Robust Network Capacity Expansion with Non-Linear Costs

Francis Garuba[1]
Department of Management Science, Lancaster University, United Kingdom
f.garuba@lancaster.ac.uk

Marc Goerigk
Network and Data Science Management, University of Siegen, Germany, Germany
marc.goerigk@uni-siegen.de

Peter Jacko
Department of Management Science, Lancaster University, United Kingdom
p.jacko@lancaster.ac.uk

Abstract

The *network capacity expansion problem* is a key network optimization problem practitioners regularly face. There is an uncertainty associated with the future traffic demand, which we address using a scenario-based robust optimization approach. In most literature on network design, the costs are assumed to be linear functions of the added capacity, which is not true in practice. To address this, two non-linear cost functions are investigated: (i) a linear cost with a fixed charge that is triggered if any arc capacity is modified, and (ii) its generalization to piecewise-linear costs. The resulting mixed-integer programming model is developed with the objective of minimizing the costs.

Numerical experiments were carried out for networks taken from the SNDlib database. We show that networks of realistic sizes can be designed using non-linear cost functions on a standard computer in a practical amount of time within negligible suboptimality. Although solution times increase in comparison to a linear-cost or to a non-robust model, we find solutions to be beneficial in practice. We further illustrate that including additional scenarios follows the law of diminishing returns, indicating that little is gained by considering more than a handful of scenarios. Finally, we show that the results of a robust optimization model compare favourably to the traditional deterministic model optimized for the best-case, expected, or worst-case traffic demand, suggesting that it should be used whenever computationally feasible.

2012 ACM Subject Classification Networks; Theory of computation → Network optimization; Theory of computation → Continuous optimization

Keywords and phrases Robust Optimization, Mobile Network, Network Capacity Design & Expansion, Non-linear Cost, Traffic and Transport Routing

Digital Object Identifier 10.4230/OASIcs.ATMOS.2019.5

1 Introduction

Network design and capacity planning has always been of strategic importance in most organization. This implies that it needs to be decided far ahead of time based on the estimation of future traffic demand. Projection for future traffic is usually done using traffic measurements and population statistics in combination with other marketing data. This often results in a large discrepancy between planned and actual carried traffic volume and distribution.

[1] Corresponding author

19th Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2019).
Editors: Valentina Cacchiani and Alberto Marchetti-Spaccamela; Article No. 5; pp. 5:1–5:13
OpenAccess Series in Informatics
OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

To provide a more detailed motivation and positioning of our paper, we focus on the telecommunications field (other network design applications, such as line planning for public transport, are also well within the scope of this work). Here, this discrepancy could be as large as 10% according to [3]. Hence, the re-forecasting and re-planning becomes a continuous exercise using traffic measurements and traffic optimization tools, which are often based on deterministic concepts assuming the traffic demand is estimated without error.

The demand for capacity in mobile wireless networks has seen an ever-growing trend in the last couple of decades and growth rate is expected to be even higher going into the future. This explosion in demand for data is coming at a lower cost rate. This means that in order to provide an acceptable quality of service, capacity will need to be regularly extended with optimal investment in capital expenditure. This balancing act of traffic volume, quality of service and capital expenditure has made network capacity expansion a key strategic function resulting in high global telecoms investment. Similar capacity expansion challenges are present to network designers and operators in other types of networks as well, such as transport networks. The *network capacity expansion problem* can hence be considered one of the key network optimization problems practitioners are expected to regularly face in present and future.

To have a network that is robust against uncertain estimated traffic demand, this uncertainty needs to be factored in already during the planning and design process, which we address using a scenario-based robust optimization approach. This methodology is geared towards producing results that are insensitive to the uncertain demand, by solving the problem using two separate stages. In the first stage, we determine the capacity expansion, and in the second stage, demand scenarios are realized. The resulting mixed-integer programming model is developed with the objective of minimizing costs.

In most literature on network design, costs are assumed to be linear functions of the added capacity, which is not true in practice. Real-world costs typically follow a volume discount regime which is reflected by a non-linear function and can be attributed to bulk buy. To address this, two non-linear cost functions are investigated in this paper: (i) a linear cost with a fixed charge that is triggered if any arc capacity is modified, and (ii) its generalization that is piecewise-linear in added capacity.

To the best of our knowledge, this is the first paper that includes non-linear cost functions in the robust network capacity planning problem. This extension leads to a more computationally-demanding model than the one with linear cost. The contributions of our paper are as follows: We show that networks of realistic sizes can be designed using non-linear cost functions in a practical amount of time within negligible suboptimality. We present the benefits of considering a robust optimization model (even with two scenarios) instead of the traditional deterministic model, and present the benefits of considering non-linear costs instead of the usual linear costs. It is illustrated that including additional scenarios approximately follows the law of diminishing returns, indicating that little is gained by considering more than a handful of scenarios. Finally, we show that the results of a robust optimization model compare favourably to the traditional deterministic model optimized for the best-case, expected, or worst-case traffic demand, suggesting that it should be used whenever computationally feasible.

