

THOMAS BAUSCHERT¹, CHRISTINA BÜSING²,
FABIO D'ANDREAGIOVANNI^{3 4}, ARIE M.C.A. KOSTER⁵,
MANUEL KUTSCHKA⁶, UWE STEGLICH⁷

Network Planning under Demand Uncertainty with Robust Optimization

¹Chair of Communication Networks, Technical University Chemnitz, Reichenhainer Str. 70 09126 Chemnitz, Germany

²Chair of Operations Research, RWTH Aachen University, Kackertstr. 7, 52072 Aachen, Germany

³DFG Research Center MATHEON, Technical University Berlin, Straße des 17. Juni 135, 10623 Berlin, Germany

⁴Dept. of Optimization, Zuse-Institut Berlin (ZIB), Takustr. 7, 14195 Berlin, Germany

⁵Lehrstuhl II für Mathematik, RWTH Aachen University, Wüllnerstr. 7, 52056 Aachen, Germany

⁶Lehrstuhl II für Mathematik, RWTH Aachen University, Wüllnerstr. 7, 52056 Aachen, Germany

⁷Chair of Communication Networks, Technical University Chemnitz, Reichenhainer Str. 70 09126 Chemnitz, Germany

Herausgegeben vom
Konrad-Zuse-Zentrum für Informationstechnik Berlin
Takustraße 7
D-14195 Berlin-Dahlem

Telefon: 030-84185-0
Telefax: 030-84185-125

e-mail: bibliothek@zib.de
URL: <http://www.zib.de>

ZIB-Report (Print) ISSN 1438-0064
ZIB-Report (Internet) ISSN 2192-7782

Network Planning under Demand Uncertainty with Robust Optimization

Thomas Bauschert, Chemnitz University of Technology
Christina Büsing, RWTH Aachen University
Fabio D'Andreagiovanni, Zuse Institute Berlin
Arie M. C. A. Koster, RWTH Aachen University
Manuel Kutschka, RWTH Aachen University
Uwe Steglich, Chemnitz University of Technology

This work was supported by the German Federal Ministry of Education and Research (BMBF grants 05M10PAA, 05M10OCA, and 05M10ZAA), project *ROBUKOM – Robust Communication Networks*, <http://www.robukom.de>.

Abstract

The planning of a communication network is inevitably depending on the quality of both the planning tool and the demand forecast used. In this article, we show exemplarily how the emerging area of Robust Optimization can advance the network planning by a more accurate mathematical description of the demand uncertainty. After a general introduction of the concept and its application to a basic network design problem, we present two applications: multi-layer and mixed-line-rate network design. We conclude with a discussion of extensions of the robustness concept to increase the accuracy of handling uncertainties.

Introduction

Mathematical tools play a vital role in the design and operation of communication networks as, for example, the *Handbook of Optimization in Telecommunications* [13] shows. The concept of (directed) graphs and elementary algorithms for computing a shortest path or a spanning tree are core components of communication networking [11]. Many new innovations in technology and network management are first rendered precisely by a mathematical model of the optimization problem (e.g., an integer linear

program), that needs to be solved. The network planner is then assisted by mathematical software tools in solving such models. In addition, the challenges to plan more and more complicated communication networks have been one of the main driving forces of new solution methods in the mathematical optimization community.

This interaction between theory and practice received a new impulse by the technical progress to collect large amount of historical data. Until recently, the most successful applications (in telecommunications and beyond) of mathematical optimization involved a deterministic estimation of all relevant parameters like traffic demand values between core router locations. In a time that traffic can be logged in very small time intervals, network planning based on a single traffic matrix seems outdated. The emerging branch of Robust Optimization addresses this issue, by taking into account the uncertainty of the input parameters beyond estimations.

