-Integer Problems*” (TULIP). In the first step, we identify a representative yet relatively small subset of scenarios with [15, Algorithm 2.4] and solve the root node of the corresponding problem. In the second step, we use the information gathered in the first step to accelerate the solution of the original problem, corresponding to the whole scenario tree. By applying this method to the Stochastic Capacitated Vehicle Routing Problem (SCVRP) [6, Section 1.5] and the Two-Stage Stochastic Steiner Forest Problem (2S-SSFP) [18], we show the computational gains that our method yields compared to solving the whole problem at once.

TULIP performs well given two assumptions. First, we assume that the corresponding model contains integer decision variables, which yields a (M)ILP. Second, we assume that the formulation of the problem at hand contains an exponential number of constraints both in the first and second stages, which we add dynamically using branch-and-cut.

Contribution. In this work, we contribute to the existing literature as follows:

- We propose a novel combination of methods, called TULIP, to solve large-scale instances of (mixed-)integer two-stage stochastic programming models efficiently to optimality.
- We perform an extensive computational study to analyze the generic performance of this framework for two benchmark problems, SCVRP and 2S-SSFP, and show that TULIP outperforms the benchmark methods.

Outline. The remainder of this paper is structured as follows. We describe the relevant literature in Section 2, and introduce the notation and present our method in Section 3. Then, in Sections 4 and 5, we describe the two problems, SCVRP and 2S-SSFP, on which we then test the performance of our algorithm compared to benchmark method(s) and analyze TULIP’s robustness. Finally, we summarize our findings and give directions for future research in Section 6.

2. Related literature

As mentioned in Section 1, the literature describes several methods that deal with a large number of scenarios. Our proposed method does the same by building upon several foundational approaches in the literature, specifically those introduced by [15, 19, 20]. We elaborate on these methods to provide a context for our contribution and demonstrate how they collectively inform our approach.

Scenario reduction – also referred to as scenario generation in the literature – for stochastic programming problems can generally be approached using either distribution-based (e.g., [15, 21]) or problem-based methods [16, 22]. Distribution-based approaches focus on replicating the true distribution independent of the specific problem being solved. This makes them broadly applicable and suited to our needs, as we aim for a generic optimization method for two-stage (mixed-)integer programs. Problem-based scenario generation methods, on the other hand, are tailored to specific problems and can result in smaller scenario trees while maintaining solution quality. However, they require problem-specific insights, which can complicate their application. Our work uses the fast-forward selection method from [15] as it strikes a good balance between computational efficiency and accuracy. Next to fast-forward selection, the authors also presented backward reduction, to reduce the number of scenarios while retaining a good approximation of the original distribution. The algorithms leverage the Fortet-Mourier probability metric [23] to evaluate stability and computational feasibility, showing significant improvements compared to earlier methods.

In [19], the authors propose a warm start technique for improving the efficiency of solving large-scale stochastic programming problems using interior point methods. Their approach involves generating an initial solution – also referred to as a warm start point – by solving a smaller version of the original stochastic problem and mapping its solution to the full problem. The authors demonstrated considerable gains in run time and the number of iterations needed for convergence. Although subsequent work by [24] further expanded on these ideas, the warm start technique remains underexplored in the literature, which highlights the potential for new advancements, especially in other solving methods than interior point methods.

Our approach benefits from the findings of [20], which improved upon the integer L-shaped method [10] by introducing two key strategies, of which we use one in our proposed method. The authors alternated between linear relaxations and mixed-integer subproblems to evaluate second-stage costs, allowing for faster elimination of non-optimal solutions. This improvement significantly enhanced computational efficiency and convergence speed, particularly for large-scale problems. We use this concept tailored to the integer L-shaped method in the first step of TULIP as we explain in more detail in Section 3.2.

In our proposed method, see Section 3 for an elaborate description, we incorporate and extend these strategies as follows. Based on the ideas from [19, 20], we reuse cuts generated in the root node of the reduced problem as constraints to warm start the optimization of the original (M)ILP. Yet, we apply this method to the branch-and-cut method to solve large-scale two-stage (mixed-)integer recourse models. Additionally, we use fast forward selection from [15] to efficiently reduce the scenario tree.

3. Methodology

We first give a brief introduction to two-stage stochastic programming in Section 3.1 and then describe our method in Section 3.2.

3.1. Stochastic programming framework

Nowadays many integer decision-making problems are solved through deterministic mathematical optimization (DO), which entails minimizing (or maximizing) an objective function under a set of fixed constraints, see [25] for an introduction. Mathematically, we denote the linear variant as follows, for a problem with $n = n_1 + n_2$ decision variables and m constraints, where $n_1 > 0$:

$$\begin{aligned} \text{(DO)} \\ \min_x \quad & c^\top x \end{aligned} \quad (1a)$$

$$\text{s.t.} \quad x \in X^{(\text{DO})}, \quad (1b)$$

where

$$X^{(\text{DO})} = \left\{ \begin{array}{l} Ax \geq b, \\ x \in \mathbb{Z}_+^{n_1} \times \mathbb{R}_+^{n_2} \end{array} \right\} \quad (2)$$

captures the set of feasible solutions for the DO problem, $x \in \mathbb{Z}_+^{n_1} \times \mathbb{R}_+^{n_2}$ is the vector of decision variables, $c \in \mathbb{R}^n$ is the cost vector, $A \in \mathbb{R}^{m \times n}$ represents the technology matrix, and $b \in \mathbb{R}^m$ is the right-hand-side vector. This setup assumes that every parameter in A , b , and c is precisely known upfront.

However, in reality, uncertainty is an omnipresent factor that significantly influences decision quality. A suitable approach for this problem is Stochastic Programming (SP), which is a framework that models decision processes under uncertainty, see [6, 7] for an introduction. SP allows decision-makers to optimize not just for a single outcome but for a spectrum of possible future, uncertain states, which we will elaborate on in the following paragraphs.

The benefit of SP compared to DO is that it can deal with unforeseen parameter variations that might influence a solution's feasibility and optimality. Corrective actions are executed after the parameter uncertainty is revealed to restore feasibility and optimality. This adaptability is captured within SP models, which can be split into two- and multi-stage models.

In this work, we focus on solving large-scale two-stage (mixed-)integer SP models. These models are designed to optimize decisions across two sequential stages: in the first stage, decisions are made before the uncertainty is realized; in the second stage, once uncertainty has been revealed, corrective actions (second-stage decisions) are employed to adapt to the new circumstances.

We refer to the finite set of uncertain parameter outcomes in the second stage as scenarios. We capture the indices corresponding to the finite set of scenarios in the set $\mathcal{S} = \{1, \dots, S\}$. Each scenario corresponds to a unique combination of uncertain parameters, thus necessitating

separate second-stage decisions. The objective of an SP problem is then to find optimal decisions that perform well across all scenarios.

In line with (1), we formulate SP mathematically as follows, given a problem with n_s decision variables per scenario $s \in \mathcal{S}$:

$$\begin{aligned} \text{(SP)} \\ \min \quad & c^\top x + \sum_{s \in \mathcal{S}} p^{(s)} q_{(s)}^\top y^{(s)} \end{aligned} \quad (3a)$$

$$\text{s.t.} \quad x \in X^{(\text{DO})} \quad (3b)$$

$$y^{(s)} \in Y_{(s,x)}^{(\text{DO})} \quad \forall s \in \mathcal{S}. \quad (3c)$$

We denote parameters and decision variables corresponding to a specific scenario $s \in \mathcal{S}$ with sub- or superscript (s). We represent the probability corresponding to scenario $s \in \mathcal{S}$ with $p^{(s)} \in \mathbb{R}_+^S$ and $\sum_{s \in \mathcal{S}} p^{(s)} = 1$, whereas $q_{(s)} \in \mathbb{R}^{n_s}$ and $y^{(s)} \in \mathbb{Z}_+^{n_s}$ capture the second-stage cost vector and the second-stage decision vector describing recourse actions corresponding to scenario $s \in \mathcal{S}$, respectively. Consequently, $\sum_{s \in \mathcal{S}} p^{(s)} q_{(s)}^\top y^{(s)}$ represents the recourse costs, i.e., the costs of the corrective actions. We capture the feasible set for the second-stage decision variables corresponding to scenario $s \in \mathcal{S}$ given the first-stage decision $x \in X^{(\text{DO})}$ in the set $Y_{(s,x)}^{(\text{DO})}$. The objective in (3a) shows that the recourse costs both depend on the uncertainty corresponding to the scenarios and the first-stage decision vector.

One of the primary challenges of SP lies in computational tractability. As the number of scenarios increases, the problem can quickly yield an exceedingly large (mixed-) integer linear program [16]. To address this issue, we propose a novel method that combines scenario reduction and warm starting, which we will elaborate on in the following.

3.2. Introducing TULIP

After providing our assumptions, we will describe our method, “*Two-step warm start method Used for solving Large-scale stochastic mixed-Integer Problems*” (TULIP), and elaborate on how it can be integrated with traditional branch-and-cut strategies to solve two-stage recourse problems more efficiently.

Method assumptions. As we want to provide a framework that speeds up the traditional branch-and-cut method, we make the following two assumptions:

1. We assume that the corresponding model contains integer decision variables, which yields a (mixed-) integer linear program, making the problem non-convex and hence not solvable through the simplex method.
2. We assume that the problem contains an exponential number of constraints both in the first and second stages.

