Konrad-Zuse-Zentrum
für Informationstechnik Berlin

Takustraße 7
D-14195 Berlin-Dahlem
Germany

ANDREAS BLEY

Routing and Capacity Optimization for IP Networks

ZIB-Report 07-33 (September 2007)

Routing and Capacity Optimization for IP Networks*

Andreas Bley
Konrad-Zuse-Zentrum für Informationstechnik
Takustr. 7, 14195 Berlin, Germany

September 2007

The world-wide Internet is a huge, virtual network comprised of more than $13,000$ distinct networks, which all rely on the Internet Protocol (IP) for data transmission. Shortest path routing protocol such as OSPF or IS-IS control the traffic flow within most of these networks. The network administrator can manage the routing in these networks only by supplying a so-called *routing metric*, which specifies the link lengths (or routing weights) used in the shortest path computation.

The simplicity of this policy offers many advantages in practice. It permits the use of decentralized and distributed routing algorithms, it has very good scaling properties with respect to the network size, and it typically leads to less administrative overhead than connection oriented routing schemes. From the network planning perspective, however, shortest path routing is extremely complicated. As all routing paths depend on the same shortest path metric, it is not possible to configure the end-to-end routing paths for different communication demands individually. The routing can be controlled only indirectly and only as a whole by modifying the routing metric. Additional difficulties arise if each traffic demand must be sent unsplit via a single path – a requirement that is often imposed in practice to simplify network management and to avoid out-of-order packets and other undesired effects of traffic splitting. In this routing variant, the metric must be chosen such that all shortest paths are uniquely determined.

In this paper, we describe the main concepts and techniques that have been developed in [7] to solve dimensioning and routing optimization problems for such networks. We first discuss the problem of deciding if a given path set corresponds to an unsplittable shortest path routing, the fundamen-

*The book version of this extended abstract appeared as *Routing and Capacity Optimization for IP Networks*, Cuvillier Verlag, 2007.

tal properties of such path sets, and the computational complexity of some basic network planning problems for this routing type. Then we describe an integer-linear programming approach to solve such problems in practice, which has been used successfully in the planning of the German national education and research network for several years.

1 Metrics and Routing Paths

Given a digraph $D = (V, A)$ and a set K of directed commodities, an *unsplittable shortest path routing (USPR)* is a set of flow paths P^*_{st}, $(s,t) \in K$, such that there exists a *compatible metric* $\lambda = (\lambda_a) \in \mathbb{Z}^A_+$ with respect to which each P^*_{st} is the unique shortest (s,t)-path. One of the elementary problems in planning shortest path networks is to decide whether a given path set $\mathcal{S}$ is an USPR and, if so, to find a *compatible* routing metric λ.

If there is no upper bound on the length values λ_a, this so-called Inverse Unique Shortest Paths problem can be solved very efficiently with linear programming techniques. We denote with $s(P)$ and $t(P)$ the start and end node of a path P, respectively, and with $\mathcal{P}(s,t)$ the set of all (s,t)-paths in D. It is not difficult to see that there exist an integer-valued metric compatible with $\mathcal{S}$ if and only if the following linear program has a solution [2]:

$$
\begin{aligned}
& \min \lambda_{\max} \\
\sum_{a \in P'} \lambda_a - \sum_{a \in P} \lambda_a \geq 1 \qquad & \forall\, P \in \mathcal{S},\ P' \in \mathcal{P}\big(s(P), t(P)\big) \setminus \{P\} \qquad (1) \\
1 \leq \lambda_a \leq \lambda_{\max} \qquad & \forall\, a \in A,
\end{aligned}
$$

Although this linear program contains exponentially many inequalities of type (1) it can be solved (or proven infeasible) in polynomial time; the separation problem for these inequalities reduces to $|\mathcal{S}|$ many 2-shortest path computations. Its (possibly fractional) optimal solution λ^* easily can be turned into an integer-valued, compatible metric by multiplying all values λ_a with a sufficiently large number and then rounding them to the nearest integer. As shown in [2], this approach yields a metric whose lengths exceed the lengths of the smallest possible integer-valued metric by a factor of at most $\min\left(|V|/2,\ \max\{|P| : P \in \mathcal{S}\}\right)$. For real-world problems, the lengths obtained this way are small enough to easily fit into the data formats of current routing protocols. In theory, however, the problem of finding a compatible routing metric with integer lengths as small as possible or bounded by a given constant is $\mathcal{NP}$-hard [7, 6].

