Optimization of IP Routing by Link Cost Specification

Stefan Köhler, Dirk Staehle
Institute of Computer Science
University of Würzburg, Germany
[koehler, staehle]@informatik.uni-wuerzburg.de

Ute Kohlhaas
Department of Stochastics
Technical University of Aachen
u.kohlhaas@stochastik.rwth-aachen.de

Abstract

Routing is one of the key issues in IP networks. However, few methods exist to optimize routing for a particular network ([1], [2]); most effort has been invested in improving the routing protocol itself ([3], [4], [5]). In this work, a method based on the solution to two mixed integer programs, is presented to specify appropriate values for the link costs for a given network. The obtained link costs can be directly translated into values suitable for the metrics of the currently most important interior gateway routing protocols like IS-IS, EIGRP or OSPF in today's Internet. With this method a homogeneous distribution of traffic in IP-based networks can be achieved without changing the existing routing protocols or hardware.

Keywords: IP, IGP, Link Cost, OSPF, EIGRP, Routing, Optimization

1 Introduction

In this paper we focus attention on the problem of introducing optimal routing methods for IP traffic. IP datagrams are currently used to transmit traffic of all different types of network services. The rapid explosion in the use of the Internet for web browsing, telephony and video services, as well as the more traditional services such as mail and ftp, has resulted in a massive increase in traffic load. With this increase in traffic volume, there has been a concurrent significant increase in congestion due to the lack of network resources and due to the unflexibility of the Interior Gateway Protocols (IGPs). The IGPs do not adapt the routing on the actual traffic situation and because of this, some parts of the network are heavily loaded where other parts have free capacity. To enable Quality of Service in IP networks, efficient routing of IP packets becomes a crucial issue both from the provider and the user point of view.

Most effort concerning routing optimization is focused on improvements and changes of the routing protocol ([3], [4], [5]) itself, only few methods deal with the optimization of the routing of a particular network. Those methods are mainly known from telephone networks with fixed circuit switched connections ([6], [7], [8]). Recently, due to emerging congestion in the growing Internet there has been an increased interest in optimization methods for routing in IP networks. In [1] a method is presented using heuristics to determine optimal link costs in OSPF networks. In [2] network optimization including routing with OSPF is performed with mixed integer programming (MIP).

In contrast to telephony or other connection oriented networks, in IP networks the possibilities to optimize routing are more restricted due to the existing routing protocols. IP is connectionless and therefore a new routing decision is made independently at each router. In a router all IP packets with the same destination are routed on the same path independent of their source (destination based routing). Inside an intranet or within the network of an internet service provider (ISP), the chosen path is the shortest path following a certain metric specified by the routing protocols. The next section gives an overview of the metrics implemented in the most important IP routing protocols. These metrics can be reduced to the sum of link costs. The cost values are traffic independent parameters, which can be set in the router's database. This property is used in the proposed optimization method.

The objective of our method is to determine link costs such that an optimal traffic distribution is achieved. These link costs are computed using two mixed integer programs (MIPs). The first one describes a method to specify optimum paths for given flows in a given network. These paths satisfy the condition that link costs exist such that the received paths correspond to the shortest paths. Link costs corresponding to these paths are then computed using the second program.

In Section 2 the key issues of routing in IP networks are described. Section 3 defines the inputs and the objective of the first optimization problem in detail. Sections 4 and 5 contain the formulation of the two MIPs. Finally, the results for some example networks are presented in Section 6.

2 Routing in IP networks

In IP networks routing is performed according to different unicast IP routing strategies employed by router manufacturers (e.g. EIGRP from CISCO) or the IETF (Internet Engineering Task Force). Fig. 1 gives an overview of the most important ones. Basically, they are classified as:

- Interior Gateway Protocols (IGP) and
- Exterior Gateway Protocols (EGP).

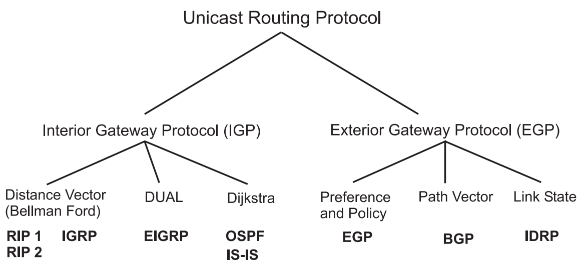

Figure 1: Overview of Unicast Routing Protocols

The IETF defines interior gateway protocols as protocols used for "routing networks that are under a common network administration". The most popular interior routing protocols are the following ones:

- Intermediate System to Intermediate System Routing Protocol (IS-IS)
- Enhanced Internet Gateway Routing Protocol (EIGRP)
- Open Shortest Path First (OSPF) Protocol

More information about the interior routing protocols can be found in [9], [10] or [11].