We now explain in detail the steps of the TULIP method, which are also schematically summarized in Figure 1. The first step corresponds to a reduced version of the problem and yields two substeps, whereas the second step corresponds to the original problem.

Step 1a: Scenario reduction. The blue box in Figure 1 corresponds to the original problem. We reduce the number of scenarios by selecting a subset of the original scenarios while preserving the problem’s inherent structure and complexity using the fast forward scenario selection technique from [15, Algorithm 2.4]. This algorithm iteratively selects one scenario to include in a subset of scenarios that best represents the original scenario tree until a fixed number of desired scenarios is reached. This selection is made by minimizing a cost function that measures the difference between the reduced and the full scenario sets. The method uses a distance metric $d(i, j)$ to quantify the difference between scenarios $i \in \mathcal{S}$ and $j \in \mathcal{S}$. This metric depends on the optimization problem at hand and is discussed in more detail in the case studies in Sections 4 and 5. The probabilities of the excluded scenarios are reassigned to the nearest included scenario, computed with the distance metric so that the probabilities of the reduced scenarios still sum to one.

Step 1b: Solve root node. We solve the root node of the reduced problem and store the generated dynamically added constraints, also referred to as cuts. We choose the root node as it can be solved much faster than the whole corresponding ILP while still providing many cuts that become useful in the final step of our method.

Step 2: Warm start. We add the stored constraints that are tight in the root node of the reduced problem as constraints to the ILP of the original problem, marked by the blue area in Figure 1, which we call *warm starting* in this work, and solve it.

With this approach, we do not discard any feasible or optimal solutions from the original ILP. For the cuts on the second-stage decision variables, this claim is trivial as they can be directly transferred from the reduced to the original problem. As the cuts on the first-stage decision variables are added dynamically to the problem, independently of the second stage, this claim holds as well. We note that our method is similar to [19], but we apply it to branch-and-cut instead of interior point methods.

4. Case study I: Stochastic Capacitated Vehicle Routing Problem

The stochastic capacitated vehicle routing problem (SCVRP) [6, Chapter 1.5] is a variation of the traditional vehicle routing problem, a well-studied problem in the literature, see [26] for an overview. In [27], the authors discuss the state-of-the-art for the SCVRP, see [28] for an




Figure 1: Schematic representation of TULIP. The dashed arrows correspond to steps in this method. The gray box represents the whole TULIP framework, whereas the blue and red boxes denote the original and reduced problem, respectively.

integer L-shaped algorithm for the SCVRP. For this problem, we use the ILP formulation from [29].

4.1. Problem description

In the variant of the SCVRP considered in [6, Chapter 1.5], a truck with capacity C must visit a set of $n + 1$ cities $\mathcal{V} = \{0, 1, 2, \dots, n\}$, where city 0 is the depot. The set of cities excluding the depot is denoted by $\mathcal{V}^* = \mathcal{V} \setminus \{0\}$. The distance between city i and city j is indicated by $d_{ij} > 0$. Each city has demand, which is subject to uncertainty and needs to be fulfilled by the truck. Let \mathcal{S} be a finite set of scenarios, with scenario $s \in \mathcal{S}$ occurring with probability $p^{(s)}$. For each scenario $s \in \mathcal{S}$, we denote by $b_i^{(s)} > 0$ the uncertain demand in city $i \in \mathcal{V}$. The objective is to minimize the expected total distance traveled, taking into account the stochastic nature of the demands and ensuring that the vehicle’s capacity is not exceeded. We model the problem using two stages: the truck’s route is identified upfront, after which the uncertain demand is revealed. Then, the recourse actions in every city after a visit describe whether to make an additional trip to the depot or not.

4.2. Illustrative example

To provide some intuition for the SCVRP, we include a small illustrative example from [6, Chapter 1.5]. The truck must start at the depot, visit four cities (1, 2, 3, and 4), and end at the depot. The truck’s capacity is $C = 10$ units. We visualize this problem in Figure 2a with a complete graph showing our assumption that the truck can travel from every city to another. The demands for cities 1, 2, and 4 are known and equal 2 units each. For city 3, the demand is stochastic and equal to either one or seven units with equal probability. The distance matrix for this example is given in [6, Table 7]. The optimal route for this instance starts with going to city 3. Taking into account also the optimal recourse action, then if the demand at city 3 turns out to be one, we follow the route

as shown in Figure 2b, and otherwise, the route as shown in Figure 2c.




Figure 2: Small instance from [6, Chapter 1.5] to illustrate the SCVRP. The numbers in bold denote the cities, whereas the numbers in parenthesis represent the (uncertain) demands. The probabilities indicate that the two scenarios are equiprobable.

4.3. Model formulation

In line with the build-up of the model formulation in Section 3, we first introduce the DO model formulation of the vehicle routing problem as an ILP problem. The binary decision variable $x_{ij} \in \{0, 1\}$ with $i \in \mathcal{V}$ and $j \in \mathcal{V}$ is equal to 1 if the truck travels from city i to city j , and 0 otherwise.

$$\begin{aligned}
 \text{(DO-1a)} \quad & \min \sum_{i \in \mathcal{V}} \sum_{j \in \mathcal{V}} d_{ij} x_{ij} & (4a) \\
 \text{s.t.} \quad & \sum_{j \in \mathcal{V}, j \neq i} x_{ij} = 1 & \forall i \in \mathcal{V}^* & (4b) \\
 & \sum_{i \in \mathcal{V}, i \neq j} x_{ij} = 1 & \forall j \in \mathcal{V}^* & (4c) \\
 & \sum_{j \in \mathcal{V}^*} x_{0j} \geq 1 & (4d) \\
 & \sum_{i \in \mathcal{V}^*} x_{i0} \geq 1 & (4e) \\
 & \sum_{i \in Q} \sum_{j \in Q} x_{ij} \leq |Q| - 1 & \begin{cases} \forall Q \subset \mathcal{V}^*, \\ Q \neq \emptyset \end{cases} & (4f) \\
 & x_{ij} \in \mathbb{B} & \begin{cases} \forall i \in \mathcal{V}, \\ \forall j \in \mathcal{V}. \end{cases} & (4g)
 \end{aligned}$$

The goal in (4a) is to minimize the costs of the VRP route. We make sure that every city is visited once in (4b)-(4c) – the depot can be visited multiple times, see (4d) and (4e) – and add subtour elimination constraints (4f) to

ensure that the solution consists of routes that are connected to the depot [29]. We add this constraint through branch-and-cut by making a support graph based on x_{ij} -values and checking for each combination of depot and city if there is a subtour. If so, we add the constraint as a cut in the branch-and-cut procedure. Integrality constraints are captured in constraint (4g). For the extension of DO to SP, we introduce the second-stage binary decision variable $y_{ij}^{(s)} \in \{0, 1\}$ with $i \in \mathcal{V}$, $j \in \mathcal{V}$, and $s \in \mathcal{S}$, which is a binary variable equal to 1 if we *actually* travel from city i to city j in scenario s , and 0 if we do not. To include capacity in the subtour elimination constraints, we can change line (4f) with the well-known capacity cuts

$$\sum_{i \in Q} \sum_{j \in \mathcal{V} \setminus Q} y_{ij}^{(s)} + y_{ji}^{(s)} \geq 2 \left\lceil \frac{\sum_{i \in Q} b_i^{(s)}}{C} \right\rceil \quad \forall Q \subset \mathcal{V}, \forall Q \neq \emptyset, \quad (5)$$

see [30, Chapter 3] and [31], whose corresponding cut-form model we refer to as **(DO-1b)**. In this constraint, we ensure that, for every subset of cities, the total number of incoming and outgoing routes should be at least as big as the total demand divided by the truck's capacity times two. For integer candidate solutions, we make a support graph based on the $y_{ij}^{(s)}$ -values and check if the inequality holds for every cycle in this graph. If not, we add the constraint as a cut in the branch-and-cut procedure. In case of a candidate solution with non-integer values, we use the support graph of the x_{ij} -values and add a constraint as a cut if the inequality does not hold for this graph.

Using the introduced notation, we can formulate the extensive form of the ILP for the SCVRP, namely:

$$\begin{aligned}
 \text{(SP-1)} \quad & \min \sum_{s \in \mathcal{S}} p^{(s)} \sum_{i \in \mathcal{V}} \sum_{j \in \mathcal{V}} d_{ij} y_{ij}^{(s)} & (6a) \\
 \text{s.t.} \quad & x \in X^{(\text{DO-1a})} & (6b) \\
 & y^{(s)} \in Y_{(s,x)}^{(\text{DO-1b})} & \forall s \in \mathcal{S} & (6c) \\
 & y_{ij}^{(s)} \leq x_{ij} & \begin{cases} \forall i, j \in \mathcal{V}^*, \\ \forall s \in \mathcal{S} \end{cases} & (6d)
 \end{aligned}$$

The objective of (6) is to minimize the total distance traveled weighted over all scenarios. There should be one tour in the first stage, but additional trips to the depot can be added in the second stage. $y^{(s)}$ in (6c) represents a vector of decision variables, while $y_{ij}^{(s)}$ in (6d) denotes a single decision variable.