In this paper, we provide an introduction to Robust Optimization and its application to different communication network settings recently studied by the authors in the context of a 3-year research program supported by the German government and in collaboration with Nokia Siemens Networks and DFN-Verein

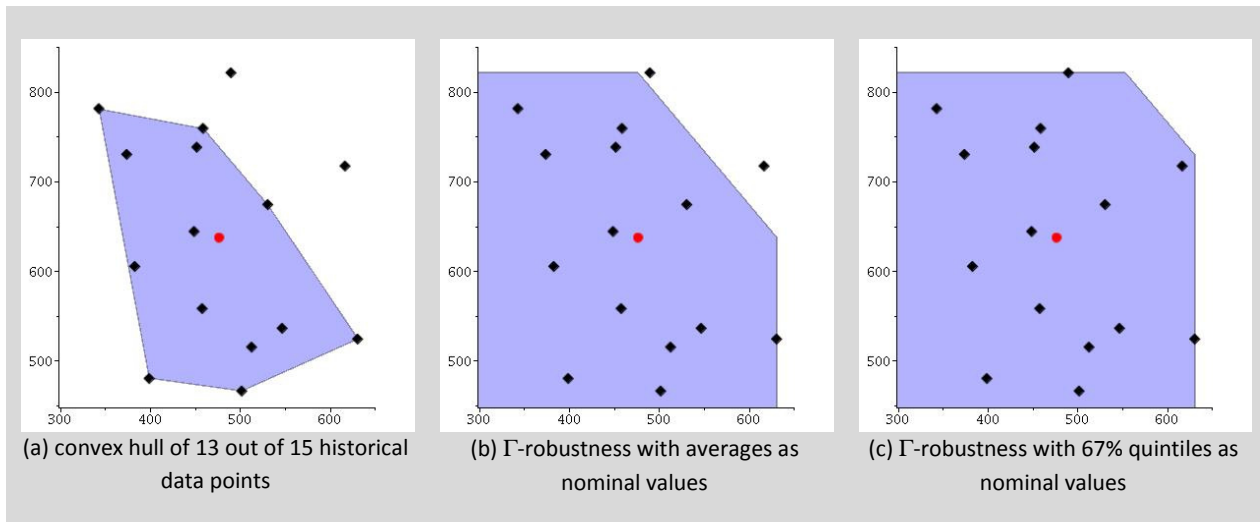


Figure 1: Possible uncertainty sets for two traffic flows covering.

(the operator of the German national research and education network).

Robust Optimization

For simplicity, let us consider a single link in a communication network and two traffic flows that can use this link. Historical data for both flows can be represented by a (x,y) -point in two-dimensional space. Figure 1(a) shows the traffic values during 15 points in history. The average traffic values are 477 Mbps (x-axis) and 637 Mbps (y-axis) as displayed by the red point. Taking those values and a link capacity of 1024 Mbps, one observes that on average 92% of the traffic can be routed across the link. Or alternatively, 100% of the first traffic flow and almost 86% of the second traffic flow. These solutions correspond to solutions of the following linear program:

$$\begin{aligned} \max \quad & x_1 + x_2 \\ \text{s.t.} \quad & 477x_1 + 637x_2 \leq 1024 \\ & 0 \leq x_1, x_2 \leq 1 \end{aligned}$$

where x_1, x_2 define the fraction of traffic routed. Feasible solutions are $(x_1, x_2) = (0.92, 0.92)$ and $(x_1, x_2) = (1.00, 0.86)$ (the latter solution is optimal). However, if we consider the

historical data, only 8 out of 15 traffic flows do not exceed the capacity in the first solution and, whereas in the second solution 10 out of 15 flows can be routed. Thus, the probability that the network link is overloaded is 46% in the first case and still 33% in the second case.

If we would like to have a solution such that in less than 15% of the historical cases the link is overloaded, we have to solve a robust linear program. Clearly the two coefficients are uncertain and taking the average traffic volume does not suffice. Robust Optimization offers an adequate way to incorporate uncertainties into our model: The uncertain coefficients are considered as random variables drawn from an *uncertainty set*. This uncertainty set describes all possible interactions between the uncertain coefficients and might look like the polyhedra in Figure 1(a)–(c). In fact, the polyhedron in Figure 1(a) is the convex hull of 13 out of 15 historical data points.

The task of Robust Optimization is to find a solution that is feasible for all considered realizations of the uncertain coefficients (from the uncertainty set) and maximizes the objective among these solutions. In case of our example, the solution $(x_1, x_2) = (1.00, 0.73)$ satisfies the

constraint regardless the values drawn from the uncertainty set and maximizes the sum among all robust feasible solutions. Accordingly, its usage would lead to a probability of overloading in about 13% of the historical cases.