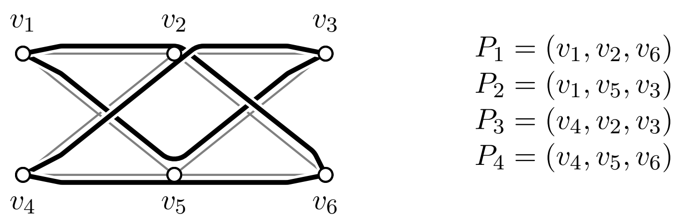

Figure 1: Four paths that cannot occur together in a unique shortest path routing, but any subset of at most three of these paths can.

If the given path set $\mathcal{S}$ is no unsplittable shortest path routing, then the above linear program is infeasible. Using standard greedy techniques, one then can construct from the final dual solution a subset $\mathcal{R}$ of the given paths, such that the paths in $\mathcal{R}$ cannot occur together in any USPR, but any proper subset of the path in $\mathcal{R}$ can. Figure 1 shows such an inclusion-wise minimal conflict set $\mathcal{R}$ consisting of four paths.

These minimal conflict sets are of great practical importance. For every given digraph $D = (V, A)$, the family of all path sets that comprise a valid USPR forms an independence system (or hereditary family), and the circuits of this independence system are exactly these minimal conflict set. Any path set $\mathcal{S}$ that is not an USPR contains at least one of these minimal conflict sets. In a routing optimization framework, it hence is sufficient to ensure that none of these elementary conflicts occurs in the set of chosen routing paths to guarantee the these paths indeed from a valid USPR.

Several types of such elementary conflicts have been studied in the literature. The simplest one is a violation of the so-called *Bellman-* or *subpath-condition* [2, 10]: Two paths P_1 and P_2 can occur together in an USPR only if their (u, v)-subpaths $P_1[u, v]$ and $P_2[u, v]$ – if existent – coincide for all node pairs u, v. All elementary conflicts that involve only two paths are violations of the Bellman-condition. Generalizations of this condition are discussed in [2, 7], another type of necessary conditions has been studied in [13].

However, none of these combinatorial conditions yields a complete combinatorial description of all unsplittable shortest paths routings in a given digraph. In general, the minimal conflict sets can be very complex and arbitrarily large. Given an arbitrary path set $\mathcal{S}$, it is $\mathcal{NP}$-hard to approximate the size $|\mathcal{R}|$ of the smallest conflict set $\mathcal{R} \subseteq \mathcal{S}$ within a factor less than 7/6. The contrary problem of finding the largest subset $\mathcal{R} \subseteq \mathcal{S}$ that still comprises an USPR is computational hard as well. This problem cannot be approximated within a factor less than 8/7, unless $\mathcal{P} = \mathcal{NP}$ [7].

An alternative linear programming approach for finding compatible metrics and an independence system description with similar properties exist also for arc-routing based descriptions of unsplittable shortest path routings.

2 Hardness and Approximability

Network design and routing optimization problems with unsplittable shortest path routing are very difficult – from both the theoretical and the practical point of view. In [7] three basic problem versions are thoroughly analyzed.

In the congestion minimization problem MIN-CON-USPR, we are given a digraph $D = (V, A)$ with fixed arc capacities u_a and a set K of directed commodities with demand values d_{st}, and we seek for an USPR that minimizes the peak congestion (i.e., the maximum flow to capacity ratio over all arcs). This problem corresponds to the task of finding an efficient USPR in an existing network. The peak congestion is a good measure for the service quality network. In general, this problem is $\mathcal{NP}$-hard to approximate within a factor of $\mathcal{O}(|V|^{1-\epsilon})$ for any $\epsilon > 0$, but polynomially approximable within $\min(|A|, |K|)$. This implies that the problem is substantially harder than the corresponding unsplittable flow problem, for example, for which constant factor approximations are known.