4.4. Experimental setup

We now explain how we generate benchmark instances for the SCVRP and how we have set up our experiments. We build upon instances for the Capacitated Vehicle Routing Problem from [32], whose properties are listed in Table 1. We keep the edge costs from these instances and

generate demand per node per scenario using a continuous distribution. Hence, between different scenarios, the graph remains the same, but the demand varies. As proposed in [33], for every node $i \in \mathcal{V}^*$ we generate this stochastic demand $b_i^{(s)}$ with an expected value $\mathbb{E}[b_i^{(s)}]$ equal to the deterministic demand B_i and variance $\text{V}[b_i^{(s)}] = \alpha \cdot \mathbb{E}[b_i^{(s)}]$, where $\alpha > 0$ is a factor taking three values, 0.05, 0.25, and 0.75, to describe situations with low, medium and high variance, respectively. We assume the demand is lognormally distributed, a common choice for demand data [33, Section 6]. To prevent city demand from exceeding the truck’s capacity, we enforce that such a capacity is always at least as high as the largest stochastic demand after having generated all the scenarios. We measure the distance between scenario $i \in \mathcal{S}$ and $j \in \mathcal{S}$ with the L_1 -norm $d(i, j) = \sum_{v \in \mathcal{V}} |b_v^{(i)} - b_v^{(j)}|$, which we choose over the other norms as absolute differences are the most intuitive in the context of demand.

Table 1: Instance overview for CVRP.

Instance name	$ V $	$\min(B_i)$	$\max(B_i)$	Capacity C
eil7	7	1	1	3
eil13	13	1100	1900	6000
eil22	22	100	2500	6000
eil23	23	60	4100	4500
eil30	30	100	3100	4500
eil31	31	1	123	140
eil33	33	40	4000	8000
att48	48	1	1	15
eil51	51	3	41	160
eilA76	76	1	37	140
eilB76	76	1	37	100
eilC76	76	1	37	180
eilD76	76	1	37	220
eilA101	101	1	41	200
eilB101	101	1	41	112
gil262	262	0	100	500

We benchmark our proposed method with solving (6) at once using branch-and-cut. We generate 5 instances for each parameter setting and consider the following parameters and their values: $\alpha \in [0.05, 0.25, 0.75]$ and $S \in [25, 50, 100, 250]$. So in total, we conduct 1920 runs ($16 \cdot 5$ instances, 2 methods, 3 α -values, 4 S -values) with a maximum runtime of 2 hours each. In the numerical experiments, we reduce the scenario tree to 10% of its original size as we found that the method is out-of-sample stable after that point in initial experiments.

We run the experiments on a cluster with 2.4GHz CPU with 8GB RAM, single-threaded, using Gurobi [34] in Python. Some of the 1920 runs did not finish due to out-of-memory issues, especially for the benchmark methods. For a fair comparison, we only include runs that finished for both the benchmark and TULIP, which lead to 1302 finalized

runs. The code and benchmark instances are available on GitHub after publication.

4.5. Results

We compare TULIP with solving the ILP (6) at once using branch-and-cut. Figure 3 shows the time until optimality for both the benchmark and TULIP method on all instances and indicates that our method outperforms the benchmark method as it solves many instances faster. Additionally, TULIP solves more instances optimally than the benchmark method.




Figure 3: Cumulative number of SCVRP instances solved optimally progressively over time by the ILP (6) and TULIP.

Table 2 shows the average run time, the average optimality gap¹, the number of cases in which both lower and upper bounds could be computed (i.e., in which the optimality gaps are smaller than infinity, also referred to as non-infinity optimality gaps), and the number of optimally solved instances. To keep the table compact, we aggregate the instances in three groups based on the graph size: small ($|\mathcal{V}| < 25$), medium ($25 \leq |\mathcal{V}| < 50$), and large ($|\mathcal{V}| \geq 50$). It is clear from this table that TULIP outperforms the benchmark method when the difficulty of the instance increases in terms of the number of scenarios and graph size. In the medium group, especially with many scenarios, TULIP outperforms the benchmark method as it solves more instances optimally with faster run times. However, this difference is less evident in the other two groups. In the small group with few scenarios, TULIP consumes relatively a lot of overhead time, making TULIP redundant. In the large group and especially with many scenarios, TULIP still finds non-infinity gaps whereas the benchmark method does not find any solution. However, this performance difference decreases as

¹We take the instances that yield non-infinity gaps for the benchmark method and compute the average gaps over them for TULIP.

the number of scenarios increases, indicating that TULIP has reached its limit.

As the α -values indicate the variability in the stochastic demand of the SCVRP-instances, our hypothesis was that TULIP would thrive especially under low variability as many scenarios are similar to each other. Table 3 shows the run time in seconds, and the optimality gap in percentages per α -value. From this table, we conclude that the benefits of TULIP exists for all considered α -values.

When comparing TULIP with the benchmark method in terms of the number of capacity and subtour elimination cuts that are added before finding the optimal solution, we find that TULIP yields considerably more cuts, even when we discard the cuts that are not tight after the first step. When focusing on the ratio of tight cuts over the total number of cuts generated in the first step of TULIP, we observe that this ratio decreases as the instance difficulty grows (0.61, 0.55, and 0.31 for small, medium, and large, respectively). For difficult instances, relatively many cuts can be discarded as TULIP only uses tight cuts in the warm start, which saves time in solving the ILP of the original problem. This is a reason why TULIP outperforms the benchmark method, especially for difficult instances.

5. Case study II: Two-Stage Stochastic Steiner Forest Problem

We introduce the Two-Stage Stochastic Steiner Forest Problem (2S-SSFP), compare TULIP performance on this problem with two benchmark methods, and finally address the robustness of TULIP.

The 2S-SSFP has many applications in different industries such as telecommunication [1, 35] and maritime design [18] and is a generalization of the Steiner Forest Problem (SFP) [36]. The SFP, which is itself a generalization of the well-known Steiner Tree Problem (STP) [37], seeks to find a minimum-cost subgraph spanning one or more sets of vertices, which we refer to as terminals.

In the Stochastic Steiner Tree Problem (SSTP) [35] and Stochastic Steiner Forest Problem (SSFP) [38], the edge costs and terminals are affected by uncertainty. The decision maker can connect vertices using edges in the first and second stages. In the first stage, it is unknown which set of terminals must be connected in the second stage, as these are revealed only in the second stage.

In the 2000s, researchers focused on approximation algorithms for the SSTP [39, 40, 41, 42, 43, 44, 38, 45] whereas exact methods have been studied afterwards. In [46], the authors describe an exact model that uses a two-stage branch-and-cut algorithm based on Benders' decomposition. Different ILP models for the SSTP are studied in [47], whereas [1] describe a two-stage branch-and-cut algorithm based on a decomposed model. In [35], the authors suggest a new decomposition model, which is the current state-of-the-art for solving SSTP to optimality. The authors show that their methods considerably outperform these benchmarks using three procedures for computing lower bounds:

dual ascent, Lagrangian relaxation, and Benders' decomposition. In [48], different ILP models for the SFP are studied, which [18] uses to describe and model the 2S-SSFP.

5.1. Problem description

The 2S-SSFP introduced in [18], which we will use to test our method, differs from other variants of the Stochastic Steiner Forest problem (SSFP) by considering: 1) sets of terminals that must be connected already in the first-stage solution; and 2) multiple types of connections per edge that can, e.g., correspond to different pipes or cables in ship design or telecommunications, respectively.

The 2S-SSFP considers an undirected graph $G = (\mathcal{V}, \mathcal{E})$ and a set of connection types \mathcal{M} . For each scenario $s \in \mathcal{S} \cup \{0\}$, where $s = 0$ indicates the first stage, a subset of edges $E^{(s)} \subseteq \mathcal{E}$ and connection types $M^{(s)} \subseteq \mathcal{M}$ can be used. The first-stage costs $c_{me}^{(0)} \geq 0$ are defined for each edge $e \in \mathcal{E}$ and connection type $m \in \mathcal{M}$. With the connection types $M^{(0)}$, we can connect the first-stage terminals groups $\mathcal{T}^{(0)} = (T_k^{(0)})_{k \in \mathcal{K}^{(0)}}$, $T_k^{(0)} \subseteq \mathcal{V}$, $\mathcal{K}^{(0)} = \{1, \dots, K^{(0)}\}$, $K^{(0)} \in \mathbb{N}$. Similarly, second-stage costs $c_{me}^{(s)} \geq 0$, $e \in \mathcal{E}$, $m \in \mathcal{M}$, and second-stage terminal groups $\mathcal{T}^{(s)} = (T_k^{(s)})_{k \in \mathcal{K}^{(s)}}$, $T_k^{(s)} \subseteq \mathcal{V}$, $\mathcal{K}^{(s)} = \{1, \dots, K^{(s)}\}$, $K^{(s)} \in \mathbb{N}$, are considered for each scenario $s \in \mathcal{S}$ which occurs with probability $p^{(s)} \in (0, 1]$, $\sum_{s \in \mathcal{S}} p^{(s)} = 1$.