A major challenge in Robust Optimization is the construction of a reasonable uncertainty set, in our example, the set depicted in Figure 1(a). Bertsimas and Sim [2] developed a generic uncertainty set that can be adjusted by a parameter $\Gamma \geq 0$. For each uncertain coefficient a_i we define a nominal value \bar{a}_i and a maximum deviation $\hat{a}_i \geq 0$. The Γ -robust uncertainty set is now defined as values $a_i \in [0, \bar{a}_i + \hat{a}_i]$ such that the sum of the relative excesses $\frac{a_i - \bar{a}_i}{\bar{a}_i + \hat{a}_i}$ of the nominal values is at most Γ . In Figure 1(b) and 1(c), the average traffic volumes are taken as nominal values and the difference to the maximum values as deviations. The uncertainty set in Figure 1(b) corresponds to $\Gamma = 1$, in Figure 1(c) to $\Gamma = 1.5$ (if more coefficients can deviate, typically integer values are taken for Γ). As the graphics show, 13 out of 15 historical data points are included in the uncertainty set induced by $\Gamma = 1$, whereas all of them are part of the set induced by $\Gamma = 1.5$. Here, the advantage of Robust Optimization comes into play: robust feasible solutions can be found without setting the input parameters to their most conservative estimation, i.e., the maximum values. By varying the Γ -value, the protection level against traffic fluctuations can be adapted to the needs of the planner. By comparing the network cost and the robustness level, network planners can exploit this trade-off for decision support.

Two major advantages of the Γ -robustness concept of Bertsimas and Sim [2] are:

- 1) as long as the uncertain coefficients are independently and symmetrically distributed, the probability that the constraint is violated by an optimal solution can be bounded by a function depending on the number of uncertain coefficients and the parameter Γ , i.e., given a value $\varepsilon > 0$, a value Γ can be chosen such that the probability of constraint satisfaction of the actual values is at least $1 - \varepsilon$ (see Bertsimas and Sim [2] for details);
- 2) the mathematical description of robust feasible solutions can be reformulated so that the size of the linear program is increased moderately, yielding a *compact* model, i.e., a model that is polynomial in the network size. Accordingly, the complexity increase of solving the linear program is bounded. This property will be explained exemplarily in the next section for the network design problem under demand uncertainty.

In the following sections, we give several examples from network design where a robust approach was successfully applied.

Network Design under Demand Uncertainty

The core of traffic engineering in a (backbone) communication network is the following technology-independent question: how to route the traffic flows from sources to destinations across the links, such that the capacity of those links is not exceeded? In the network design problem, this question is accompanied by the decision on the capacity granularities at the links with the aim to find a solution with minimum capacity installation cost.

Until recently, network planning was based on a single traffic matrix consisting of the forecasted traffic demands between every




Figure 2: Traffic fluctuations for three node-pairs in the US Abilene Internet2 network in time intervals of 5 minutes during one week (publicly available via [10]). In (a) the light colors represent the spare capacity for every single node pair, if capacity is reserved by the maximum traffic volume. In (b) the spare capacity is computed by the maximum sum of the traffic volumes. The purple line represents the sum of the 90% quintiles of the traffic volumes.

pair of network nodes. To avoid congestion in the designed network due to the dynamic traffic fluctuations (that frequently happen in modern communication networks) as shown in Figure 2, traffic estimates for every node-pair have to be very conservative. However, traffic peaks do not occur simultaneously for all traffic flows using the same link, and thus an unnecessarily high amount of resources are installed by such an approach.

The Γ -robustness concept provides a valuable alternative in this case. Instead of a single traffic forecast for every pair (s, t) of source and target nodes, a nominal demand \bar{d}^{st} and a deviation \hat{d}^{st} are defined for every node-pair. Let C be the installable capacity batch size, f_{ij}^{st} be the decision variable determining the fraction of the traffic flow between s and t via the link between nodes i and j , and x_{ij} the integer decision variable representing the number of capacity granularities to be installed.

Now, the capacity constraint for the link between nodes i and j is given by

$$\sum_{(s,t)} \bar{d}^{st} f_{ij}^{st} + \text{DEV}(f, \Gamma) \leq C x_{ij},$$

where $\text{DEV}(f, \Gamma)$ is the total capacity that has to be reserved to cope with the realized traffic values above the nominal values if the Γ -robust uncertainty set is used. $\text{DEV}(f, \Gamma)$ can be computed by the following linear program:

$$\begin{aligned} \text{DEV}(f, \Gamma) = \max \quad & \sum_{(s,t)} \bar{d}^{st} f_{ij}^{st} y^{st} \\ \text{s.t.} \quad & \sum_{(s,t)} y^{st} \leq \Gamma \\ & 0 \leq y^{st} \leq 1 \end{aligned}$$

By linear programming duality, the term $\text{DEV}(f, \Gamma)$ can be replaced by a linear function (on new variables) and additional linear constraints. Accordingly, the robust network design problem is formulated as an integer linear program, slightly more complicated than the network design problem for a single traffic matrix.