The rest of this paper is organized as follows. Section 2 presents a literature review of related research. In Section 3, we then introduce the problem description of robust network capacity expansion and mathematical models. Experimental results using networks from the SNDLib (see [21]) are discussed in Section 4. Finally, Section 5 concludes our work and points out future research directions.

2 Literature Review

2.1 Robust Optimization in Network Design

In robust optimization, we assume that all possible data scenarios are given in form of an uncertainty set. For general surveys, we refer, e.g., to [13, 14]. The classic approach aims at finding a solution that is feasible for all scenarios from the uncertainty set, while optimizing a worst-case performance. This approach is relaxed through two-stage robust optimization, where not all decisions need to be taken in advance, see [6]. Instead, one distinguishes between "here and now" decisions that need to be fixed in advance, and "wait and see" variables that are determined once a scenario has been revealed. Two-stage robust optimization problems are also known as adjustable robust counterparts.

Adjustable robust optimization has been applied to wireless telecommunication services in the area of network design and expansion. This helps to model decisions that are delayed in time, e.g., traffic needs to be routed only once the demand scenario is known. Three closely related problems are the radio network design problem, the radio network loading problem and the virtual private network problem [17].

In telecoms, the long term strategic network planning can be viewed as the first stage "here and now" decision making, while the traffic redistribution that occurs after the realisation of the traffic demand pattern would be the second stage "wait and see" adjustment decision. Unrestricted second stage recourse in robust network design is called dynamic routing, see [7]. Most applications of adjustable robust optimization have focused on approximations that put a restriction on the recourse.

A special type of recourse restriction based on a specific type of uncertainty model (Hose model) has been proposed independently by [11] and [12] for an asynchronous transfer mode and broadband traffic network. They also introduced the concept of static routing, which [5] applied under their generalized polyhedral uncertainty model using a column and constraint generation algorithm. [20] investigated network capacity expansion under demand and cost uncertainty and recently, [23] used a cutting plane algorithm while taking into consideration the outsourcing costs for unmet demand. Some papers use an affine decision rule to restrict the recourse decisions, thus creating a tractable robust counterpart. [22] introduced affine routing in their robust network capacity planning model, while [24] and [3] used polyhedral uncertainty sets. On the other hand, [2] study the problem in detail by exploiting the underlying network structure.

2.2 Related Work on Non-linear Cost Functions

In general, routing costs, transportation costs or capacity costs can be a non-linear functions of traffic flows. In the following, we review literature on fixed-charge costs and piecewise-linear costs.

2.2.1 Fixed-Charge Cost Models

In a network with fixed-charge costs, an initial outlay cost is incurred to make an arc available. In this setting, one needs to pay a fixed initial cost in addition to the arc expansion cost. The fixed costs could be the installation costs, cabinet outlay costs, additional energy or utility costs and line replacement costs. Applications are found in wide areas of network design problems and not limited to energy networks, transportation and communication. A survey is provided by [16] that demonstrate many applications in logistics, transportation and communications. The fixed-charge cost network design problem (FCND) has been found to be NP-hard, see [16, 19].

Literature on the FCND has concentrated on solution algorithms for the different model variants. [8] addressed the multi-commodity capacitated FCND using a cutting plane algorithm with an improvement on the mixed-integer programming (MIP) formulation. [9] presented a detailed survey on the use of Benders decomposition to solving a wide range of FCND's which includes two facility networks. This can be viewed as a two-commodity network with a variant that introduces a quality of service measure. In [1], a heuristic approach for separating and adding violated partition inequalities was implemented. [26] solved a FCND using a variant of Benders decomposition which they referred to as the Bender-and-cut technique. The closest work to our model is [18]. Here, they formulate a robust network design problem with both transportation cost and demand uncertainty. Investment in arc capacity is modeled as a binary decision (i.e., expansion or no expansion). The model is approximated using an affine decision rule.

2.2.2 Piecewise-Linear Cost Models

The piecewise-linear cost model (PLC) can be used to model costs with economies of scale. In general, optimization problems involving PLC arise in domains including transportation, communications networks, large scale integrated circuits, supply chain management and logistics planning. They are usually modeled as MIPs, see [25]. The problem has been proven to be NP-hard for general concave cost objective functions, see [15].