Two extremal versions of designing and dimensioning an USPR network are expressed as the fixed charge network design problem FC-USPR and as the capacitated network design problem CAP-USPR, respectively. In both problems we are given a digraph with arc capacities and arc costs and a set of directed commodities with demand values. In FC-USPR the capacities are fix, and the goal is to find a minimum cost subgraph that admits an USPR of the commodities. This problem is $\mathcal{NPO}$-complete even if the underlying graph is an undirected ring or a bidirected cycle. In the capacitated network design problem CAP-USPR, we consider the given arc capacities as basic capacity units and seek a minimum cost installation of integer multiples of these basic capacity units, such that the resulting capacities admit an USPR of the given commodities. This problem cannot be approximated within a factor of $\mathcal{O}(2^{log^{1-\epsilon}|V|})$ in the directed and within a factor of $2 - \epsilon$ in the undirected case, unless $\mathcal{P} = \mathcal{NP}$. For various special cases, however, better approximation algorithms can be derived. For the case where the underlying network is an undirected cycle or a bidirected ring, for example, MIN-CON-USPR and CAP-USPR are approximable within constant factors [5].

The very restricted possibilities to configure the routing make unsplittable shortest path routing problems not only theoretically very hard, they are also an inherent drawback compared to other routing schemes in practice. In certain cases, these restrictions necessarily lead to unbalanced traffic flows with some highly congested links. In [4], we present a class of examples where the minimum congestion that can be obtained with unsplittable shortest path routing exceeds the congestion achievable with multicommodity flow, unsplittable flow, or shortest multi-path routing by a factor of $\Omega(|V|^2)$.

3 Solution Approaches

Traditional planning approaches for shortest path networks use local search, simulated annealing, or Lagrangian relaxation techniques with the lengths of the routing metric as primary decision variables [1, 3, 8, 14, 15, 16, 17, 19, 20]. Basically, these approaches generate or iteratively modify numerous routing metrics and evaluate the resulting routings. The search for promising metrics is guided by the subgradients observed at the solution or other simple, local improvement criteria. The main challenges are to speed up the evaluation of the generated solution candidates to avoid the creation of poor candidates. The major drawbacks of these approaches are that they deliver no or only very weak quality guarantees for the computed solutions and that they perform well only for "easy" problems, where a globally efficient routing metric actually can be found by iterating simple local improvements.

In order to compute provenly optimal solutions, we propose a solution approach that – similar to Bender's decomposition – decomposes the routing subproblem into the two tasks of first finding the optimal end-to-end routing paths and then, secondly, finding a compatible routing metric for these paths.

In the master problem, we consider only the decisions concerning the design and dimensioning of the network and the choice of end-to-end routing paths. This part is solved using combinatorial methods and advanced integer linear programming techniques, which finally guarantees the optimality of the solution by this approach.

The client problem consists in finding a compatible routing metric for the end-to-end paths computed in the master problem. Whenever during the solution of the master problem an integer routing is constructed, we solve the client problem to determine whether the corresponding path set is a valid routing or not. This is done using the linear programming techniques illustrated in Section 1. If the current path set is an USPR, then we have found a incumbent solution for the master problem and the client problem's solution yields a compatible metric. Otherwise the client problem yields a minimal conflict among the current paths, which leads to an inequality that is valid for all USPRs, but violated by the current routing. Adding this inequality to the master problem, we cut off the current invalid solution and re-optimize the master problem. This approach was first described in [10] and refined and adapted to similar routing problems in [7, 12, 18, 21].

To illustrate this approach, consider the Min-Con-USPR problem introduced in the previous section. With $\mathcal{C}$ denoting the family of all (inclusion-wise) minimal path sets that cannot occur in an USPR, this problem can be formulated as in integer programming problem as follows:

$$
\begin{aligned}
& \min L \\
& \sum_{P \in \mathcal{P}(s,t)} x_P = 1 && \forall\, (s,t) \in K \\
& \sum_{(s,t) \in K} \; \sum_{P \in \mathcal{P}(s,t):\, a \in P} d_{st}\, x_P \le u_a L && \forall\, a \in A \\
& \sum_{P \in \mathcal{S}} x_P \le |\mathcal{S}| - 1 && \forall\, \mathcal{S} \in \mathcal{C} \\
& x_P \in \{0,1\} && \forall\, P \in \bigcup_{(s,t) \in K} \mathcal{P}(s,t) \\
& L \ge 0
\end{aligned}
\tag{2}
$$