Figure 3: Price of robustness (relative to $\Gamma = 0$) for robust network design on the basis of historical data for the pan-European research backbone network GÉANT with observed means as nominal values and subject to different quintiles for the peak demand values (nominal + deviation).

Source: Koster and Kutschka [6].

The function of the network cost in dependence of the parameter Γ is known as the “price of robustness” [2]. It describes the additional cost of increasing the protection (and thus reducing the violation probability) by increasing Γ . Figure 3 shows the price of robustness for a computation based on historical data of a 22-node network.

What remains is the choice of the nominal and deviation values. Experiments with historical data have shown that the mean as nominal and the 95% quintile as peak value (nominal plus deviation) are good choices (Koster, Kutschka, Raack [7], see also [8]).

Robust Multi-Layer Network Design

General multilayer problem

The design problem outlined in the previous section describes a single-layer network problem. However, many communication networks are nowadays consisting of two and more technological layers, e.g. the Internet Protocol (IP) layer, the Multiprotocol Label Switching (MPLS) or MPLS Transport Profile (MPLS-TP) layer, the Optical Transport Network (OTN) layer and the Dense Wavelength Division Multiplex (DWDM) layer. Additionally there is a logical demand layer,

which induces traffic demand for arbitrary end-to-end connections. A wide range of technologically feasible layer configurations and possibilities for transporting the traffic demand through the layers exist, see Figure 4. Common layer configurations are for instance IP-over-DWDM or IP-over-MPLS-over-OTN-over-DWDM. A multi-layer network optimization formulation has to incorporate all technological and logical layers that should be part of the potential solution space.

Considering all constraints of a multi-layer network design problem in a generic mathematical formulation is a very challenging task. A too abstract layer model might neglect important technological constraints. On the other hand, a fine-grained formulation of the layers might lead to a huge computational complexity of the multi-layer model.

A comprehensive multi-layer modeling should integrate:

- layer model, e.g. multi-layer structure and feasible layer interconnections;
- technological restrictions, e.g. capacity granularities of interfaces (IF) and sub-

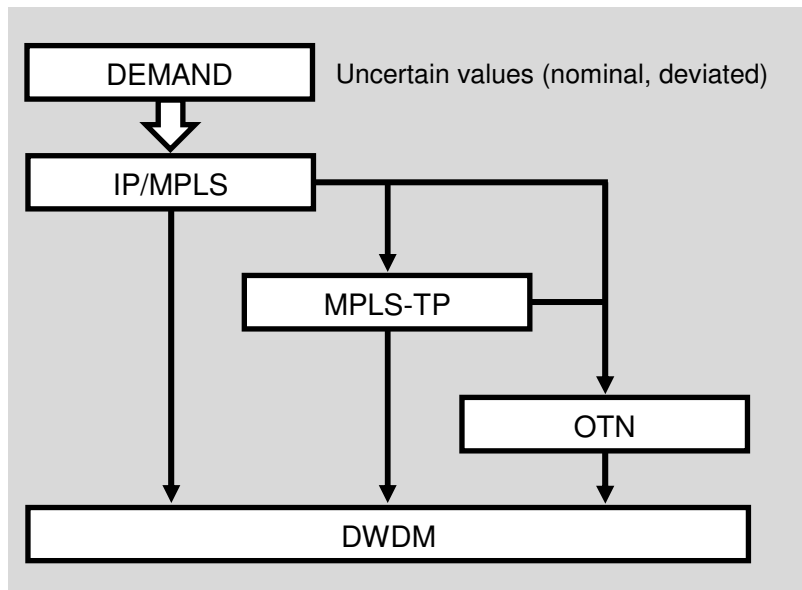


Figure 4: Feasible multi-layer interconnections and resulting layer configurations

interfaces, number of interface card slots, multiplexing capabilities;

- cost model, considering e.g. capital expenditures (CAPEX), operational expenditures (OPEX), energy consumption;
- traffic demand model (with or without demand uncertainty);
- model of resilience mechanisms, e.g. 1+1, 1:1 protection, re-routing .