As is the case for fixed-charge costs, most literature in this domain tends to focus on solution algorithms, see [10]. A continuous relaxation technique for solving network design with piecewise-linear costs was presented by [19]. [15] noted that exact techniques based on dynamic programming and branch and bound are only efficient for specific subclasses of the problem. A number of MIP model formulations exist for piecewise-linear functions. The names for these were unified in [27], which also provides a performance comparison. In terms of execution speed, they recommended the use of Multiple Choice Model (MCM) by [4] or the Incremental approach for a small number of segments.

3 Problem Formulation

We consider a multi-commodity network design problem where capacities are to be added on top of existing ones on a subset of arcs, with the aim of minimizing the total cost involved and so that routing of traffic for the different commodities over the arcs subject to design and network constraints is possible. We call this problem the *Robust Network Capacity Expansion Problem* (RNCEP). We first introduce the basic problem version with linear costs, before introducing two non-linear cost extensions.

3.1 RNCEP with Linear Costs

A communications network topology can be represented by a directed connected graph $G = (\mathcal{V}, \mathcal{A})$. Each of the arcs $a \in \mathcal{A}$ has an original capacity u_a. The original capacity on each arc a can be expanded at a cost c_a per each additional unit of capacity. A set of commodities $\mathcal{K}$ represents potential traffic demands. A commodity $k \in \mathcal{K}$ corresponds to node pair $(s^k, t^k) \in \mathcal{V} \times \mathcal{V}$ and a demand $d^k \geq 0$ for traffic from s^k to t^k. The actual demand values are considered to be uncertain and depend on random scenarios $\xi \in \Xi$. We assume a finite set $\Xi = \{\xi^1, \dots, \xi^N\}$ of possible demand scenarios and write $d^k(\xi)$ for the demand of pair (s^k, t^k) in scenario ξ.

The robust network capacity expansion problem is to find a minimum-cost installation of additional capacities while satisfying all traffic demands $d^k(\xi)$ for all $k \in \mathcal{K}$ and all $\xi \in \Xi$. In this respect, RNCEP is a two-stage robust program. The additional capacity we install on arc $a \in \mathcal{A}$ is denoted by x_a and is a first stage decision variable, which has to be fixed before observing a demand realization $\xi \in \Xi$. Once the demand scenario ξ becomes known, traffic is routed through a multi-commodity flow with variables $f_a^k(\xi)$.

Let $\delta^+(v)$ and $\delta^-(v)$ denote the sets of outgoing and incoming arcs at node $v \in \mathcal{V}$, respectively. The problem can now be formulated as the following linear program.

$$\min \sum_{a\in\mathcal{A}} c_a x_a \tag{1}$$

$$\text{s.t.} \sum_{a\in\delta^-(v)} f_a^k(\xi) - \sum_{a\in\delta^+(v)} f_a^k(\xi) = \begin{cases} -d^k(\xi) & \text{if } v = s^k \\ d^k(\xi) & \text{if } v = t^k \\ 0 & \text{otherwise} \end{cases} \quad \forall v \in \mathcal{V}, k \in \mathcal{K}, \xi \in \Xi \tag{2}$$

$$\sum_{k\in\mathcal{K}} f_a^k(\xi) \le u_a + x_a \quad \forall \xi \in \Xi, a \in \mathcal{A} \tag{3}$$

$$f_a^k(\xi) \ge 0 \quad \forall k \in \mathcal{K}, \xi \in \Xi, a \in \mathcal{A} \tag{4}$$

$$x_a \ge 0 \quad \forall a \in \mathcal{A} \tag{5}$$

Objective function (1) is to minimize the total cost of capacity expansion subject to flow conservation constraint (2), while constraint (3) imposes that the amount of flow does not exceed the sum of existing and added arc capacity.

3.2 RNCEP with Fixed-Charge Costs

We now introduce an extension of the previous model, where a fixed charge occurs if the capacity of an arc is modified. To this end, let p_a be this fixed charge associated with arc $a \in \mathcal{A}$.

We introduce a new variable $h_a \in \{0, 1\}$ to denote if the capacity of arc a is modified. The *RNCEP with fixed-charge costs* can then be formulated as the following mixed-integer program:

$$\min \sum_{a\in\mathcal{A}} (c_a x_a + h_a p_a) \tag{6}$$

$$\text{s.t. } x_a \le M_a h_a \quad \forall a \in \mathcal{A} \tag{7}$$

$$h_a \in \{0, 1\} \quad \forall a \in \mathcal{A} \tag{8}$$

$$\text{Constraints (2)} - \text{(5)} \tag{9}$$

Here, M_a for all a are constants that are sufficiently large not to restrict the solution. For instance, taking any $M_a \ge \max_{\xi\in\Xi} \sum_{k\in\mathcal{K}} d^k(\xi)$ for all a is valid.