These protocols have to be configured in their respective network environment by setting certain parameters in the router's database. Router manufacturers publish detailed instruction about router configuration, although they do not provide guidance on specific parameter settings. All interior routing protocols follow the same principle. They define the cost of a link with a protocol dependent metric and determine the shortest path according to these values. The shortest path in EIGRP additionally depends on values which are not link but path dependent such as the reliability of a path.

The world's largest router manufacturer, CISCO, defines the metrics of OSPF and EIGRP as shown in Table 1. In the metric definition, C_i denotes the capacity of link i and accordingly $min_i(C_i)$ is the bottleneck bandwidth of the considered path. R stands for the reliability of the path, K_1 to K_5 are scaling parameters and $delay_i$ is the physical delay of link i. In the default configuration of ther routers the metric of EIGRP is reduced to:

$$M = K_1 \cdot \frac{1}{min_i(C_i)} + K_3 \cdot \sum_i delay_i$$

Protocol	Metric
EIGRP	$\left[K_1 \cdot \frac{1}{min_i(C_i)} + \frac{K_2 \cdot \frac{1}{min_i(C_i)}}{256-\text{load}} + K_3 \cdot \sum_i delay_i\right] \cdot \frac{K_5}{\text{R}+K_4}$ if $K_4 \neq 0$ and $K_5 \neq 0$
OSPF	$M = \sum_i \frac{100{,}000{,}000\text{bps}}{C_i}$

Table 1: Metric of EIGRP and OSPF.

with $K_1 = 1$ and $K_3 = 1$.

The metric of OSPF is not specified in an RFC. Thus, router manufacturers may define their own metric; they have to observe solely the rule that link costs are in the range between 1 and 65535.

As seen in Table 1 the metric for both routing protocols is based on time independent parameters, e.g. physical delay or capacity of links[1]. All parameters can be configured in the router database separately. To adapt the EIGRP metric to the OSPF metric we choose $K_1 = 0$ and $K_3 = 1$. Thus, the optimized solution can be used for EIGRP as well as for OSPF.

The idea behind our optimization is to determine the routing parameters, for a measured end-to-end traffic matrix, with the goal of achieving a homogeneous traffic distribution. The results of the optimization are values which represent link costs. They could easily be transformed to the routing parameters and then entered directly into the router database to achieve an optimized routing.

3 Problem definition

The aim of the routing optimization in this paper is to achieve a routing that balances the utilization of the links in an IP network. A network with N routers is defined by a capacity matrix C of size $N \times N$ which comprises the link capacities c_{ij} of link (i, j) between each pair of routers i and j. If no link exists the entry in the capacity matrix is set to zero. The matrix is not restricted to be symmetrical since, for example, ADSL links are asymmetric.

The optimization is performed for traffic flows, such as the maximum flows during the busy hour. Like the capacities they are defined by an end-to-end traffic matrix F which comprises an entry f_{uv} for each flow $u \rightarrow v$ between nodes u and v. The matrix specifies the volume of the data stream that has to be transmitted from u to v.

Furthermore, a third matrix D is introduced which describes the physical delays d_{ij} of the links. These physical delays restrict the set of routing possibilities between two nodes. Of course, one would not send packets on a

[1]The full EIGRP metric is time dependent because of the factor load, but the reduced metric which is used in practice is time independent.

path which due to its physical delay is many times longer than the shortest path. A intelligent preselection of the possible paths also speeds up the computation time of the problem.

So the objectives of the routing optimization are to

1. minimize the maximum link utilization

2. minimize the average network utilization

3. keep the physical delays within a specified bound

4 Linear programming formulation for path optimization

Solving routing optimization problems by linear/integer optimization is well-known in the literature ([12], [1], [2], [13]). Mixed Integer programs or more general LPs consist of two parts, the objective function and the constraints.

In the case of the IP routing optimization the solutions have to fulfill several constraints. First, a path from source to destination has to be found for each traffic flow. In Section 2 the routing possibilities are restricted to single path routing. So in a possible solution exactly one path has to be defined between each pair of routers, one of which is a source and the other a destination. The second condition for a solution is that the amount of data flowing over a certain interface does not exceed the link's capacity. At least two different modeling approaches can be used to formulate those conditions in terms of constraints for a linear program. The first one is path-oriented, the second one is link flow-oriented. The flow-oriented approach was chosen in this paper, due to the smaller number of variables.