Incorporation of Robustness in Multi-Layer Planning

Robustness in multi-layer networks can be defined in various ways. Frequently it is understood in the context of network resilience. Here the key question is which resilience concepts should be deployed in which layer and how different resilience mechanisms should interact. Incorporating resilience in multi-layer optimization requires a sophisticated modeling of the potential failure cases and the respective sequence of resilience mechanisms that react on these failures. For next generation networks employing IP-over-DWDM, an evaluation of different robust models (wrt. resilience) was done by Kubilinskas, Pióro and Nilsson [9].

Another interpretation of robustness in multi-layer network design is the ability to cope with uncertain traffic demand. Traffic demand

fluctuations can occur in temporal and spatial manner. The temporal effects can be classified into short-, mid- and long-term fluctuations. In particular the mid- and long-term effects such as the daytime usage behavior as depicted in Figure 2 are relevant for robust network design. On the other hand spatial traffic demand fluctuations are either caused by day of time traffic shifts (in large networks spanning over multiple time zones) or by effects outside the own network like Border Gateway Protocol (BGP) route flaps or dynamic server selection policies of Content Delivery Networks (CDN).

Concepts like the previously described Γ -robustness can be applied in multi-layer network design similar to the single-layer case. However, the complexity, model size and computation time are substantially increased by introducing Γ -robustness in multi-layer network optimization as shown in Steglich et al. [14]. Uncertainty in traffic affects the capacity dimensioning of all subjacent technological layers. In the lower layers traffic demand uncertainty is smoothed by *multiplexing* traffic from higher layers.

Layer (Technology) Selection and Optimization

Further potential challenges in multi-layer network design are the determination of the layers (technologies) that should be used given a set of potential networking technologies and the determination of the optimum connectivity (topology graph) within each layer.

Regarding the first challenge, *layer-skipping* is an option to reduce the network CAPEX. Although interfaces for connecting higher layers to lower ones (e.g., IP to DWDM IFs) are more expensive [12], the overall CAPEX might be cheaper than establishing an intermediate layer with further interfaces. The result of the optimization should reveal which particular layers are used and which layers are omitted. For this, layer configurations with possible layer sequences have to be included into the multi-layer optimization model.

To cope with the second challenge, flexible path sets (per layer) are included in the multi-layer network optimization. These path sets contain three types of paths: *opaque* paths (calculated by a k -shortest path algorithm), *transparent* paths (with no intermediate nodes), and specific paths where some of the intermediate nodes of opaque paths might be omitted. As a result of the multi-layer optimization the cheapest (in terms of the optimization objective) paths are selected, thus leading to *shortcuts* in some layers. The inclusion of such path sets (allowing the determination of shortcuts) influences the size of the multi-layer network optimization model significantly.

The well-known *IP router offloading problem* can be considered as a combination of the layer skipping and shortcut determination.

Results from ROBUKOM

In the ROBUKOM project, a multi-layer network design model with traffic demand uncertainty has been developed. This model

applies Γ -robustness to model traffic uncertainty. Moreover, aspects like layer-skipping, shortcuts and router-offloading are included. First computational results with off-the-shelf solvers are provided for small-, mid- and large-scale networks in Steglich et al. [14]. The introduction of Γ -robustness increases the CAPEX costs. For a 5-node network without layer-skipping, securing at most ten demands ($\Gamma=10$) is 23.0% more expensive compared to a non-robust network design. With layer-skipping it is 25.6% more expensive to consider traffic uncertainty. The uncertainty parameter Γ shows an even higher influence for the GÉANT network: CAPEX is raised here by 117.2% ($\Gamma=0$ vs. $\Gamma=10$).

In our future work, we intend to apply other techniques (like meta-heuristics) in order to reduce the computation times and memory requirements when dealing with large-scale network design.

Mixed-Line-Rate Optical Networks

In an optical network, lightpaths are used for transporting traffic flows. Mixed-line-rate optical networks allow for a more resource-efficient handling of small and large traffic volumes by the simultaneous configuration of lightpaths with different bitrates (e.g., 10 Gbps, 40 Gbps, and 100 Gbps).

Given a potential network topology and commodities with (uncertain) demand values, a cost-minimal hardware configuration (line-rate used for each demand, installed transponders, amplifiers, and regenerators) and an optimum routing have to be determined. Additional survivability requirements may exist.