3.3 RNCEP with Piecewise-Linear Cost

We further extend the RNCEP by introducing a piecewise-linear cost function. To this end, we apply the multiple choice model (MCM) as mentioned in the literature review. We assume that for every arc, there are up to S segments with different slopes in the cost function. Let us write $\mathcal{S} = \{1, \dots, S\}$. For every arc a and segment s, let b_a^s denote the load breakpoint, with an additionally defined $b_a^0 := 0$. Let c_a^s denote the cost slope of segment s, and p_a^s its y-intercept.

In addition to the variables of RNCEP, we introduce two new sets of auxiliary variables. Variables h_a^s are binary variables that select the cost segment where the added capacity x_a falls in. Variables x_a^s denote the amount of capacity that is added within each cost segment. This gives the following mixed-integer programming formulation for the *RNCEP with piecewise-linear costs*:

$$\min \sum_{s\in\mathcal{S}}\sum_{a\in\mathcal{A}}(c_a^s x_a^s + h_a^s p_a^s) \tag{10}$$

$$\text{s.t. } x_a = \sum_{s\in\mathcal{S}} x_a^s \qquad \forall a \in \mathcal{A} \tag{11}$$

$$b_a^{s-1} h_a^s \le x_a^s \le b_a^s h_a^s \qquad \forall a \in \mathcal{A}, s \in \mathcal{S} \tag{12}$$

$$\sum_{s\in\mathcal{S}} h_a^s \le 1 \qquad \forall a \in \mathcal{A} \tag{13}$$

$$x_a \le M_a \sum_{s\in\mathcal{S}} h_a^s \qquad \forall a \in \mathcal{A} \tag{14}$$

$$x_a^s \ge 0 \qquad \forall a \in \mathcal{A}, s \in \mathcal{S} \tag{15}$$

$$h_a^s \in \{0,1\} \qquad \forall a \in \mathcal{A}, s \in \mathcal{S} \tag{16}$$

$$\text{Constraints } (2) - (5) \tag{17}$$

4 Experimental Study

We implemented the fixed-charge cost model and the piecewise-linear cost model using instances from the SNDLib library by [21]. Network parameters characteristics on the four considered networks from SNDLib are presented in Table 1.

Table 1 Network parameters characteristics (rounded to integers).

Network	Janos26	Janos39	Sun27	Node39
$\lvert\mathcal{V}\rvert$	26	39	27	39
$\lvert\mathcal{A}\rvert$	84	122	102	172
$\lvert\mathcal{K}\rvert$	650	1,482	67	1,471
d^k (mean±SD)	123±198	69±243	28±16	5±2
u_a (mean±SD)	64±0	1,008±0	40±0	160±0
c_a (mean±SD)	468±225	313±162	19±10	23±11

Models were implemented using Julia and Gurobi version 7.5 on a Lenovo desktop machine with 8 GB RAM and Intel Core i5-6500 CPU at 2.50Ghz on 4 Cores. We used a time limit of 4000s for each problem instance and optimality is achieved once the optimality gap is below 0.01%.

4.1 Experimental Setup

Both the fixed-charge cost and the piecewise-linear cost models were implemented with one scenario (single-scenario) and with two scenarios (double-scenario). The base demand scenario was provided from the SNDLib library, which we randomly modified to generate additional demand scenarios. The amount of modification is controlled by a parameter λ, the maximum deviation of modified demand from the base demand. The parameter λ is chosen to be a fraction of the mean base demand $\widehat{d}$; we consider $\lambda = \text{round}(0.3\widehat{d})$ and

Table 2 Experimental setup for generating 120 problem instances for each network.

Parameters	# options	Options
Number of scenarios	2	1 (single) / 2 (double)
Scenario variability λ	2	$0.3\widehat{d}$ / $0.6\widehat{d}$
Fixed-charge factor P	3	0 / 10 / 100
Number of runs	10	—

Table 3 Proportion of instances not solved to optimality within the time limit (rounded to one decimal).

Network	Janos26	Janos39	Sun27	Node39
Total	0.0%	24.2%	35.0%	66.7%
$P = 0$	0.0%	0.0%	0.0%	0.0%
$P = 10$	0.0%	0.0%	12.5%	100.0%
$P = 100$	0.0%	72.5%	92.5%	100.0%
Single-scenario	0.0%	15.0%	28.3%	66.7%
Double-scenario	0.0%	33.3%	41.7%	66.7%

$\lambda = 2 \cdot \text{round}(0.3\widehat{d})$, corresponding to small uncertainty and large uncertainty, respectively. The value is then used as a bound for uniformly generating the modified demands around the base demand of every arc.