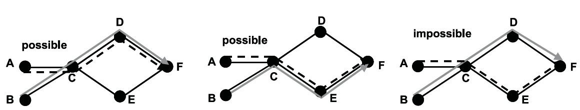

Figure 2: IP Routing restriction

Furthermore, the routing algorithm specified by the applied routing protocol has to be considered. As mentioned above routing protocols like IS-IS, OSPF or EIGRP, currently use shortest path routing, however with different metrics and different algorithms. This implies that the routing between any two routers is identical for all flows unless multipath routing is considered. However, this possibility is considered to be inappropriate in this optimization. In Fig. 2 a simple example for this restriction is shown. The shortest path - independent of the used metric - between router C and router F is either over D or over E. Consequently, both flows A-F and B-F are routed first over C and then either over D or over E, but not one over D and one over E.

4.1 Variables

Before formulating the different constraints and the objective function the variables have to be specified. Since a flow-oriented approach is used for each flow with $f_{uv} \geq 0$ and for each link with $c_{ij} > 0$ a boolean variable x_{ij}^{uv} is introduced. This variable is set to one if flow $u \rightarrow v$ is routed over link (i, j) and to zero, otherwise. The variable t is an upper bound for the utilization of all links.

4.2 Objective function

The objective function quantifies the aim of the routing optimization. In Section 3 it was stated that we wish to achieve a traffic distribution as homogeneous as possible in the entire network; this means that ideally all links should be utilized equally at a level as low as possible. This aim is obtained by minimizing the maximum link utilization as far as possible. But once this maximum utilization is found for a certain link, the traffic of all other links shall be minimized without increasing the obtained maximum value. Therefore, the objective function comprises two additive parts:

$$a_t \cdot t + \sum_{ij} \sum_{uv} \frac{f_{uv} x_{ij}^{uv}}{c_{ij}}. \tag{1}$$

In the first part the maximum link utilization is minimized. As we will see in Eqn. (7) a constraint is formulated that forces all utilizations below the value of t. The second part reduces the average link utilization. The actual term in the objective is the sum of all link utilizations, which is proportional to the average network link utilization.

The parameter a_t defines the importance of a small maximum utilization versus the importance of a small average utilization. If it is set sufficiently large the prime aim of reducing the maximum utilization by directing traffic onto less utilized links is achieved.

4.3 Constraints

The constraints are used to define the set of possible solutions. They can be subdivided into

- Transport Constraints: provide a loopfree path between all origin destination pairs

- Capacity Constraints: keep the link utilizations under a certain limit

- IP Routing Constraints: provide IP conforming routing

- Delay Constraints: avoid paths that are too long

4.3.1 Transport Constraints

The transport constraints guarantee that for each origin to destination flow $u \rightarrow v$ exactly one loopfree path is specified by the resulting values of the variables x_{ij}^{uv}. This is obtained with four types of constraints for each flow:

1. Only one link leading out of a router i may carry traffic of flow $u \rightarrow v$. For all flows $u \rightarrow v$ and all routers i add the inequalities:

$$\sum_{j=1,\ c_{ij}>0}^{N} x_{ij}^{uv} \leq 1 \tag{2}$$

2. Exactly one link (u, i) from router u to another router i has to carry flow $u \rightarrow v$:

$$\sum_{i=1,\ c_{ui}>0}^{N} x_{ui}^{uv} - \sum_{i=1,\ c_{iu}>0}^{N} x_{iu}^{uv} \geq 1 \tag{3}$$

3. For all flows $u \rightarrow v$ and for each router $i \notin \{u, v\}$ the sum of incoming links used by flow $u \rightarrow v$ equals the sum of outgoing links used by $u \rightarrow v$. This is realized by adding the following constraints:

$$\sum_{j=1,\ c_{ij}>0}^{N} x_{ij}^{uv} - \sum_{j=1,\ c_{ji}>0}^{N} x_{ji}^{uv} = 0 \tag{4}$$

4. Exactly one link incident on router v has to carry the traffic of flow $u \rightarrow v$:

$$\sum_{i=1,\ c_{iv}>0}^{N} x_{iv}^{uv} - \sum_{i=1,\ c_{vi}>0}^{N} x_{vi}^{uv} \geq 1 \tag{5}$$

These constraints, together with the second term in the objective function (see Figure 3), imply that exactly one loopfree path from u to v is specified in each solution.

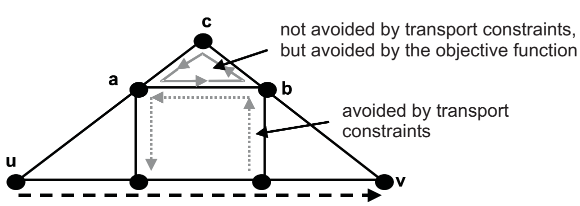

Figure 3: Loops despite Transport Constraints

4.3.2 Capacity Constraints

The capacity constraints guarantee that the traffic over a link does not exceed certain limits. Two constraints are required for each link (i, j). The first one ensures that the link utilization stays below a fixed limit given by the parameter a_c. This is achieved by the following constraint:

$$\sum_{uv} x_{ij}^{uv} f_{uv} \leq a_c c_{ij} \quad \text{, for all links } (i, j). \tag{6}$$

Within IP networks it is possible that links are offered more traffic than they are able to handle. Therefore a_c is not restricted to values between zero and one. Nevertheless, by default a_c is set to one.