In Duhovniko et al. [5] a mixed integer linear programming formulation for the design of mixed-line-rate networks with uncertain demands is given. In addition to the modeling of Γ -robustness, its main feature consists of

the computation of the nominal and peak demand values. In contrast to single-line-rate planning, the nominal and deviation values depend on the line-rate of the lightpath used for a particular demand. If small demands are routed on a lightpath with a high bit-rate, on the one hand additional lightpaths for absorbing traffic peaks are not needed, but on the other hand, the resources are not used efficiently. If lightpaths with a low bit-rate are used instead, traffic peaks might exceed the capacity reserved by the lightpaths for the nominal demand, and additional spare lightpaths have to be reserved to handle these peaks. Hence, depending on the line rate used, different nominal and deviation values have to be used. Figure 5(a) shows an example with a nominal demand (in 1 Gbps) of 65 and a deviation of 30. In case a line rate of 10 Gbps is chosen, 7 lightpaths have to be reserved for the nominal demand, and an additional 3 for peak values. In case 40 Gbps is chosen, 2 lightpaths are needed for the nominal demand

and another one for the peak. However, if 100 Gbps is chosen, a single lightpath provides enough capacity for the nominal as well as the peak demand, and thus no further deviation value is needed in this case.

Figure 5(b) shows exemplarily the cost of a robust mixed-line-rate optical network with GÉANT data, with and without 1+1 protection for different values of Γ . The costs are normalized to the case without protection and without robustness ($\Gamma = 0$). Not surprisingly the costs are more than doubled if 1+1 protection is implemented, but the price of robustness for unprotected cases is rather low. For robust designs with 1+1 protection in particular, the transponder cost increase significantly with increasing Γ , which can be explained by the need to use more and more high bit rate transponders since the number of lightpaths per fiber is limited.




Figure 5: Mixed-line-rate Optical Network Design. (a) Bandwidth requirement depending on the operated line rate for an example with nominal demand value of 65 Gbps and a deviation of 30 Gbps. (b) Cost of multi-line-rate optical networks normalized to $I=0$ and no protection. The costs are broken down with respect to transponder (TP) cost, amplifier (AM) cost, and regenerator (REG) cost.

Multiband Robust Network Design

In the previous sections, we showed the benefit of adopting Robust Optimization to tackle traffic uncertainty in the design of communication networks. A Robust Optimization approach guarantees a moderate dimensioning of the installed capacity and of the established routing paths that are able to accommodate (bad) traffic variations with respect to expected traffic values. Our primary robustness tool has been the model proposed by Bertsimas and Sim [2], essentially based on adopting a single deviation band for each uncertain parameter and an upper bound on the number of data that can simultaneously deviate from their nominal value.

The Bertsimas-Sim model has attracted high attention also outside the Mathematical Programming community and has been applied in many different contexts (see Bertsimas et al. [1] for an overview). Key factors of this success have been undoubtedly its elegant simplicity and its straightforward application: its use just requires fixing the upper bound Γ on the number of deviating data. Different solutions according to this parameter can then be compared by the corresponding price of robustness allowing an evaluation of the trade-off between stability and price by the decision maker.

However, the adoption of a single deviation band to represent the uncertain value of a coefficient may greatly limit the possibility of modeling uncertainty: in many real-world problems, the deviations distribute asymmetrically within often non-symmetric

intervals. In such cases, adopting a single deviation band and thus focusing just on the extreme deviations, neglects the inner uncertainty behavior and may lead to over conservative robust solutions that overestimate the impact of variances. It is thus desirable to increase the resolution of the model.

Increasing the resolution of the Bertsimas-Sim model can be done by a simple operation: partitioning the single deviation band into multiple bands, each with its own upper bound on the number of data falling into that band. Moreover, to further increase the power of modeling uncertainty, we can also introduce a lower bound on the number of deviations falling in each band: this simple trick allows to explicitly take into account also good deviations that in a Bertsimas-Sim approach are neglected, but that in reality are actually present with the effect of reducing the impact of bad deviations. We call an uncertainty set based on multiple deviation bands a *Multiband Set* and *Multiband Robustness* the resulting Robust Optimization model.

Multiband Robustness looks particularly attractive in real-world applications, where it is common to have historical data that shows the past behavior of the uncertainty. These data can be used to define histograms representing the (discrete) distribution of the uncertainty in the past and form a basis to build Multiband Sets, which are now strongly data-driven. We refer to Figure 6 for a visual representation of the differences between a single and a multiband representation of the uncertainty.