A solution to the 2S-SSFP consists of a set of first-stage connection type-edge pairs $\bar{E}^{(0)} \times \bar{M}^{(0)} \subseteq \mathcal{E} \times \mathcal{M}$ and second-stage connection type-edge pairs $\bar{E}^{(s)} \times \bar{M}^{(s)} \subseteq \mathcal{E} \times \mathcal{M}$ for each scenario $s \in \mathcal{S}$ such that the subgraph(s) induced by $(\bar{E}^{(0)} \times \bar{M}^{(0)}) \cup (\bar{E}^{(s)} \times \bar{M}^{(s)})$, connects $\mathcal{T}^{(0)}$ and $\mathcal{T}^{(s)}$ and the expected costs

$$\sum_{(e,m) \in \bar{E}^{(0)} \times \bar{M}^{(0)}} c_{me}^{(0)} + \sum_{s \in \mathcal{S}} p^{(s)} \sum_{(e,m) \in \bar{E}^{(s)} \times \bar{M}^{(s)}} c_{me}^{(s)}$$

are minimized.

In the context of ship design [18], vertices \mathcal{V} correspond to ship rooms containing engines or fuel tanks (subsets of which need to be connected by appropriate pipes), and scenarios correspond to different fuel types (each of which require different pipe types [49]).

Now, we introduce some notation that will become useful when explaining the 2S-SSFP ILP model. We denote the set of arcs of the bi-direction of G by $\mathcal{A} := \{(u, v) : u \in \mathcal{V}, v \in \mathcal{V}, \{u, v\} \in \mathcal{E}\}$. For a given terminal set $T_k^{(i)}$ with $i = \{0\} \cup \mathcal{S}$ and $k \in \mathcal{K}^{(i)}$, a vertex $v \in \mathcal{V} \setminus T_k^{(i)}$ is called a Steiner node. For $s \in \{0\} \cup \mathcal{S}$, we define the set of Steiner nodes by: $\mathcal{Q}^{(s)} = \mathcal{V} \setminus \mathcal{T}^{(s)}$. Next, we introduce the set $\mathcal{R}^{(s)}$, which denotes the root vertices; $\mathcal{R}^{(s)} = \{r^1, \dots, r^K\}$, where $r^k \in T_k$ for terminal group $k \in \mathcal{K}$. Note that the root vertex is chosen arbitrarily for each terminal group. $\tau(t)$ corresponds to the index of the terminal group to which terminal t belongs. Finally, the set $\mathcal{T}_r^{k \dots K}$ is the set of some terminal sets: $\mathcal{T}^{k \dots K} = (T^k)_{k \in \{k, \dots, K\}}$. Then,

Table 2: Comparison TULIP and ILP (6) for SCVRP between different instance groups, based on the number of vertices and the number of scenarios. We report the average run time in seconds, the average optimality gap in percentages, the number of non-infinity optimality gaps, and the number of optimally solved instances. Per column and in every row, we represent in bold whether TULIP or ILP (6) has the best value.

Group	Scenarios	Instances	Time (sec)		Gap (%)		Non-infinity gaps		Optimal solutions	
			ILP (6)	TULIP	ILP (6)	TULIP	ILP (6)	TULIP	ILP (6)	TULIP
Small	25	60	6	7	0.00%	0.00%	60	60	60	60
	50	60	15	16	0.00%	0.00%	60	60	60	60
	100	60	117	60	0.00%	0.00%	60	60	60	60
	250	60	588	572	0.00%	0.00%	60	60	45	60
Medium	25	56	1470	1437	6.14%	4.91%	54	53	45	45
	50	58	2059	1707	0.00%	0.00%	45	45	45	45
	100	53	2961	1760	0.05%	0.15%	45	45	38	43
	250	49	4715	4110	4.60%	6.67%	30	45	25	30
Large	25	83	3911	2748	0.04%	0.05%	77	81	76	56
	50	65	7200+	5845	2.90%	0.47%	19	52	0	37
	100	42	7200+	7200+	1.73%	2.32%	10	11	0	0
	250	5	7200+	7200+	NAN	NAN	0	4	0	0


Table 3: Comparison of TULIP and ILP (6) for different α -values.

α	Time (sec)		Gap (%)	
	ILP (6)	TULIP	ILP (6)	TULIP
0.05	2527	2060	0.83%	0.48%
0.25	2860	2322	1.36%	1.34%
0.75	2740	2342	0.96%	1.06%


$\mathcal{T}_r^{k\dots K}$ represents the set of some terminal sets without the corresponding root vertex r^k : $\mathcal{T}_r^{k\dots K} = \mathcal{T}^{k\dots K} \setminus \{r^k\}$. For $W \subset \mathcal{V}$, let $\delta^+(W) := \{(u, v) \in \mathcal{A} : u \in W, v \in \mathcal{V} \setminus W\}$ be the outgoing arc set. For some $s \in \{0\} \cup \mathcal{S}$, $k \in \mathcal{K}^{(s)}$ and $l \in \mathcal{K}^{(s)}$, we say that a cut-set $\mathcal{H} \subseteq \mathcal{V}$ is relevant for r^k and $T_l^{(s)}$ if $r^k \in \mathcal{H}$ and some terminal $t \in T_l^{(s)}$ is not in \mathcal{H} . The set of all cut-sets that are relevant for r^k and $T_l^{(s)}$ is written by $H_{kl}^{(s)}$.

5.2. Illustrative example


To explain the 2S-SSFP, we include a small illustrative example. We consider a graph whose first-stage edge costs for connection type 1 are shown in Figure 4a and second-stage costs are twice as high as the first-stage costs. In the first stage, our terminal set consists of vertices A and D . This set changes with probability $p^{(2)}$ into $\{A, B\}$ in the second stage and remains the same with probability $p^{(1)}$, where $p^{(1)} + p^{(2)} = 1$. Hence, we consider two scenarios whose indices are captured in $\mathcal{S} = \{1, 2\}$. We consider two connection types $\mathcal{M} = \{1, 2\}$, of which we can use all for the first scenario ($M^{(0)} = M^{(1)} = \{1, 2\}$) and one for the second scenario ($M^{(2)} = \{2\}$). For simplicity, we assume that the second connection type is twice as expensive as the first connection type. All edges are admissible and can be used to connect vertices. As shown in Figure 4, the deterministic solution suggests installing connection type 1 on the edge connecting A and D , ensuring only a first-stage connection, whereas the stochastic solution installs



(a) Graph with the first stage edge costs.



(b) DO solution.



(c) SP solution for $p^{(2)} = 0.4$.

Figure 4: Small example that illustrates the 2S-SSFP, based on [50, Figure 2]. Red and blue denote connection types 1 and 2, respectively. Solid and dotted lines correspond to the first and second stages, respectively.

two connections in the first stage and one in the second, which yields the lowest expected costs in case $p^{(2)} = 0.4$.

5.3. Model formulation

We compare our method with two similar benchmark methods, a flow- and cut-based ILP. To that end, we first describe the DO cut-based ILP for the 2S-SSFP, which is similar to the flow-based ILP as proposed by [18] and provided in Appendix A. After that, we provide the SP formulation.

The ILP for the 2S-SSFP contains four types of decision variables, denoted by a $(s) \in \{0\} \cup \mathcal{S}$ superscript, which clarifies if it entails a first- ($s = 0$) or second-stage ($s \geq 1$) decision variable. The binary decision variable $x_{muv}^{(s)}$ equals 1 if we install a connection type $m \in \mathcal{M}$ on edge $(u, v) \in \mathcal{E}$. Binary decision variable $z_{kl}^{(s)}$ is 1 when the root of the terminal group k sends flow to all terminals of

the terminal group l , and 0 else. When $z_{kl}^{(s)} = 1$, binary decision variable $y_{kmuv}^{(s)}$ equals 1 when flow from the root of terminal group k is sent over arc (u, v) through connection type m , and 0 else. Lastly, binary decision variable $y_{muv}^{(s)}$ equals 1 when connection type m at arc (u, v) is used

to send flow over by any of the created arborescences (a directed tree), and 0 else. As we describe the DO ILP for the 2S-SSFP, we only consider the first stage and therefore only use (0) superscripts for the decision variables and sets in the model's description.

(DO-2a)