Furthermore the value of the variable t is another upper bound for the utilization of the links. While a_c defines a fixed upper bound, the percentage value t is a variable upper bound which is minimized in the objective function. This is realized by the following constraint:

$$\sum_{uv} x_{ij}^{uv} f_{uv} \leq \frac{t}{100} c_{ij} \quad \text{, for all links } (i, j). \tag{7}$$

4.3.3 Specific Routing Constraints in IP Networks

Up to now the constraints specify that each solution found provides a path for each flow such that the given link utilizations are not exceeded. The routing in an IP network is more restricted due to the functionality of the routing protocols. The crucial point here is that all IP packets are always routed on the shortest path between their origin and their destination. The shortest path is determined following a certain metric for each link. This metric depends on the used routing protocol.

As a consequence all flows with routers i and j in the same sequence on their path have to be routed in the same way between i and j. This restriction is further illustrated in Fig. 2. The IP conforming routing is achieved by adding the following inequalities to the MIP:

1. For all flows $u \rightarrow v$, routers $i \notin \{u, v\}$, and links (s, t) add:

$$x_{ui}^{uv} + x_{st}^{iv} - x_{st}^{uv} \leq 1 \tag{8}$$

2. For all flows $u \rightarrow v$, routers $j \notin \{u, v\}$, and links (s, t) add:

$$x_{jv}^{uv} + x_{st}^{uj} - x_{st}^{uv} \leq 1 \tag{9}$$

Eqn. (8) can be interpreted in the following way:

$$x_{ui}^{uv} = 1 \Rightarrow x_{st}^{iv} \leq x_{st}^{uv}$$

If the first hop of flow $u \rightarrow v$ leads to router i, all links (s, t) used by $i \rightarrow v$ have to be used by $u \rightarrow v$ as well. So $u \rightarrow v$ and $i \rightarrow v$ are routed on the same path from router i to router v.

Eqn. (9) can be interpreted similarly. If the last hop of $u \rightarrow v$ goes out of router j, all links (s, t) used by flow $u \rightarrow j$ have to be used by $u \rightarrow v$ as well. So $u \rightarrow v$ and $u \rightarrow j$ are routed on the same path from router u to router j.

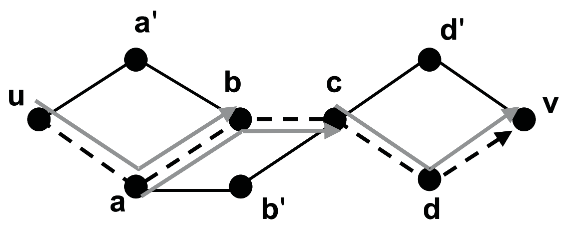

Figure 4: IP Constraints

If we conceive an iterated application of these equations we notice that for each two routers i and j the routing between them is identical for all flows through these nodes. And so the solutions correspond to IP routing implementations.

In Fig. 4 the essence of the IP Routing constraints is illustrated. Assume that flows $u \to b$, $a \to c$, and $c \to v$ are routed as indicated by the grey arrows. Due to Eqns. (8) and (9) the $u \to v$ has to be routed along the dashed black arrow.

If we look at flow $u \to c$ first, two alternatives to the path along the dashed black arrow exist, either over a' or over b'. Assume $u \to c$ runs over b'. Then link (u, a) is the first link of this flow and due to Eqn. (8) flows $u \to c$ and $a \to c$ have to use the same path between a and c. For the other alternative path over a' the case is analogous by (Eqn. 9). Thus, flow $u \to c$ is routed along the dashed black arrow.

Again, for flow $b \to v$ an alternative path over d' exists but it can not be used due to Eqn. (8). With both routing constraints applied repeatedly, $u \to v$ has to use the same path as $u \to c$ between u and c and also the same path as $b \to v$ between b and v. Altogether, $u \to v$ has to use the path indicated by the dashed black arrow.

4.3.4 Physical Delay Constraints

With the constraints and the objective function described above a shortest path conforming routing with evenly distributed traffic is achieved. Thus, only the third aim to keep the physical delays bounded is not yet fulfilled. The range of the possible physical delays is defined by the parameter a_r and the lowest possible physical delay for a flow d_{min}^{uv}. Let P^{uv} be the set of all loopfree paths from u to v and let d_p be the sum of the physical delays of the interfaces belonging to path p. Then the minimum physical delay d_{min}^{uv} is the delay of the path p from u to v with the smallest delay:

$$d_{min}^{uv} = \min_{p \in P^{uv}} (d_p). \tag{10}$$

This value can be determined by computing the physical delays of all possible loopfree paths from u to v.