We summarize the experimental setup in Table 2. For each of the four networks, we consider the single-scenario and the double-scenario case, as well as small and large uncertainty. Additionally, for fixed-cost models we use three different fixed-charge factors P. These are used to calculate the fixed charges p_a of arc a by setting $p_a = Pc_a$. With $P = 0$, we recover the basic linear cost model without fixed charge. All networks and parameter settings are run 10 times to reduce variability in the results. In total, this gives $4 \cdot 2 \cdot 2 \cdot 3 \cdot 10 = 480$ optimization problem instances that need to be solved for the fixed charge case. For the piecewise-linear case, we follow the same setup with $4 \cdot 2 \cdot 2 \cdot 10 = 160$ instances. Each arc has three cost segments where the cost of each segment is calculated as ratio of the nominal arc cost. This gives segment costs as $c_a^s = c_a \cdot r_s$ where $r \in \{1.00, 0.90, 0.75\}$.

4.2 Results for RNCEP with Fixed-Charge Cost

4.2.1 Single- and Double-Scenario Results

Table 3 summarizes the results of the 480 problem instances, reporting the proportion of instances that were not solved to optimality within the time limit. We can see the optimization performance of problem instances in total, for different values of P, and for different number of scenarios. This performance measure gives a high-level summary of the hardness of particular instances. We can conclude that the instances become harder to solve as P increases, or as the number of scenarios increases.

Other performance metrics are presented in more detail in Table 4 and Table 5, where each cell gives an average and standard deviation from a sample of 20 problem instances. *Optimality gap* refers to the sub-optimality estimated and reported by Gurobi using the built-in procedure for lower-bounding the objective. *Solution time* is the time reported by Gurobi, capped by the time limit. *Capacity added* is the overall network capacity added on top of the original capacity (which can be calculated as $|\mathcal{A}|u_a$ from Table 1).

Table 4 Single-scenario results (rounded to one decimal).

		Janos26	**Janos39**	**Sun27**	**Node39**
Optimality gap	$P = 0$	0.0%	0.0%	0.0%	0.0%
	$P = 10$	0.0%	0.0%	0.0%	7.7 ±2.9%
	$P = 100$	0.0%	0.3 ±0.6%	5.0 ±2.8%	51.9 ±4.8%
Solution time	$P = 0$	6.5 ±0.5	156.9 ±17.0	0.3 ±0.1	536.4 ±82.2
	$P = 10$	7.4 ±0.6	227.1 ±86.0	201.7 ±201.4	4,000.1 ±0.0
	$P = 100$	10.8 ±2.1	3,120.9 ±1,088.0	3,694.8 ±815.9	4,000.1 ±0.1
Capacity added	$P = 0$	268,698 ±23,970	331,864 ±57,041	3,043 ±271	1,194 ±357
	$P = 10$	270,931 ±23,195	329,330 ±54,751	2,925 ±412	1,204 ±281
	$P = 100$	275,409 ±23,476	321,808 ±53,261	3,652 ±447	1,167 ±357

Table 5 Double-scenario results (rounded to one decimal).

		Janos26	**Janos39**	**Sun27**	**Node39**
Optimality gap	$P = 0$	0.0%	0.0%	0.0%	0.0%
	$P = 10$	0.0%	0.0%	0.1 ±0.2%	11.0 ±1.8%
	$P = 100$	0.0%	1.3 ±0.5%	10.8 ±1.4%	57.1 ±3.3%
Solution time	$P = 0$	88.4 ±25.1	1,285.6 ±349.5	1.2 ±0.2	2,256.6 ±317.9
	$P = 10$	92.2 ±21.0	2,373.9 ±770.5	1,729.0 ±1,418.2	4,000.2 ±0.1
	$P = 100$	189.0 ±57.7	4,000.3 ±0.2	4,000.1 ±0.1	4,000.2 ±0.1
Capacity added	$P = 0$	278,358 ±8,988	363,225 ±26,348	4,399 ±304	1,185 ±154
	$P = 10$	278,031 ±7,857	367,324 ±18,522	4,635 ±329	1,286 ±254
	$P = 100$	282,467 ±9,830	368,547 ±19,887	5,668 ±503	1,236 ±254

Interestingly, network Sun27 shows large variability in solution time, for both single-scenario and double-scenario settings. While with $P = 0$ it is the quickest to solve out of all networks, for larger values of P it is roughly similar to Janos39, despite dealing with a smaller number of commodities. On the other hand, solution time of Janos26 is affected very little by different values of P.