$$\begin{aligned}
\min \quad & \sum_{((u,v),m) \in (\mathcal{E} \times \mathcal{M})} \left(x_{muv}^{(0)} \cdot c_{muv}^{(0)} \right) & (7a) \\
\text{s.t.} \quad & \sum_{m \in \mathcal{M}} \sum_{(u,v) \in \delta^+(\mathcal{H})} y_{kmuv}^{(0)} \geq z_{kl}^{(0)} & \left\{ \begin{array}{l} \forall k \in \mathcal{K}^{(0)}, \forall l \in \{k, \dots, K^{(0)}\}, \\ \forall \mathcal{H} \in H_{kl}^{(0)} \end{array} \right. & (7b) \\
& \sum_{k \in \mathcal{K}} y_{kmuv}^{(0)} \leq y_{muv}^{(0)} & \forall m \in M^{(0)}, \forall (u, v) \in A^{(0)} & (7c) \\
& y_{muv}^{(0)} + y_{mvu}^{(0)} \leq x_{muv}^{(0)} & \forall m \in M^{(0)}, \forall (u, v) \in E^{(0)} & (7d) \\
& \sum_{l=1}^k z_{lk}^{(0)} = 1 & \forall k \in \mathcal{K}^{(0)} & (7e) \\
& z_{kk}^{(0)} \geq z_{kl}^{(0)} & \left\{ \begin{array}{l} \forall k \in \mathcal{K}^{(0)} \setminus \{1, K^{(0)}\} \\ \forall l \in \mathcal{K}^{(0)} \text{ if } l \geq k + 1 \end{array} \right. & (7f) \\
& \sum_{m \in M^{(0)}} \sum_{u:(u,v) \in A^{(0)}} y_{muv}^{(0)} \leq 1 & \forall v \in \mathcal{V} & (7g) \\
& \sum_{m \in M^{(0)}} \sum_{u:(u,t) \in A^{(0)}} y_{kmuv}^{(0)} = 0 & \forall k \in \mathcal{K}^{(0)} \setminus \{1\}, \forall t \in \mathcal{T}^{1 \dots k-1} & (7h) \\
& \sum_{m \in M^{(0)}} \sum_{u:(u,v) \in A^{(0)}} y_{muv}^{(0)} \leq \sum_{m \in M^{(0)}} \sum_{u:(v,u) \in A^{(0)}} y_{muv}^{(0)} & \forall v \in \mathcal{Q}^{(0)} & (7i) \\
& \sum_{m \in M^{(0)}} \sum_{u:(u,v) \in A^{(0)}} y_{kmuv}^{(0)} \leq \sum_{m \in M^{(0)}} \sum_{u:(v,u) \in A^{(0)}} y_{kmuv}^{(0)} & \forall k \in \mathcal{K}^{(0)}, \forall v \in \mathcal{V} \setminus \mathcal{T}_r^{k \dots K^{(0)}} & (7j) \\
& \sum_{u:(u,r^l) \in A^{(0)}} y_{kmur^l}^{(0)} \leq z_{kl}^{(0)} & \left\{ \begin{array}{l} \forall k \in \mathcal{K}^{(0)} \setminus K^{(0)} \\ \forall l \in \mathcal{K}^{(0)} \text{ if } l \geq k + 1 \\ \forall m \in M^{(0)} \end{array} \right. & (7k) \\
& x_{muv}^{(0)} \in \mathbb{B} & \forall m \in \mathcal{M}, \forall (u, v) \in \mathcal{E} & (7l) \\
& y_{muv}^{(0)} \in \mathbb{B} & \forall m \in M^{(0)}, \forall (u, v) \in A^{(0)} & (7m) \\
& y_{kmuv}^{(0)} \in \mathbb{B} & \left\{ \begin{array}{l} \forall k \in \mathcal{K}^{(0)}, \forall m \in M^{(0)} \\ \forall (u, v) \in A^{(0)} \end{array} \right. & (7n) \\
& z_{kl}^{(0)} \in \mathbb{B} & \forall k \in \mathcal{K}^{(0)}, \forall l \in \{k \dots K^{(0)}\} & (7o)
\end{aligned}$$

With constraint (7b), we ensure connectivity between the terminals. For example, if $z_{kl}^{(0)} = 1$, the model must connect the terminals from $T_l^{(0)}$ to root r^k . If $z_{kl}^{(0)} = 0$, the constraint is automatically satisfied. In (7c), we ensure that each arc is assigned to only one arborescence. If multiple arborescences share the same arc, they are forced to merge into a single arborescence. Constraint (7d) limits flow direction to a single direction for each edge. In (7e), we enforce that each terminal group has exactly one root, whereas (7f) requires a single root per arborescence.

Constraints (7g)-(7k) are not strictly required for **(DO-2a)** to generate feasible solutions. Instead, they are introduced to improve the model's LP-relaxation, as described in [48]. In (7g), we require that each vertex receives flow through at most one connection. Since z_{kl} defines that the root r^k is responsible for terminal groups $l \geq k$, constraint (7h) prevents any connection between root r^k and terminals from groups $\mathcal{T}^{1 \dots k-1}$. Flow-balance constraints are given in (7i) and (7j), similar to those in [35, Section 2.2] for the SSTP. These constraints enforce that the

in-degree of a Steiner vertex cannot exceed its out-degree: (7i) applies this to the complete solution, while (7j) focuses on each terminal group. We ensure that the arborescence rooted at r^k can only use root r^l if $z_{kl} = 1$ in (7k). Lastly, integrality constraints are imposed in (7l)-(7o).

Typically, flow-based formulations are computationally slower than cut-based formulations. We introduce the flow-based equivalent of (7), which we refer to as **(DO-2b)**, in Appendix A.

Using the introduced notation, we can formulate the ILP for the 2S-SSFP, which is given below:

$$\begin{aligned}
\text{(SP-2)} \quad & \min \sum_{((u,v),m) \in (\mathcal{E} \times \mathcal{M})} \left(x_{muv}^{(0)} \cdot c_{muv}^{(0)} + \right. \\
& \left. \sum_{s \in \mathcal{S}} p^{(s)} \cdot (x_{muv}^{(s)} - x_{muv}^{(0)}) \cdot c_{muv}^{(s)} \right) \quad (8a) \\
\text{s.t.} \quad & x^{(0)}, f^{(0)}, y^{(0)}, z^{(0)} \in X^{(\text{DO-2a})} \quad (8b) \\
& x^{(s)}, f^{(s)}, y^{(s)}, z^{(s)} \in Y_{(s, x^{(0)})}^{(\text{DO-2a})} \quad \forall s \in \mathcal{S} \quad (8c) \\
& x_{muv}^{(s)} \geq x_{muv}^{(0)} \quad \left\{ \begin{array}{l} \forall s \in \mathcal{S}, \forall m \in \mathcal{M}, \\ \forall (u, v) \in \mathcal{E} \end{array} \right. \quad (8d)
\end{aligned}$$

We will elaborate on the types of cuts we make in our ILP (7), which is mainly based on [48, Section 4.1]. For each $s \in \{0\} \cup \mathcal{S}$, $k \in \mathcal{K}^{(s)}$ and $l \in \mathcal{K}^{(s)}$ with $l \geq k$, we compute a maximum r^k - t -flow in the support graph of y_k for each $t \in T^l$. If the flow value is strictly less than z_{kl} , the corresponding minimum r^k - t -cut induces an inequality of type (7b). To prevent generating equivalent cuts for different root-terminal pairs, we do not include cuts added by the previous root-terminal pair(s). Additionally, we use creep flows and back cuts, which increase the likelihood of generating tight cuts.

5.4. Experimental setup

We now explain how we generate benchmark instances for the 2S-SSFP and how we set up our experiments. We build upon instances used in [35], which are taken from [51], a benchmark set provided during the 11th DIMACS challenge on Steiner trees. These instances - called K100, P100, LIN01-10, and WRP - have been generated from STP instances in the SteinLib dataset. Each dataset contains up to 1000 scenarios; an instance with, e.g., 50 scenarios is a subset of the original instance with 1000 scenarios. In total, we have $S \in \{5, 10, 20, 50, 75, 100, 150, 200, 250, 300, 400, 500, 750, 1000\}$. We limited the experiments to instances of up to 200 scenarios as initial experiments showed optimality gaps of infinity for many instances with more scenarios.

We adjust these SSTP instances, see Table 4 for their properties, to 2S-SSFP as follows. First, we require the model to solve scenario $s = 1$ in the first stage as well. Second, we add two more terminal sets with five terminals each to each scenario. These numbers are chosen to

mimic a realistic ship pipe routing setting in which relatively few rooms need to be connected to each other. In [18], the authors studied the algorithmic behavior for different numbers of terminal sets and terminals per scenario, hence we do not vary the number of terminals and terminal sets in this study. Third, we add a connection type set $\mathcal{M} := \{1, 2\}$ to the dataset where the costs of $m = 2$ are twice as high as $m = 1$. Again, this parameter was chosen to mimic a realistic ship pipe routing setting where, for example, double-walled pipes are more expensive than single-walled pipes. With equal probability, we assign either $M^{(s)} \in \{\{1\}, \{2\}, \{1, 2\}\}$ to a scenario $s \in \mathcal{S}$. We set the second-stage costs twice as high as the first-stage costs to incentivize installing connections in the first stage. For simplicity, we assume that all edges can be used, i.e., $E^{(s)} = \mathcal{E} \quad \forall s \in \{0\} \cup \mathcal{S}$. For more detail, we refer to [18].

To quantify the distance between two scenarios $i \in \mathcal{S}$ and $j \in \mathcal{S}$, we introduce three distance metrics based on edge costs, terminal positions, and connection types. We start with the distance metric between scenarios $i \in \mathcal{S}$ and $j \in \mathcal{S}$ based on edge costs:

$$L_{1(ij)} = \sqrt{\sum_{(u,v) \in \mathcal{E}} \sum_{m \in \mathcal{M}} (c_{muv}^{(i)} - c_{muv}^{(j)})^2}. \quad (9)$$

Here, we take the L_2 -norm of the edge costs between two scenarios because it is an intuitive metric to quantify the distance between two points in a two-dimensional plane.

Next, we propose a distance metric based on the terminal positions:

$$L_{2(ij)} = \sum_{k_1 \in \mathcal{K}^{(i)}} \min_{k_2 \in \mathcal{K}^{(j)}} \left| \left(T_{k_1}^{(i)} \setminus T_{k_2}^{(j)} \right) \cup \left(T_{k_2}^{(j)} \setminus T_{k_1}^{(i)} \right) \right|. \quad (10)$$

For each terminal group in scenario i , we find the most similar terminal group in scenario j , and compute the difference between those two. We repeat this process for all terminal groups in scenario i .