With these two values a_r and d_{min}^{uv}, a constraint can be specified for each flow $u \to v$ that keeps the physical delay below or equal to $a_r d_{min}^{uv}$:

$$\sum_{ij} x_{ij}^{uv} d_{ij} \leq a_r d_{min}^{uv}. \tag{11}$$

4.4 Reduction of the complexity of the linear problem

In the above formulated problem the number of variables is of magnitude $N^2 M^2$ where N is the number of routers and M the number of interfaces. The number of constraints is of magnitude $N^3 M^2$, determined by the routing constraints. Accordingly, the computational effort for solving the problem increases considerably for larger networks. However, if one takes a closer look at the majority of intranets the routing decision involves a few competing paths only. Therefore, many variables can be presolved by reducing the problem to relevant paths.

The physical delay of a path defines whether it must be considered or not. A path $p \in P^{uv}$ is not relevant if its physical delay d_p is greater than a_r times the minimal physical delay d_{min}^{uv}. Then $\hat{P}^{uv}$ denotes the set of all relevant paths:

$$\hat{P}^{uv} = \{p \in P^{uv} \mid d_p \leq a_r d_{min}^{uv}\} \tag{12}$$

With the set of all relevant paths the value of the following variables can be determined:

$$x_{ij}^{uv} = \begin{cases} 0 & \text{if } \forall p \in \hat{P}^{uv} : \delta_{ij}(p) = 0 \\ 1 & \text{if } \forall p \in \hat{P}^{uv} : \delta_{ij}(p) = 1 \end{cases} \tag{13}$$

where $\delta_{ij}(p)$ equals one if p uses link (i, j) and zero otherwise. For all other variables x_{ij}^{uv} link (i, j) is used by some but not all of the relevant paths in $\hat{P}^{uv}$ and therefore their value is not specified by the presolver. The fixed variable values are substituted into the MIP and the constraints which are always satisfied by these values are omitted.

Another simplification reduces the complexity of the problem itself. The treated example networks revealed that simultaneously minimizing the maximum and average utilization increases the problem's complexity considerably. However, it is possible to omit the variable t from the objective function and also the constraints of Eqn. (7) which force the maximum utilization below t. The reduction of the maximum utilization is now obtained by reducing a_c iteratively. A first solution is found with a_c set to one or an even greater value. With this solution the maximum link utilization is computed and the MIP can be solved again with a_c set to a value smaller than the received maximum utilization. If this procedure is repeated until either the problem is identified as infeasible or the solver can not handle the problem due to its complexity, an optimal or at least a good solution is found. In Section 6 results with and without this simplification are compared.

5 Specification of the link costs

In the previous section a routing scenario was proposed which aims to minimize the maximum and average link utilizations and hence the IP packet transmission delays. Additionally, the approach is suitable for shortest path routing. According to the optimized paths, the appropriate interface costs now have to be specified such that the obtained routing corresponds to the shortest path routing.

A second MIP is formulated to compute these link costs. The objective function of the second MIP minimizes the obtained link costs which is necessary with respect to the OSPF metric where an upper bound for the link costs is given. The more important part of the MIP are the constraints. These restrict the solution space to contain only solutions where the shortest path routing is identical with the optimized routing. The variable m_{ij} represents the cost of interface (i, j). The shortest path from u to v is the path with the least cost. The path cost is the sum of the cost values of the single interfaces. For flow $u \to v$ the cost m^{uv} is given by

$$m^{uv} = \sum_{ij} m_{ij} x_{ij}^{uv}, \tag{14}$$

where x_{ij}^{uv} describe the paths of the previous routing optimization. m^{uv} has to be the minimum cost of all possible paths from u to v. Therefore, for each loopfree path p other than the optimized path p_{opt}^{uv}, a constraint has to be specified that forces the cost of p_{opt}^{uv} to be the smaller one:

$$\sum_{ij} x_{ij}^{uv} m_{ij} < \sum_{ij, \delta_{ij}(p)=1} m_{ij}. \tag{15}$$

To include all such constraints for all paths would lead to a very complex problem for larger networks. However, not all of these constraints are required to obtain a suitable solution. For a flow $u \to v$ the mandatory constraints are identified by the following algorithm:

Consider each neighbor router i of u:

if p_{opt}^{uv} and p_{opt}^{iv} have no common node except v, add the constraint:

$$\sum_{st} x_{st}^{uv} m_{st} < m_{ui} + \sum_{st} x_{st}^{iv} m_{st} \tag{16}$$

else add no constraint.