Comparing the solution time reported in Table 4 and Table 5, the double-scenario model, as expected, takes longer to solve to optimality as the goal here is to factor in robustness into the solution. On average, this double-scenario model resulted in 7.39% additional capacity across the networks for instances that were solved to optimality. The average increase in solution time across the instances that were solved to optimality is 828.24%.

We also note that capacity added is highly network dependent. The capacities of Janos26 and Janos39 are expanded dramatically due to the high variability in the demand, which for some commodities significantly exceeds the original capacity (see Table 1). On the other hand, the demands in Sun27 and Node39 are small compared to the original capacity, so the capacity added is relatively small.

Not reported elsewhere is the effect of scenario variability λ: the solution time becomes smaller if the uncertainty is larger, i.e., on the average for all the networks and parameter settings, the $0.6\widehat{d}$ variability results in lower solution times than for the $0.3\widehat{d}$ variability. This was also found to be the trend when looking at single networks. This is summarized in Table 6.

Overall, it is possible to solve most of the problem instances to optimality within the time limit, and even most of those not solved to optimality report very small optimality gap. The only settings that would significantly benefit from an increased time limit are Sun27 at $P = 100$ and Node39 at $P = 10$ and $P = 100$.

4.2.2 Effect of Number of Scenarios

While the previous discussion focused only on single- and double-scenario instances, it is also of interest to understand how an increased number of scenarios affects the performance measures. Considering more scenarios is expected to lead to a solution which in practical terms guarantees the network ability to accommodate a higher level of demand variation and provides additional capacity.

To illustrate that, we tested network Janos26 with fixed charge $P = 10$. We started with a single-scenario instance, where the base scenario considered reflects the *expected* demand (this is the original demand from SNDLib). We then generated and gradually added additional scenarios by randomly perturbating all the demands of the base scenario within $\pm\lambda$, in the large uncertainty setting.

For comparison, we also considered the *optimistic* instance, which is a single-scenario instance in which the demand is generated by subtracting λ from the expected demand on every arc. This instance expands the capacity of the network to satisfy only the smallest demand scenario, and would be almost surely unable to satisfy the realized demand. Finally, we considered the *pessimistic* instance, which is a single-scenario instance in which the demand is generated by adding λ to the expected demand on every arc. This instance expands the capacity of the network to satisfy all the possible demand scenarios.

The results are presented in Table 7. These results are representative; similar results were obtained when we replicated the experiment with other randomly generated scenarios. The key observations are as follows: By gradually expanding the set of scenarios, the cost (our minimization objective) non-decreases; the added capacity follows a similar trend, but is not necessarily monotone (cf. 8 vs 9 scenarios); the solution time (reported in seconds and as a multiple of the expected scenario instance) increases exponentially; expansion by adding more scenarios approximately follows the law of diminishing returns in both the cost and added capacity: the increase is highest when expanding from 1 (expected) scenario to 2 scenarios (which includes the expected scenario and one randomly generated), with only a minor increase when considering more than 3 scenarios, indicating the value of considering a robust optimization approach even with few scenarios; the increase in both the cost and added capacity is dramatic (36.9%) when expanding from 1 (expected) scenario to 2 scenarios (which includes the expected scenario and one randomly generated), indicating that optimizing the network based on the expected scenario (i.e. on point forecasts) only may be an inappropriate approach, leading to a large amount of unsatisfied realized demand; optimizing the network for the pessimistic scenario is very expensive (the increase in both the cost and added capacity is about 115% compared to the expected scenario), indicating the value of

Table 6 Effect of higher λ on solution time.

Solution Time	Single Scenario	Double Scenario
$\lambda = 0.3\widehat{d}$	527.31	3,010.85
$\lambda = 0.6\widehat{d}$	346.62	2,299.23
% Improvement	34.3%	23.6%

Table 7 Results on Janos26 with fixed-charge cost ($P = 10$) for different numbers of scenarios.

# Scenarios	Cost (in 10^3)	ΔCost	Added Capacity	ΔAdded Capacity	Time (sec.)	∝Time
1 (optimistic)	83,001	-10.9%	192,610	-9.2%	8	1x
1 (expected)	93,116	—	212,104	—	8	—
2	127,484	36.9%	292,893	38.1%	59	8x
3	129,804	39.4%	298,131	40.6%	376	50x
4	130,265	39.9%	300,426	41.6%	768	102x
5	130,272	39.9%	300,492	41.7%	1,080	143x
6	130,462	40.1%	300,913	41.9%	3,124	413x
7	130,753	40.4%	301,598	42.2%	2,488	329x
8	131,206	40.9%	301,936	42.4%	4,456	589x
9	131,255	41.0%	301,715	42.2%	8,869	1173x
1 (pessimistic)	200,593	115.4%	456,182	115.1%	8	1x

Table 8 Solution results for piecewise-linear cost.