In principle, our decomposition approach solves this model with a branch-and-bound approach that dynamically separates violated conflict constraints (2) via the client problem. The initial formulation of the master problem would contain only the path variables for each commodity and some of the conflict inequalities (2), for example those corresponding to the Bellman-condition. At each node of the branch-and-bound tree we solve the current LP relaxation, pricing in path variables as needed. Whenever an integer solution $\mathbf{x}$ is found, we solve the linear program for the corresponding Inverse Unique Shortest Paths to find a compatible metric for the corresponding routing. If there exists one, then $\mathbf{x}$ yields the new incumbent solution for the master problem. If there is no compatible metric, we generate a violated conflict inequality (2) from the dual solution of the Inverse Unique Shortest Paths LP, add this inequality to the formulation of the master problem, and proceed with the branch-and-bound algorithm.

From the theoretical point of view, this approach seems not very attractive. For the plain integer programming formulation illustrated above, the integrality gap of the master problem can be arbitrarily large, the separation problem for the conflict inequalities is $\mathcal{NP}$-hard, and the optimal bases of the linear relaxation may necessarily become exponentially large. Nevertheless, carefully implemented this approach works surprisingly well for real-world problems. Our software implementation uses alternatively either a formulation based on path-routing variables or a formulation based on arc-routing

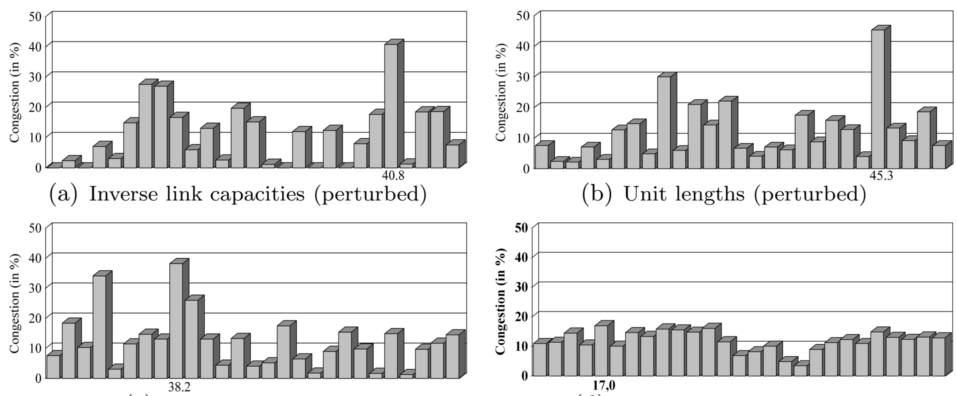

(a) Inverse link capacities (perturbed)
(b) Unit lengths (perturbed)
(c) Geographic lengths
(d) Optimized routing lengths

Figure 2: Link congestion values in G-WiN for several routing metrics.

variables for the master problem. In the branch-and-cut or branch-and-price-and-cut algorithms, we use specially tailored branching and pricing schemes as well as additional problem-specific primal heuristics and strong cutting planes. One type of these cutting planes, for example, exploits the special structure of the precedence constrained knapsacks defined by a link capacity constraint and the Bellman-condition among the paths across that link. In order to handle real-world problems, we also incorporated a very detailed and flexible hardware model, network failure resilience conditions, and various other types of technical and operational constraints into our software.

Numerous small and medium size benchmark problems could be solved optimally with this implementation. Even for large problems, for which optimality was not always achieved, our approach found better solutions than traditional metric-based methods in reasonable computation times. For several years, this software implementation has been used in the planning of the German national education and research networks B-WiN, G-WiN and X-WiN [8, 9, 11].

Figure 2 illustrates the importance of optimizing the routing in practice. It shows the different link loads that would result from the three most commonly used default settings for the routing metric and those resulting from the optimal routing metric for the G-WiN network with capacities and traffic demands of August 2001. The traffic is distributed much more evenly for the optimized metric. The peak congestion is not even half of that for the default settings, which significantly reduces packet delays and loss rates and improves the network's robustness against unforeseen traffic changes and failures.

References

[1] W. Ben-Ameur. Multi-hour design of survivable classical IP networks. *International Journal of Communication Systems*, 15:553–572, 2002.