The algorithm is explained using the example of Figure 5. Every path from u to v runs over a neighbor router of u. In the example the three neighbors of u are i, j, and additionally the first router of p_{opt}^{uv}. Three types of paths exist between a router n, which is adjacent to u, and v:

- optimized path p_{opt}^{nv} without a node in common with p_{opt}^{uv}; in the example the solid grey line between i and v
- optimized path p_{opt}^{nv} with a node in common with p_{opt}^{uv}; in the example the solid grey line between j and v
- a non-optimal path from n to v; in the example the dashed grey line

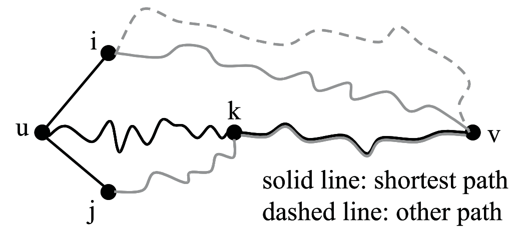

Figure 5: Example for link cost algorithm

Following the above algorithm a constraint is necessary only for the path including p_{opt}^{iv}. The constraints concerning the other paths are already fulfilled by the constraints of other flows.

In the case of the path p_g, indicated by the dashed grey line between i and v, the constraint

$$\sum_{st} x_{st}^{iv} m_{st} < \sum_{st, \delta_{st}(p_g)=1} m_{st}$$

already exists when flow $i \to v$ is considered. This constraint together with Eqn. (16) implies that the cost of the path using p_g is larger than the cost of p_{opt}^{uv}.

Considering the path between j and v, the part between k and v is common with p_{opt}^{uv}. According to the routing principles we know that the shortest path from u to k is routed along the black line. For this flow $u \to k$ the constraint exists that the path along the black line p_{opt}^{uk} has to be shorter than the path over j and then along the grey line. But this constraint implies that p_{opt}^{uv} is shorter than the path over j and then along the grey line.

6 Results

In this section results of the path optimization and the link cost specification are presented. They are demonstrated using two example networks of different size. These networks are shown in Fig. 6. The network on the left with only six routers was chosen because its size allows us to depict the resulting paths. The networks with fourteen routers were selected because of the more complex structure which makes many different path choices possible. Hence, the routing optimization is rather complex.

For all two example networks a flow exists between each pair of nodes. Consequently the flow matrix is dense with the exception of the main diagonal. For the six router network it is given in Table 2. Table 3 shows the link capacities of this network.

The hop counts were taken as physical delays for all three networks, that is d_{ij} equals one for each link (i, j).

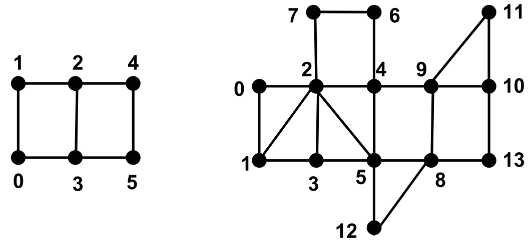

Figure 6: Example networks with six and fourteen routers

	0	1	2	3	4	5
0	0	5	6	7	4	7
1	4	0	6	4	8	5
2	6	8	0	4	6	4
3	8	9	4	0	5	6
4	5	7	8	5	0	3
5	5	7	6	8	4	0

Table 2: Flow matrix of the 6-router network

Furthermore, in each optimization run the parameter a_r was set at three. Thus, for the 14-router network 4714 of the 8008 variables were presolved and set to zero. Furthermore, the number of constraints was reduced to less than the half of the possible number.

	0	1	2	3	4	5
0	0	120	0	100	0	0
1	110	0	80	0	0	0
2	0	90	0	100	0	70
3	130	0	80	0	90	0
4	0	0	0	70	0	120
5	0	0	75	0	125	0

Table 3: Capacity Matrix of the 6-router network

Optimizations with several parameter settings were performed for all example networks. First, the optimization was performed without an effective upper bound for the link utilizations. This was achieved by omitting the constraints keeping the link utilizations below t. Also the parameter a_c was set to a value of 10. Thus, the load of the links was permitted to be ten times greater than their capacities.

The resulting routing scenario for the 6-router network is depicted in Fig. 7. The link with the highest utilization of 42.9% is $(4, 3)$. The value received for the minimized average utilization is 22.4%. As described above the value of parameter a_c was then repeatedly decreased below the previous obtained maximum utilization. The settings for a_c were 0.4, 0.375, and 0.36. The result for $a_c = 0.36$ is identical with the result of the default optimization.

The parameter a_t was set to 1000 in the default optimization. With this setting the maximum link utilization is weighted more strongly than the average utilization. The paths that have changed by minimizing the maximum utilization are shown in Fig. 8. The left figure contains the paths for the minimum average utilization and the right figure the corresponding paths for the minimized maximum utilization. The link utilizations which differ in the two graphs are also shown.