Single-Scenario	Sun27	Janos26	Janos39	Node39
Optimality Gap	0.00%	2.90%	10.43%	22.43%
Solution time	653.67 ±640.84	4000.22 ±0.11	4000.22 ±0.06	4000.16 ±0.04
Capacity Added	2,863 ±539	276,172 ±26,036	335,258 ±58,895	1,472 ±574
Double-Scenario				
Optimality Gap	1.43%	6.73%	37.44%	77.99%
Solution time	4000.04 ±0.01	4000.21 ±0.23	4000.10 ±0.03	4000.12 ±0.04
Capacity Added	4,380 ±278	296,354 ±11,398	472,889 ±110,491	4,117 ±2,601

considering a robust optimization approach even with few scenarios; optimizing the network for the optimistic scenario leads to savings (the decrease in both the cost and added capacity is about 10% compared to the expected scenario), but may not be acceptable in practice if the consequences of having practically no satisfied realized demand are non-negligible.

These results provide an indication of the ability of our model to become more robust by including more demand scenarios. We note that Gurobi was able to deal with up to approximately 200 scenarios for this network without giving an out-of-memory error, however, it would be unlikely to compute a close-to-optimal solution in a reasonable amount of time.

4.3 Results for RNCEP with Piecewise-Linear Costs

Next we consider the robust network capacity expansion problem with piecewise-linear costs. Overall, 12.5% of all problem instances were solved to optimality within the time limit, 77.5% returned a non-optimal solution, and 10% were timed out already during the root relaxation. None of the double-scenario problem instances reached optimality within the time limit. Only one of the networks, Sun27, reached optimality and this was for all the problem instances in the single-scenario case. Two networks, Janos39 and Node39, had instances timing out under the root relaxation phase.

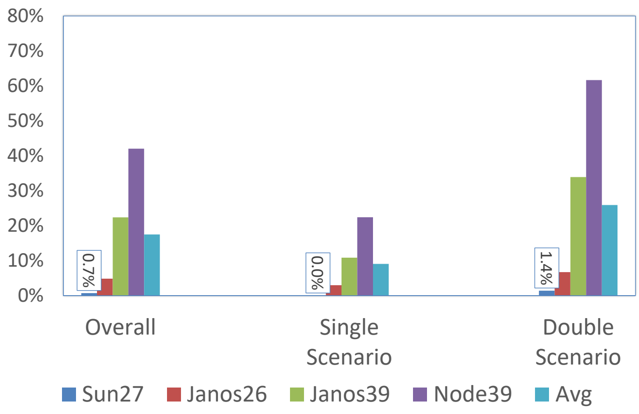

Figure 1 Optimality gap for piecewise-linear cost.

Table 8 presents more detailed results of this model for each network. The optimality gap is further illustrated in Figure 1, indicating that the optimality gap may be acceptable because of small values and small variability for Sun27 and Janos26 in the single-scenario setting and for Sun27 in the double-scenario setting. Better solutions can of course be achieved by increasing the time limit, which would be recommendable in the remaining settings.

The optimality gap provides insight into the increased difficulty of solving these problem instances, which also translates into longer solution time. It takes at least 512% more time to solve the double-scenario models compared to the single-scenario using Sun27 network, which is the easiest setting considering its very low optimality gap of 1.43% for the double-scenario instances. A further analysis was performed on the solution time using the paired sample t-Test which indicates no significant difference between solution time returned by $0.3\widehat{d}$ and $0.6\widehat{d}$ with a t-statistic of -0.2047 and a p-value 0.8423.

5 Conclusions

In this paper, a robust approach to network capacity expansion with non-linear cost functions was investigated. We developed robust models with fixed-charge costs and with piecewise-linear costs. They were implemented on four networks taken from the SNDlib, [21], with results compared to using linear costs. In the experimental setup, a number of possible parameter configurations was considered, including different demand variability and fixed-charges.

When further increasing the number of scenarios, we found that results follow a law of diminishing returns. While objective values and added capacity change little beyond five scenarios, computation times increase considerably. This is an indicator that already few scenarios suffice to find solutions that are robust against uncertainty in demand. The next pursuit will be to further improve the solution time for these models testing a path-based flow formulation and by developing specialized algorithms based on column generation and Benders decomposition.

References

1 Y. K. Agarwal and Y. P. Aneja. Fixed charge multicommodity network design using p-partition facets. *European Journal of Operational Research*, 258(1):124–135, 2017.

2 A. Atamturk and M. Zhang. Two-Stage Robust Network Flow and Design Under Demand Uncertainty. *Operations Research*, 55(4):662–673, August 2007.