[2] W. Ben-Ameur and E. Gourdin. Internet routing and related topology issues. *SIAM Journal on Discrete Mathematics*, 17:18–49, 2003.

[3] A. Bley. A Lagrangian approach for integrated network design and routing in IP networks. In *Proceedings of the 1st International Network Optimization Conference (INOC 2003), Paris, France*, pages 107–113, 2003.

[4] A. Bley. On the approximability of the minimum congestion unsplittable shortest path routing problem. In *Proceedings of the 11th Conference on Integer Programming and Combinatorial Optimization (IPCO 2005), Berlin, Germany*, pages 97–110, 2005.

[5] A. Bley. Approximability of unsplittable shortest path routing problems. ZIB Preprint ZR-06-02, Konrad-Zuse-Zentrum für Informationstechnik Berlin, 2006. To appear in *Networks*.

[6] A. Bley. Inapproximability results for the inverse shortest paths problem with integer lengths and unique shortest paths. *Networks*, 50:29–36, 2007.

[7] A. Bley. *Routing and Capacity Optimization for IP Networks*. PhD thesis, Technische Universität Berlin, 2007.

[8] A. Bley, M. Grötschel, and R. Wessäly. Design of broadband virtual private networks: Model and heuristics for the B-WiN. In N. Dean, D.F. Hsu, and R. Ravi, editors, *Robust Communication Networks: Interconnection and Survivability*, volume 53 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 1–16. American Mathematical Society, 1998.

[9] A. Bley and T. Koch. Optimierung des G-WiN. *DFN-Mitteilungen*, 54:13–15, 2000.

[10] A. Bley and T. Koch. Integer programming approaches to access and backbone IP-network planning. ZIB Preprint ZR-02-41, Konrad-Zuse-Zentrum für Informationstechnik Berlin, 2002. Also in *Proceedings of 3rd International Conference on High Performance Scientific Computing, Hanoi, Vietnam*, 2006.

[11] A. Bley and M. Pattloch. Modellierung und Optimierung der X-WiN Plattform. *DFN-Mitteilungen*, 67:4–7, 2005.

[12] N. Bourquia, W. Ben-Ameur, E. Gourdin, and P. Tolla. Optimal shortest path routing for Internet networks. In *Proceedings of the 1st International Network Optimization Conference (INOC 2003), Paris, France*, pages 119–125, 2003.

[13] P. Broström and K. Holmberg. Stronger necessary conditions for the existence of a compatible OSPF metric. Technical Report LiTH-MAT-R-2004-08, Linköping University, May 2004.

[14] L. Buriol, M. Resende, C. Ribeiro, and M. Thorup. A hybrid genetic algorithm for the weight setting problem in OSPF/IS-IS routing. *Networks*, 46:36–56, 2005.

[15] M. Ericsson, M. Resende, and P. Pardalos. A genetic algorithm for the weight setting problem in OSPF routing. *Journal of Combinatorial Optimization*, 6:299–333, 2002.

[16] B. Fortz and M. Thorup. Internet traffic engineering by optimizing OSPF weights. In *Proceedings of the 19th IEEE Infocom 2000, Tel-Aviv, Israel*, pages 519–528, 2000.

[17] B. Fortz and M. Thorup. Increasing Internet capacity using local search. *Computational Optimization and Applications*, 29:13–48, 2004.

[18] K. Holmberg and D. Yuan. Optimization of Internet protocol network design and routing. *Networks*, 43:39–53, 2004.

[19] F. Lin and J. Wang. Minimax open shortest path first routing algorithms in networks suporting the SMDS service. In *Proceedings of the IEEE International Conference on Communications 1993 (ICC'93), Geneva, Suisse*, volume 2, pages 666–670, 1993.

[20] M. Pióro, A. Szentesi, J. Harmatos, and A. Jüttner. On OSPF related network optimization problems. In *8th IFIP Workshop on Performance Modelling and Evaluation of ATM & IP Networks*, pages 70/1–70/14, Ilkley, UK, 2000.

[21] M. Prytz. *On Optimization in Design of Telecommunications Networks with Multicast and Unicast Traffic*. PhD thesis, Royal Institute of Technology, Stockholm, Sweden, 2002.