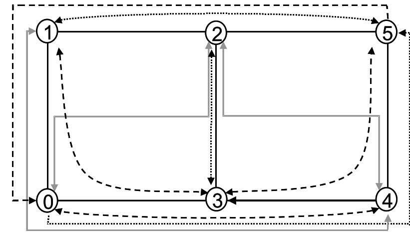

Figure 7: Optimized routing for minimized average utilization

The most utilized link in the right figure is still $(4, 3)$, however it's utilization was decreased from 42.9% to 35.7%. As a compensation, the average utilization increased by only 0.3% to a value of 22.7%. We see that flow $5 \rightarrow 0$ was removed from link $(4, 3)$ and is now routed over routers 2 and 1. Consequently, the path of flow $2 \rightarrow 0$ was changed as well. With these changes only, link $(3, 4)$ would have been utilized with 40%, so flow $2 \rightarrow 4$ was directed over node 5 instead of node 3.

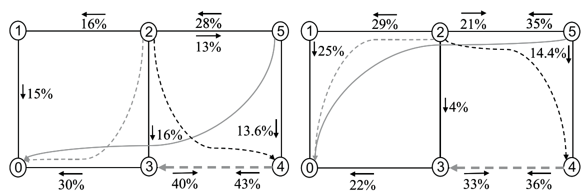

Figure 8: Changed paths through minimizing the maximum utilization

Link costs as small as possible were specified with the second MIP such that the shortest path routing equalizes the routing defined by the first MIP. In Fig. 9 the network with the received costs is depicted. They are shown in the boxes next to the links.

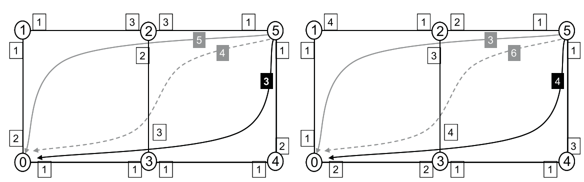

Figure 9: Link costs for minimized average and maximum utilization

Additionally, the pictures show the three possible paths for flow $5 \rightarrow 0$ represented with three arrows. The boxes next to the arrows contain the total path cost. We can see that in the left figure the minimum cost is 3 for the black path. The two other possible paths either run over routers 2 and 1 or over routers 2 and 3. The associated path costs are 5 and 4, respectively. Thus, the shortest path is equal to the desired path. This holds also for the other flows.

In the right figure showing the results for the minimized maximum utilization, the smallest cost of flow $5 \to 0$ is again 3, here for the solid grey path. The alternative paths over nodes 4 and 3 or over nodes 2 and 3 have higher costs of 4 and 6, respectively. Again, the shortest path corresponds to the desired path.

For the other example network the procedure was similar. First, the average utilization was minimized without restriction for the maximum utilization. Then the maximum utilization was decreased iteratively. Finally, the optimization was performed with the objective of obtaining the minimal possible maximum utilization. The maximum utilization was reduced to 38.8% in comparison to 52.8% with results from default (shortest path respectively shortest hop) routing.

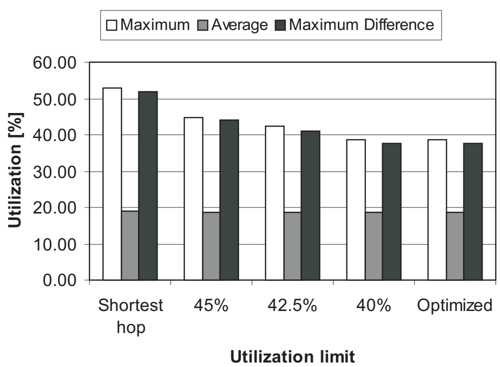

Figure 10: Utilization changes for the 14-router network

In Fig. 10 the results for the path optimizations with the chosen upper bound for the link utilization can be seen. The white bars show the maximum utilization and the grey bars the average utilization. The maximum utilization is decreased conspicuously, whereas the average utilization stays almost unchanged. The black bars show the difference between the utilization of the most and the least utilized links. This difference indicates whether the traffic is evenly distributed over the network or not. The

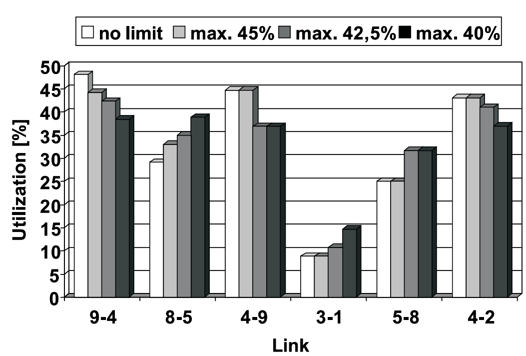

Figure 11: Selected interfaces of the 14-router network

graphs show that the traffic distribution, as the routing is optimized, becomes more homogeneous with a stricter upper bound for the utilization of the links.