Figure 6: Visual comparison of a single and a multiband uncertainty set defined over the same overall deviation range. (a) is the histogram of deviations built upon the historical data and (b),(c) are possible single and multiband representation of the histogram.

Within the project ROBUKOM, we started to investigate the theoretical properties of Multiband Robustness. Here, we recall the main theoretical results that we have obtained and we refer the reader to Büsing and D'Andreagiovanni [3,4] for a complete and detailed overview of them. Given an uncertain Mixed-Integer Linear Program (MILP) and assuming that we represent uncertainty by a Multiband Set:

- the robust counterpart of a MILP is equivalent to a *compact* Mixed-Integer Linear Program, whose size grows linearly with the number of deviation bands of the Multiband Set and quadratically by the number of variables and constraints;
- verifying if a solution of MILP is robust w.r.t. the Multiband Set can be operated efficiently by solving a min-cost flow problem (note that this result can be used as basis to define an efficient cutting-plane solution algorithm, see [3] for details);
- if the uncertain MILP includes only binary variables and the uncertainty just affects the objective function, then a robust optimal solution can be obtained by solving a polynomial number of original MILPs with modified objective coefficients.

The application of Multiband Robustness to Network Design with demand uncertainty implies that the overall range of deviation

$[\bar{d}^{st} - \hat{d}^{st}, \bar{d}^{st} + \hat{d}^{st}]$ of each demand associated with a source-target pair (s,t) is partitioned into a number $K > 1$ of non-overlapping sub-bands. Each of these bands is then associated with a lower and an upper bound on the number of deviations that may fall in it (these should be derived from the historical data).

We carried out preliminary experiments about the adoption of Multiband Robustness in Network Design, referring to the well-known US Abilene Internet2 network instances. The number of deviation bands was fixed to seven and the extremes of the bands were defined according to the 50th, 70th, 75th, 80th, 85th, 90th, 95th percentile demand values, derived from historical data. The used bounds of each band took into account the probability of realization of the demands in each band. In comparison to a single band approach using a comparable and optimistic Γ parameter, the multiband approach granted a percentage reduction in the price of robustness between 1% and 5%, while maintaining the same computational performance (no significant increase in solution time). This is due to the refined representation of the uncertainty, which reduce conservatism of robust solutions. These preliminary results have encouraged ongoing investigations about a better tuning of the parameters of the

Multiband Set (number, bounds and width of the bands).

Conclusions

Robust optimization is an emerging mathematical optimization technique to deal with uncertain input parameters. In recent years, the methodology has also been applied to communication networks in various settings. Its potential has been clearly shown by those case studies and deserves a further integration in network planning tools in practice. Moreover, driven by the availability of historical data, the methodology is developed further as well to allow the usage of more accurate models.