3 F. Babonneau, J.-P. Vial, O. Klopfenstein, and A. Ouorou. Robust capacity assignment solutions for telecommunications networks with uncertain demands. *Networks*, 62(4):255–272, 2013.

4 A. Balakrishnan and S. C. Graves. A composite algorithm for a concave-cost network flow problem. *Networks*, 19(2):175–202, 1989.

5 W. Ben-Ameur and H. Kerivin. Routing of Uncertain Traffic Demands. *Optimization and Engineering*, 6(3):283–313, September 2005.

6 A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351–376, March 2004.

7 C. Chekuri, F.B. Shepherd, G. Oriolo, and M.G. Scutellá. Hardness of robust network design. *Networks*, 50(1):50–54, 2007.

8 M. Chouman, T. G. Crainic, and B. Gendron. Commodity Representations and Cut-Set-Based Inequalities for Multicommodity Capacitated Fixed-Charge Network Design. *Transportation Science*, 51(2):650–667, 2017.

9 A. M. Costa. A Survey on Benders Decomposition Applied to Fixed-Charge Network and Design Problems. *Computers and Operations Research*, 32(6):1429–1450, June 2005.

10 K. L. Croxton, B. Gendron, and T. L. Magnanti. Variable Disaggregation in Network Flow Problems with Piecewise Linear Costs. *Operations Research*, 55(1):146–157, 2007.

11 N. G. Duffield, P. Goyal, A. Greenberg, P. Mishra, K. K. Ramakrishnan, and J. E. Van der Merive. A flexible model for resource management in virtual private networks. *ACM SIGCOMM Computer Communication Review*, 29(4):95–108, October 1999.

12 J. A. Fingerhut, S. Suri, and J. S. Turner. Designing Least-Cost Nonblocking Broadband Networks. *Journal of Algorithms*, 24(2):287–309, August 1997.

13 V. Gabrel, C. Murat, and A. Thiele. Recent advances in robust optimization: An overview. *European Journal of Operational Research*, 235(3):471–483, June 2014.

14 M. Goerigk and A. Schöbel. Algorithm engineering in robust optimization. In *Algorithm engineering*, pages 245–279. Springer International Publishing, 2016.

15 G. M. Guisewite and P. M. Pardalos. Minimum concave-cost network flow problems: Applications, complexity, and algorithms. *Annals of Operations Research*, 25(1):75–99, December 1990.

16 T. L. Magnanti and R. T. Wong. Network Design and Transportation Planning: Models and Algorithms. *Transportation Science*, 18(1):1–55, February 1984.

17 S. Mattia. The robust network loading problem with dynamic routing. *Computational Optimization and Applications*, 54(3):619–643, August 2012.

18 S. Mudchanatongsuk, F. Ordóñez, and J. Liu. Robust solutions for network design under transportation cost and demand uncertainty. *Journal of the Operational Research Society*, 59(5):652–662, May 2008.

19 A. Nahapetyan and P. Pardalos. Adaptive dynamic cost updating procedure for solving fixed charge network flow problems. *Computational Optimization and Applications*, 39(1):37–50, January 2008.

20 F. Ordóñez and J. Zhao. Robust capacity expansion of network flows. *Networks*, 50(2):136–145, 2007.

21 S. Orlowski, R. Wessäly, M. Pióro, and A. Tomaszewski. SNDlib 1.0-Survivable Network Design Library. *Networks.*, 55(3):276–286, May 2010.

22 A. Ouorou and J.-P. Vial. A model for robust capacity planning for telecommunications networks under demand uncertainty. In *Design and Reliable Communication Networks, 2007. DRCN 2007. 6th International Workshop on*, pages 1–4. IEEE, 2007.

23 A. A. Pessoa and M. Poss. Robust Network Design with Uncertain Outsourcing Cost. *INFORMS Journal on Computing*, 27(3):507–524, August 2015.

24 M. Poss and C. Raack. Affine recourse for the robust network design problem: Between static and dynamic routing. *Networks*, 61(2):180–198, September 2012.

25 S. Sridhar, J. Linderoth, and J. Luedtke. Locally ideal formulations for piecewise linear functions with indicator variables. *Operations Research Letters*, 41(6):627–632, 2013.

26 V. Sridhar and J. S. Park. Benders-and-cut algorithm for fixed-charge capacitated network design problem. *European Journal of Operational Research*, 125(3):622–632, September 2000.

27 J. P. Vielma, S. Ahmed, and G. Nemhauser. Mixed-Integer Models for Nonseparable Piecewise-Linear Optimization: Unifying Framework and Extensions. *Operations Research*, 58(2):303–315, 2010.