Nevertheless, since the greatest utilization was decreased, some utilizations have to be increased as well. This can be seen in Fig. 11. The utilization of selected links is shown for the different upper bounds. We can see that the utilizations of the links that are initially highly utilized are decreasing with a reduction of the upper bound, for example links $(9, 4)$, $(4, 9)$ and $(4, 2)$. And as a compensation the utilizations of links $(8, 5)$, $(5, 8)$ and $(3, 1)$ are increasing.

7 Conclusions

In the last section we presented results for networks of different sizes. If we look at the structure of the fourteen-router network one has to notice that it is quite complex. Between most pairs of nodes a lot of possible paths with the same or a similar physical delay exist. Of course, actual intranets mostly comprise more that fourteen routers. However, in many cases a main part of the network can be identified which is relevant for routing optimization and consists of a number of routers for which the MIPs can be solved. An example for such a scenario is the $AT\&T$ network. Here more than 100 routers build the original network (Fig. 12 left), but for optimization, only the core out of 25 routers has to be considered (Fig. 12 right). The routing optimization was also performed for this network and arbitrary flows were assumed. Sometimes it was not possible to obtain an the optimal solution for special traffic matrices, but again the maximum utilization was decreased, and consequently the traffic distribution was homogenized. The results are too numerous to present them here.

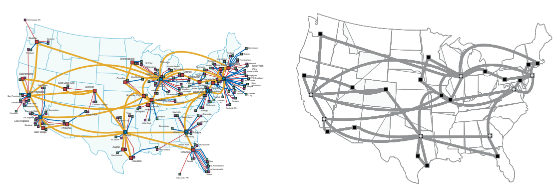

Figure 12: $AT\&T$ IP Backbone Network [14]

The major issue of this work was to find a possibility of optimizing the routing for existing routing protocols. It was shown that the results obtained by the two MIPs are applicable in the currently most important routing protocols IS-IS, OSPF and EIGRP.

In general, our routing optimization works well for instantaneous or peak traffic flows. However, the real internet traffic is not static but varies over time. The quality of our result, applied to varying traffic flows, still has to be investigated. This leads to the question how to compare two different routing possibilities and also how to evaluate the quality of a routing decision. In our work the results were compared by the average and maximum link utilizations. However, the effects on actual transmission delays packets in the network experience, are not compared. As well, the influence of these delays on TCP connections would be an interesting point.

Acknowledgments

The authors would like to thank Prof. Tran-Gia and Prof. Berry for the fruitful discussions.

References

[1] M. Thorup, B. Fortz, "Internet Traffic Engineering by Optimizing OSPF Weights," in *Proc. IEEE Infocom 2000, Tel Aviv, Israel*, 2000.

[2] A. Bley, M. Grötschel, R. Wessäly, "Design of Broadband Virtual Private Networks: Model and Heuristics for the B-WiN," Tech. Rep. SC-98-13, Konrad-Zuse-Zentrum für Informationstechnik Berlin (ZIB), 1998.

[3] E. Crawley et al , "A Framework for QoS-based Routing in the Internet," IETF RFC 2386, August 1988.

[4] A. Orda, D. Lorenz, "Qos routing in networks with uncertain parameters," in *Proc. IEEE Infocom 1998, San Francisco, USA*, 1998.

[5] G. Apostolopoulos et al, "QoS Routing Mechanisms and OSPF Extensions." IETF RFC 2676, August 1999.

[6] R.P. Murray, G.R. Ash, R.H. Cardwell, "Design and optimisation of networks with dynamic routing," in *Bell System Technical Journal*, vol. 60, no. 8, October 1981.

[7] K.R. Krishnan, G.R. Ash, A.H. Kafker, "Servicing and real time control of networks with dynamic routing," in *Bell System Technical Journal*, vol. 60, no. 8, October 1981.

[8] R. Gallager, D. Bertsekas, *Data Networks*, Prentice Hall, 2^{nd} edition, 1992.

[9] S. Köhler, R. Harris, "Possibilities for qos in existing internet routing protocols," Tech. Rep. 228, University of Würzburg, May 1999.

[10] J. Moy, "OSPF Version 2," IETF RFC 2328, April 1988.

[11] C. Huitema, *Routing in the Internet*, Prentice Hall, 2^{nd} edition, 1992.

[12] F. Barahona, "Network design using cut inequalities," in *SIAM J. on Optimization*, vol. 6, no. 3, 1996.

[13] T. Magnanti et al , "Modeling and solving the two-facility capacity network loading problem," in *Operations Research*, vol. 43, no. 1, Jan-Feb 1995.

[14] "AT&T IP backbone network, 3Q2001," Available at http://www.ipservices.att.com/backbone/bbone.pdf.