References

- [1] Bertsimas, D., Brown, D., and Caramanis, C. (2011): Theory and Applications of Robust Optimization. In: *SIAM Review* 53 (3), pp. 464–501.
- [2] Bertsimas, D. and Sim, M. (2004): The Price of Robustness. In: *Operations Research* 52 (1), pp. 35–53.
- [3] Büsing, C. and D'Andreagiovanni F. (2012a): New Results about Multi-band Uncertainty in Robust Optimization. In: *Proc. of SEA 2012*, LNCS 7276, pp. 63–74.
- [4] Büsing, C. and D'Andreagiovanni F. (2012b): Robust Optimization under Multi-band Uncertainty - Part I: Theory. In: *arXiv abs/1301.2734*.
- [5] Duhovniko, S., Koster, A.M.C.A., Kutschka, M., Rambach, F., and Schupke, D. (2013): Γ -Robust Network Design for Mixed-Line-Rate-Planning of Optical Networks. In: *Optical Fiber Communication - National Fiber Optic Engineers Conference (OFC/NFOEC)*.
- [6] Koster, A.M.C.A. and Kutschka, M. (2011): Network Design under Demand Uncertainties: A Case Study on the Abilene and GÉANT network data. In: *Proceedings of 12. ITG-Fachtagung Photonische Netze*, pp. 154–161.
- [7] Koster, A.M.C.A., Kutschka, M., and Raack, C. (2011): On the robustness of optimal network designs. In: *Proceedings IEEE International Conference on Communications, ICC 2011*, S. 1–5.
- [8] Koster, A.M.C.A., Kutschka, M. and Raack, C. (2013): Robust Network Design: Formulations, Valid Inequalities, and Computations. In: *Networks* 61 (2), pp. 128–149.
- [9] Kubilinskas, E., Nilsson, P., & Pióro, M. (2005): Design Models for robust multi-layer next generation internet core networks carrying elastic traffic. In: *Journal of Network and Systems Management* 13 (1), pp. 57–76.
- [10] Orłowski, S., Wessäly, R., Pióro, M., Tomaszewski, A. (2010): SNDlib 1.0—survivable network design library. In: *Networks* 55(3), pp. 276–286.
- [11] Pióro, M. and D. Mehdi (2004): Routing, Flow, and Capacity Design in Communication and Computer Networks. Elsevier.
- [12] Rambach, F., Konrad, B., Dembeck, L., Gebhard, U., Gunkel, M., Quagliotti, M., Serra, L., and López, V. (2013): A multilayer cost model for metro/core networks. In: *Journal of Optical Communications and Networking* 5(3), pp. 210–225.
- [13] Resende, M.G.C. and Pardalos, P.M. (2006), Handbook of Optimization in Telecommunications. Springer.
- [14] Steglich, U., Bauschert, T., Büsing, C., and Kutschka, M. (2013): A Generic Multi-layer Network Optimization Model with Demand Uncertainty. In: *19th EUNICE Workshop on Advances in Communication Networking*.

Biographies

Thomas Bauschert (thomas.bauschert@etit.tu-chemnitz.de) received his Dipl.-Ing. and Dr.-Ing. degrees from TUM in 1990 and 1997, respectively. From 1997 to 2007 he was with Siemens and NSN in Munich, responsible for network planning, strategic product management, and research projects in the fields of highspeed fixed networks and mobile networks. In 2007 he was appointed full professor and head of the Chair for Communication Networks at Chemnitz University of Technology.

Christina Büsing (buesing@or.rwth-aachen.de) is a Senior Researcher in the Dept. of Operations Research at the RWTH Aachen University. She received her PhD in 2011 at the Technical University of Berlin and is working on Robust Optimization and its extension to Recoverable Robustness. In 2013 she was appointed as a member of the Junges Kollege North Rhine-Westphalian Academy of Sciences.

Fabio D'Andreagiovanni (d.andreagiovanni@zib.de) is a Senior Researcher in the Dept. of Optimization at Zuse-Institut Berlin since 2010. He has received his PhD in Operations Research at Sapienza Università di Roma (2010) and been Research Scholar in the Dept. of Industrial Engineering and Operations Research at Columbia University in the City of New York (2008–2009). His research on Network Design has received several awards, such as the Accenture M.Sc. Prize 2006, the INFORMS Telecom Doctoral Dissertation Award 2010 and the ESF-JSPS Excellence Award "Mathematics for Innovations" 2012.

Arie M. C. A. Koster (koster@math2.rwth-aachen.de) is since April 2009 Professor of

Network Planning under Demand Uncertainty with Robust Optimization

Mathematics at RWTH Aachen University. He holds a M.Sc. degree in technical mathematics (1995) from Delft University of Technology a Ph.D. degree in Mathematics of Operations Research (1999) from Maastricht University. From 1999 to 2007 he was a senior researcher with the Department of Optimization at Zuse Institute Berlin, from 2007 to 2009 Assistant Professor at Warwick Business School and the Centre for Discrete Mathematics and its Applications (DIMAP) of the University of Warwick. At RWTH Aachen University his research currently focuses on Robust Optimization and its applications. He coordinates the BMBF-project ROBUKOM - Robust Communication Networks.

Manuel Kutschka (Kutschka@math2.rwth-aachen.de) is a PhD student in the Discrete

Optimization group at RWTH Aachen University. He received his Diploma degree in business mathematics in 2007 from TU Berlin. His PhD thesis is about Robust Optimization, in particular the recoverable robust knapsack problem and the robust network design problem.

Uwe Steglich (uwe.steglich@etit.tu-chemnitz.de) is a PhD student at the Chair for Communications Networks at Chemnitz University of Technology. He received his Diploma degree in Information and Communications Technology in 2010 from Chemnitz University of Technology. His PhD thesis is about robust multi-layer network design considering traffic demand fluctuations.