Finally, we introduce a distance metric based on the connection types:

$$L_{3(ij)} = \left| \left(M^{(i)} \setminus M^{(j)} \right) \cup \left(M^{(j)} \setminus M^{(i)} \right) \right|. \quad (11)$$

Here, we check the difference between the used connections in scenario i and j . Let $d(i, j) = \beta_1 L_{1(ij)} + \beta_2 L_{2(ij)} + \beta_3 L_{3(ij)}$, where $d(i, j)$ represents the distance between scenario i and scenario j , $\beta_i \in (0, 1)$ is a weight with $i \in \{1, 2, 3\}$ and $\sum_{i=1}^3 \beta_i = 1$. Unless explicitly stated otherwise, we set $\beta_1 = 1$ in our experiments.

We use 40 instances from [51] (K100, P100, LIN01-10, and WRP), adjusted to the 2S-SSFP, with eight different scenarios each ($S \in \{5, 10, 20, 50, 75, 100, 150, 200\}$). We do not include more scenarios as it typically leads to out-of-memory issues in preliminary experiments. For three methods, this yields 960 runs. To counter out-of-memory

Table 4: Properties of the SSTP benchmark instances, similar to [35, Table 1].

Dataset	Instances [#]	\mathcal{V}			\mathcal{E}			\mathcal{S}	
		min	avg	max	min	avg	max	min	max
K100	154	22	31	45	64	115	191	5	1000
P100	70	66	77	91	163	194	237	5	1000
LIN01-10	140	53	190	321	80	318	540	5	1000
WRP	196	10	194	311	149	363	613	5	1000

issues and improve the run time for larger instances (those 100 scenarios or more), we used “nodefiles” in Gurobi to store branch-and-bound nodes on disk instead of in memory, which is particularly useful for handling large MILP problems. Yet, all the other code and settings remain the same.

We run the experiments, single-threaded, on the same cluster, yet with 2GB RAM per run. Due to out-of-memory issues, some of the 960 runs did not finish. For a fair comparison, we only include the runs that finish for all three methods (flow-based ILP (A.1), cut-based ILP (7), and TULIP) up to 100 scenarios. As ILP (A.1) runs out-of-memory for each instance after 100 scenarios, we only include runs that finished for both ILP (7) and TULIP. This leads to 314 runs in total on which we conduct an analysis. The code and benchmark instances are available on GitHub after publication.

5.5. Results

Figure 5 shows the cumulative number of instances solved optimally over time for two benchmark methods, the flow-based ILP (A.1) and the cut-based ILP (7), and TULIP. The figure indicates that our method outperforms the benchmark methods as it solves more instances optimally in the same amount of time. As expected, ILP (A.1) solves the least instances optimally, especially when the number of scenarios grows. The ILP (7) and TULIP have a comparable performance up to 1000 seconds, after which TULIP outperforms this benchmark method.

Table 5 shows more elaborate data from the same experiment. To keep the table concise, we aggregate based on the number of scenarios: an instance belongs to the small group if it contains fewer than 100 scenarios, and large otherwise. We see no significant advantage of TULIP over the benchmark methods for the first group. However, for the second group, we find a considerable gain in run times for TULIP in instances from K100, LIN01-10, and P100. Additionally, we see that TULIP finds more optimal solutions than the benchmark methods, indicating that this method becomes more suitable when the instance difficulty in terms of scenarios increases.

5.5.1. Robustness of TULIP

Next, we want to study how small the subset of scenarios can be while staying close to the solution of the original problem. To that end, we pick an instance with 50 scenarios from each dataset in Table 4 and select the




Figure 5: Cumulative number of 2S-SSFP instances solved optimally progressively over time by the ILPs (7) and (A.1) and TULIP.

$\mathcal{S}' \in \{1, 3, 5, 8, 10, 20, 35, 50\}$ scenarios that yield the best representation of the original scenario tree using fast forward scenario selection. We fix the resulting first-stage decision in the original problem, solve it, and store the objective value. We visualize this by plotting \mathcal{S}' on the x-axis and the objective value on the y-axis. We also apply the same procedure for random scenario selection - executed 25 times to account for randomness - to study the impact of the chosen scenario selection method. We visualize the result with the gray shaded area denoting the minimum and maximum objective values obtained through random sampling in Figure 6. The objective values converge considerably fast, meaning that a small subset of scenarios is enough to represent the whole scenario tree. The blue and orange lines converge approximately simultaneously, which indicates that fast forward selection and random sampling, on average, yield comparable results. However, fast forward selection performs considerably better than the worst sample.

In the next experiment, we measure the impact of different distance metrics on the performance of fast forward selection. For illustration purposes, we discuss the result of one instance with 50 scenarios but remark that we see the same pattern on other instances as well. We vary the weights β_i such that either distance metric $L_{1(ij)}$, $L_{2(ij)}$, or $L_{3(ij)}$ is used, and observe that the set of selected sce-

Table 5: Comparison TULIP and benchmark methods for 2S-SSFP between different instance groups.

Scenarios	Instance group	Instances	Time (sec)			Gap (%)			Non-infinity gaps			Optimal Solutions		
			ILP (A.1)	ILP (7)	TULIP	ILP (A.1)	ILP (7)	TULIP	ILP (A.1)	ILP (7)	TULIP	ILP (A.1)	ILP (7)	TULIP
Small	K100	42	68	122	119	0,00%	0,00%	0,00%	42	42	42	33	30	33
	LIN01-10	19	2603	962	885	0,00%	0,01%	0,00%	14	19	19	12	17	16
	P100	9	170	136	184	0,00%	0,00%	0,00%	9	9	9	3	6	6
	WRP	16	2390	1151	1184	0,00%	0,00%	0,00%	12	16	16	3	7	7
Large	K100	19		3019	2078	0,00%	0,00%			16	18		6	9
	LIN01-10	5		7007	3657		0,32%	0,00%		3	5		1	4
	P100	3		7111	6183		0,01%	0,00%		1	1		0	1
	WRP	1		2215	2216		0,00%	0,00%		1	1		0	0

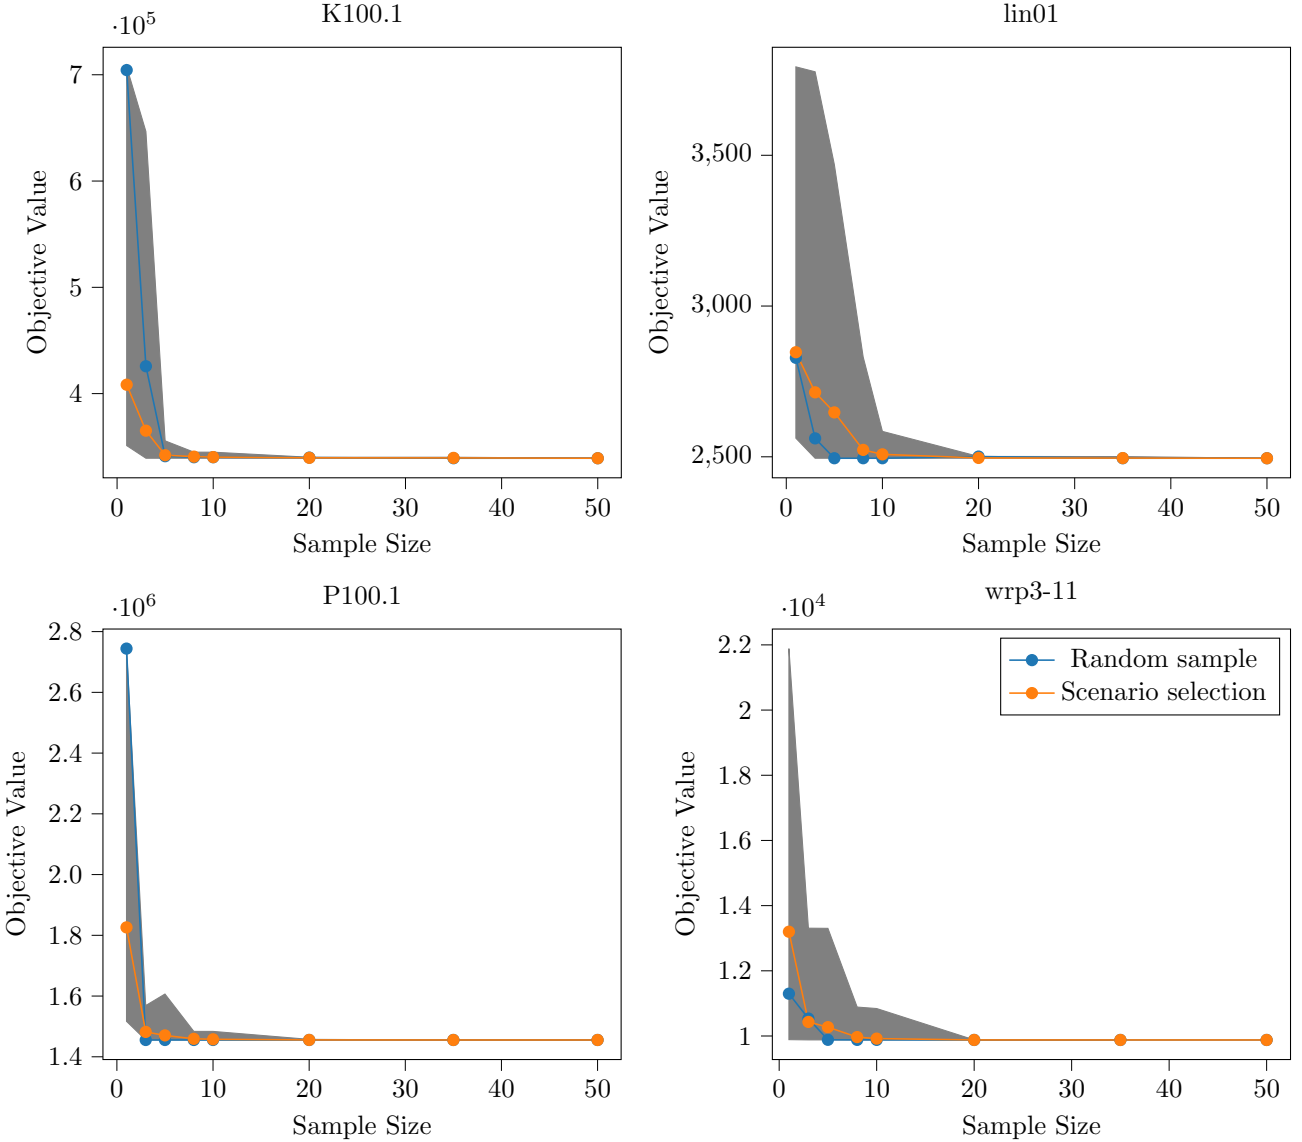


Figure 6: Comparison of random and fast forward scenario selection. Gray shaded area denotes the minimum and maximum objective value found by random sampling.

narios differs considerably between the different distance metrics. When fixing the first-stage solution of the reduced problem in the original problem, we see that the distance metric does not significantly affect the objective value (2495, 2501, 2501, respectively.) Hence, fast forward selection seems to be robust against the different distance metrics for the 2S-SSFP.

6. Conclusion

To solve large-scale instances of (mixed-)integer two-stage stochastic programming models with an exponential number of constraints to optimality, we propose TULIP, a novel combination of methods consisting of two steps. In the first step, we drastically reduce the scenario set, rep-

representing the original scenario tree as well as possible, and solve the root node of the corresponding ILP. In the second step, the tight cuts from the first step are re-used as constraints when solving the original problem with an ILP through branch-and-cut. We show that TULIP is generically applicable to different kinds of problems by testing it on two benchmark problems: the Stochastic Capacitated Vehicle Routing Problem (SCVRP) and the Two-Stage Stochastic Steiner Forest Problem (2S-SSFP).

The results of our experiments clearly indicate that TULIP outperforms the benchmark methods when the instance difficulty in terms of scenarios and graph size increases. In the first case study with the SCVRP, TULIP solves more instances optimally and faster than the benchmark method. This advantage becomes more pronounced as the difficulty of problem instances increases, particularly with the growth in the number of vertices or scenarios. However, despite expecting TULIP to perform better under low variability in stochastic demand, our experiments indicate that TULIP performs well under all levels of variability in stochastic demand.

A key reason for the efficiency of TULIP lies in its handling of cuts. Our method saves considerable time in the second step by only adding cuts that are tight in the first step as constraints to the ILP in the second step. This efficiency becomes more crucial as the problem difficulty increases, explaining why TULIP scales better than the benchmark method.

In the second case study with the 2S-SSFP, TULIP again outperforms the benchmark methods confirming that it is efficient in handling larger and more complex instances. The fast-forward scenario selection approach also proves to be efficient, as a small subset of scenarios was sufficient to represent the entire scenario tree. Additionally, fast-forward selection demonstrated robustness against different distance metrics, further supporting its usefulness.

Besides these successes, the TULIP method faces some limitations. Attempts to speed up the method using multiple warm starts yield no significant performance gains. However, it could be interesting to explore variations of this idea further for different problems. Additionally, we still solve both the SCVRP and the 2S-SSFP without a decomposition method as the L-shaped method [10], so there is a limit on the instance difficulty that TULIP can solve. Hence, studying decomposition methods in combination with TULIP could be an interesting topic for future research.

Future research could also test TULIP on other (mixed-) integer problems from the literature to gain more insight into its applicability. As TULIP is meant to be a generic framework, we use a distribution-based scenario generation method in its first step. However, it could be valuable to test the influence of problem-based scenario generation [16] in TULIP for specific problems. Finally, extending TULIP to multi-stage SP could also be an interesting topic for future research.

Acknowledgements

We thank Ruurd Buijs for the insightful discussions during this research. Additionally, we thank SURF (www.surf.nl) for the support in using the National Supercomputer Snelius. This publication is part of the project READINESS with project number TWM.BL.019.002 of the research program *Topsector Water & Maritime: the Blue route* which is partly financed by the Dutch Research Council (NWO).

References

- [1] Ivana Ljubić, Petra Mutzel, and Bernd Zey. Stochastic survivable network design problems: Theory and practice. *European Journal of Operational Research*, 256(2):333–348, January 2017.
- [2] B. Zahiri, S. Ali Torabi, M. Mohammadi, and M. Aghabegloo. A multi-stage stochastic programming approach for blood supply chain planning. *Computers & Industrial Engineering*, 122:1–14, August 2018.
- [3] Stein W. Wallace and Stein-Erik Fleten. Stochastic Programming Models in Energy. In *Handbooks in Operations Research and Management Science*, volume 10, pages 637–677. Elsevier, 2003.
- [4] J. L. Midler and R. D. Wollmer. Stochastic programming models for scheduling airlift operations. *Naval Research Logistics Quarterly*, 16(3):315–330, September 1969.
- [5] Hajnalka Vaagen, Stein W. Wallace, and Michal Kaut. Modelling consumer-directed substitution. *International Journal of Production Economics*, 134(2):388–397, December 2011.
- [6] John R. Birge and François Louveaux. *Introduction to Stochastic Programming*. Springer Series in Operations Research and Financial Engineering. Springer New York, New York, NY, 2011.
- [7] Willem K. Klein Haneveld, Maarten H. van der Vlerk, and Ward Romeijnders. *Stochastic programming: modeling decision problems under uncertainty*. Graduate texts in operations research. Springer, 2020.
- [8] Sujin Kim, Raghu Pasupathy, and Shane G. Henderson. A Guide to Sample Average Approximation. In Michael C Fu, editor, *Handbook of Simulation Optimization*, volume 216, pages 207–243. Springer New York, New York, NY, 2015. Series Title: International Series in Operations Research & Management Science.
- [9] Alan J. King and Stein W. Wallace. *Modeling with Stochastic Programming*. Springer Series in Operations Research and Financial Engineering. Springer International Publishing, Cham, 2024.
- [10] Gilbert Laporte and François Louveaux. The integer L-shaped method for stochastic integer programs with complete recourse. *Operations Research Letters*, 13(3):133–142, April 1993.
- [11] Ragheb Rahmaniani, Teodor Gabriel Crainic, Michel Gendreau, and Walter Rei. The Benders decomposition algorithm: A literature review. *European Journal of Operational Research*, 259(3):801–817, June 2017.
- [12] J. N. Hooker. Integer Programming: Lagrangian Relaxation. In Panos M. Pardalos and Oleg A. Prokopyev, editors, *Encyclopedia of Optimization*, pages 1–7. Springer International Publishing, Cham, 2024.
- [13] Mikhail A. Bragin. Survey on Lagrangian relaxation for MILP: importance, challenges, historical review, recent advancements, and opportunities. *Annals of Operations Research*, July 2023.
- [14] Richard T. Wong. A dual ascent approach for steiner tree problems on a directed graph. *Mathematical Programming*, 28(3):271–287, October 1984.
- [15] Holger Heitsch and Werner Römisich. Scenario Reduction Algorithms in Stochastic Programming. *Computational Optimization and Applications*, 24(2/3):187–206, 2003.

- [16] Xiaochen Chou and Enza Messina. Problem-Driven Scenario Generation for Stochastic Programming Problems: A Survey. *Algorithms*, 16(10):479, October 2023.
- [17] Julia L. Higle. Stochastic Programming: Optimization When Uncertainty Matters. In Harvey J. Greenberg and J. Cole Smith, editors, *Emerging Theory, Methods, and Applications*, pages 30–53. INFORMS, September 2005.
- [18] B. T. Markhorst, Joost Berkhout, Alessandro Zocca, J. F. J. Pruyn, and R. D. van der Mei. Future-proof ship pipe routing: navigating the energy transition. 2023. Publisher: arXiv Version Number: 2 DOI: 10.48550/ARXIV.2312.09088.
- [19] Marco Colombo, Jacek Gondzio, and Andreas Grothey. A warm-start approach for large-scale stochastic linear programs. *Mathematical Programming*, 127(2):371–397, April 2011.
- [20] Gustavo Angulo, Shabbir Ahmed, and Santanu S. Dey. Improving the Integer L-Shaped Method. *INFORMS Journal on Computing*, 28(3):483–499, July 2016.
- [21] Ramkumar Karupiah, Mariano Martín, and Ignacio E. Grossmann. A simple heuristic for reducing the number of scenarios in two-stage stochastic programming. *Computers & Chemical Engineering*, 34(8):1246–1255, August 2010.
- [22] Benjamin S. Narum, Jamie Fairbrother, and Stein W. Wallace. Problem-based scenario generation by decomposing output distributions. *European Journal of Operational Research*, pages 154–166, April 2024.
- [23] V. M. Zolotarev. Probability Metrics. *Theory of Probability & Its Applications*, 28(2):278–302, January 1984.
- [24] Marco Colombo and Andreas Grothey. A decomposition-based crash-start for stochastic programming. *Computational Optimization and Applications*, 55(2):311–340, June 2013.
- [25] Frederick S. Hillier and Gerald J. Lieberman. *Introduction to Operations Research*. McGraw Hill LLC, New York, NY, 2024 release, international student edition edition, 2024.
- [26] Kris Braekers, Katrien Ramaekers, and Inneke Van Nieuwenhuysse. The vehicle routing problem: State of the art classification and review. *Computers & Industrial Engineering*, 99:300–313, September 2016.
- [27] Eshetie Berhan, Birhanu Beshah, Daniel Kitaw, and Ajith Abraham. Stochastic Vehicle Routing Problem: A Literature Survey. *Journal of Information & Knowledge Management*, 13(03):1450022, September 2014.
- [28] Gilbert Laporte, François V. Louveaux, and Luc Van Hamme. An Integer L -Shaped Algorithm for the Capacitated Vehicle Routing Problem with Stochastic Demands. *Operations Research*, 50(3):415–423, June 2002.
- [29] Roberto Roberti and Paolo Toth. Models and algorithms for the Asymmetric Traveling Salesman Problem: an experimental comparison. *EURO Journal on Transportation and Logistics*, 1(1-2):113–133, June 2012.
- [30] Paolo Toth and Daniele Vigo, editors. *The Vehicle Routing Problem*. Society for Industrial and Applied Mathematics, January 2002.
- [31] Hipólito Hernández-Pérez and Juan-José Salazar-González. A branch-and-cut algorithm for a traveling salesman problem with pickup and delivery. *Discrete Applied Mathematics*, 145(1):126–139, December 2004.
- [32] Reinelt, Gerhard. TSPLib, December 2024.
- [33] A. Juan, J. Faulin, S. Grasman, D. Riera, J. Marull, and C. Mendez. Using safety stocks and simulation to solve the vehicle routing problem with stochastic demands. *Transportation Research Part C: Emerging Technologies*, 19(5):751–765, August 2011.
- [34] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2023.
- [35] Markus Leitner, Ivana Ljubić, Martin Luipersbeck, and Markus Sinnl. Decomposition methods for the two-stage stochastic Steiner tree problem. *Computational Optimization and Applications*, 69(3):713–752, April 2018.
- [36] Guido Schäfer. Steiner Forest: 1995; Agrawal, Klein, Ravi. In Ming-Yang Kao, editor, *Encyclopedia of Algorithms*, pages 897–900. Springer US, Boston, MA, 2008.
- [37] Ivana Ljubić. Solving Steiner trees: Recent advances, challenges, and perspectives. *Networks*, 77(2):177–204, March 2021.
- [38] Anupam Gupta and Amit Kumar. A constant-factor approximation for stochastic Steiner forest. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 659–668, Bethesda MD USA, May 2009. ACM.
- [39] Nicole Immorlica, David Karger, Maria Minkoff, and Vahab S. Mirrokni. On the costs and benefits of procrastination: approximation algorithms for stochastic combinatorial optimization problems. pages 691–700, New Orleans, Louisiana, November 2004. Society for Industrial and Applied Mathematics.
- [40] Anupam Gupta, Martin Pál, R. Ravi, and Amitabh Sinha. Boosted sampling: approximation algorithms for stochastic optimization. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 417–426, Chicago IL USA, June 2004. ACM.
- [41] Chaitanya Swamy and David B. Shmoys. Approximation algorithms for 2-stage stochastic optimization problems. *ACM SIGACT News*, 37(1):33–46, March 2006.
- [42] Anupam Gupta and Martin Pál. Stochastic Steiner Trees Without a Root. In David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos, Dough Tygar, Moshe Y. Vardi, Gerhard Weikum, Luís Caires, Giuseppe F. Italiano, Luís Monteiro, Catuscia Palamidessi, and Moti Yung, editors, *Automata, Languages and Programming*, volume 3580, pages 1051–1063. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005. Series Title: Lecture Notes in Computer Science.
- [43] Anupam Gupta, R. Ravi, and Amitabh Sinha. LP Rounding Approximation Algorithms for Stochastic Network Design. *Mathematics of Operations Research*, 32(2):345–364, May 2007.
- [44] Anupam Gupta, MohammadTaghi Hajiaghayi, and Amit Kumar. Stochastic Steiner Tree with Non-uniform Inflation. In Moses Charikar, Klaus Jansen, Omer Reingold, and José D. P. Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 4627, pages 134–148. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007. Series Title: Lecture Notes in Computer Science.
- [45] Lisa Fleischer, Jochen Könnemann, Stefano Leonardi, and Guido Schäfer. Strict Cost Sharing Schemes for Steiner Forest. *SIAM Journal on Computing*, 39(8):3616–3632, January 2010.
- [46] Immanuel Bomze, Markus Chimani, Michael Jünger, Ivana Ljubić, Petra Mutzel, and Bernd Zey. Solving Two-Stage Stochastic Steiner Tree Problems by Two-Stage Branch-and-Cut. In Otfried Cheong, Kyung-Yong Chwa, and Kunsoo Park, editors, *Algorithms and Computation*, volume 6506, pages 427–439. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010. Series Title: Lecture Notes in Computer Science.
- [47] Bernd Zey. ILP formulations for the two-stage stochastic Steiner tree problem, November 2016. arXiv:1611.04324 [cs].
- [48] Daniel Schmidt, Bernd Zey, and François Margot. Stronger MIP formulations for the Steiner forest problem. *Mathematical Programming*, 186(1-2):373–407, March 2021.
- [49] Lloyd’s Register. Rules and Regulations for the Classification of Ships, 2023. url: <https://www.lr.org/en/knowledge/lloyds-register-rules/>.
- [50] B.T. Markhorst, J. Berkhout, A. Zocca, J.F.J. Pruyn, and R.D. Van Der Mei. Sailing through uncertainty: ship pipe routing and the energy transition. *International Marine Design Conference*, May 2024. Publisher: TU Delft OPEN Publishing.
- [51] Zey, Bernd. SSTPLib, 2024.

Appendix A. Flow-based ILP

We present the flow-based version of ILP (7) in (A.1) and introduce one new decision variable. The binary deci-

sion variable $f_{ktmuv}^{(s)}$ equals 1 if a flow is sent from the root of terminal group k to terminal t via connection type m at arc (u, v) .

(DO-2b)

$$\min \sum_{((u,v),m) \in (\mathcal{E} \times \mathcal{M})} \left(x_{muv}^{(0)} \cdot c_{muv}^{(0)} \right) \quad (\text{A.1a})$$

$$\text{s.t.} \quad \sum_{m \in M^{(0)}} \left(\sum_{u:(v,u) \in A^{(0)}} f_{ktmuv}^{(0)} - \sum_{u:(u,v) \in A^{(0)}} f_{ktmuv}^{(0)} \right) = \begin{cases} z_{kl}^{(0)} & \text{if } v = r^k \\ -z_{kl}^{(0)} & \text{if } v = t \\ 0 & \text{otherwise} \end{cases} \quad \begin{cases} \forall k \in \mathcal{K}^{(0)}, \forall t \in \mathcal{T}_r^{k \dots K^{(0)}} \\ \forall v \in \mathcal{V} \text{ with } \tau(t) = l \end{cases} \quad (\text{A.1b})$$

$$f_{ktmuv}^{(0)} \leq y_{kmuv}^{(0)} \quad \begin{cases} \forall k \in \mathcal{K}^{(0)}, \forall t \in \mathcal{T}_r^{k \dots K^{(0)}} \\ \forall m \in M^{(0)}, \forall (u, v) \in A^{(0)} \end{cases} \quad (\text{A.1c})$$

$$(7c) - (7o) \quad (\text{A.1d})$$

$$\sum_{m \in M^{(0)}} \sum_{u:(t,u) \in A^{(0)}} f_{ktmuv}^{(0)} = 0 \quad \forall k \in \mathcal{K}^{(0)}, \forall t \in \mathcal{T}_r^{k \dots, K^{(0)}} \quad (\text{A.1e})$$

$$f_{ktmuv}^{(0)} \in \mathbb{B} \quad \begin{cases} \forall k \in \mathcal{K}^{(0)}, \forall t \in \mathcal{T}_r^{k \dots K^{(0)}} \\ \forall m \in M^{(0)}, \forall (u, v) \in A^{(0)} \end{cases} \quad (\text{A.1f})$$

The objective in (A.1a) is similar to (7a). Constraints in (A.1b) guarantee that each terminal is included in an arborescence with its root at r^k for some $k \in \mathcal{K}$. An artificial flow is distributed from each root r^k to all other terminals in the respective arborescence. The decision variables f_{ktmuv} trigger y_{kmuv} in (A.1c) when a flow travels from the root r^k to terminal t using connection type m through arc (u, v) . Constraints (A.1e) ensure that flow does not leave a terminal, and is meant to improve the model's LP-relaxation. As (A.1f), which makes $f_{ktmuv}^{(0)}$ a binary decision variable, already ensures the integrality of y_{kmuv} , constraints in (7m) and (7n) may be relaxed.