13

Network Engineering

In this chapter, the influence of network engineering on the efficiency and quality of service of a network is investigated. As argued in Section 11.1, capacity expansion is the most frequent network engineering task of an INSP. Therefore we focus on capacity expansion. We start by evaluating the influence of capacity on the performance of different QoS systems in Section 13.1. Different capacity expansion strategies are evaluated in Section 13.2. We base this analysis on the results of the previous chapter by incorporating the previously found best traffic engineering algorithms into our analyses. The mutual influence of capacity expansion and traffic engineering is also analysed in that section. Finally, in Section 13.3, we investigate the effect of elastic traffic on traffic matrices in the context of capacity expansions.

13.1 Quality of Service Systems and Network Engineering*

Capacity expansion (CE) deals with increasing the network capacity of a network. Internet traffic volumes are growing very fast. Numbers presented, for example in Odlyzko (2003) indicate that the traffic volume is increasing by 70 to 150% per year. Therefore, the capacity of a network has to be adapted regularly to the growing needs.

The effect of capacity expansion on the performance of different QoS systems is analysed by the following experiment. It is based on the packet simulations that are described in detail in Chapter 8, especially Section 8.5. We repeat the experiments of Section 8.5[1] with varying levels of capacity (bandwidth and buffer), starting with half the capacity used in Section 8.5; the capacity multiplicator is depicted on the x axis of the following graphs.

The utility of the accepted flows is used as the performance measure of the overall network performance; see Chapter 8 for details. For the four different types of traffic of Chapter 8, it is depicted in Figures 13.1 and 13.2. Please note, that the maximum possible utility is 1.0.

As one can see for all QoS systems, the overall utility obviously increases with the amount of available capacity. There are, however, great differences between the different QoS systems.

* Reproduced with permission from Oliver Heckmann and Ralf Steinmetz, Capacity Expansion for MPLs Networks, Proceedings of INOC 2005.

[1] DFN topology, traffic mix A.

The Competitive Internet Service Provider: Network Architecture, Interconnection, Traffic Engineering and Network Design
Oliver Heckmann © 2006 John Wiley & Sons, Ltd. ISBN: 0-470-01293-5

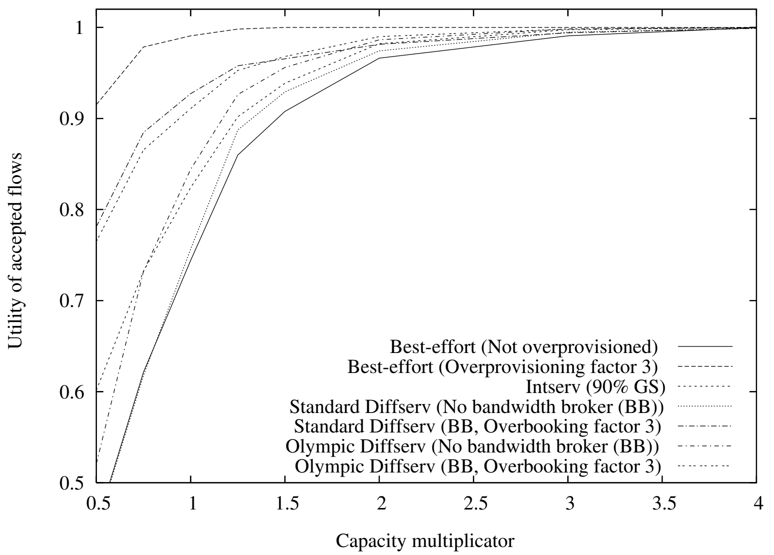

(a) Short-lived TCP Traffic

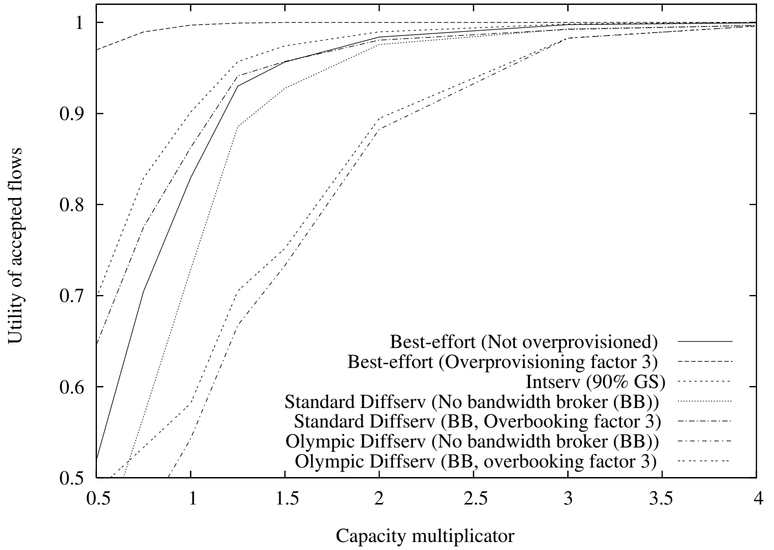

(b) Long-Lived TCP Traffic

Figure 13.1 Utility of the TCP Flows for Different QoS Systems as a Function of the Capacity

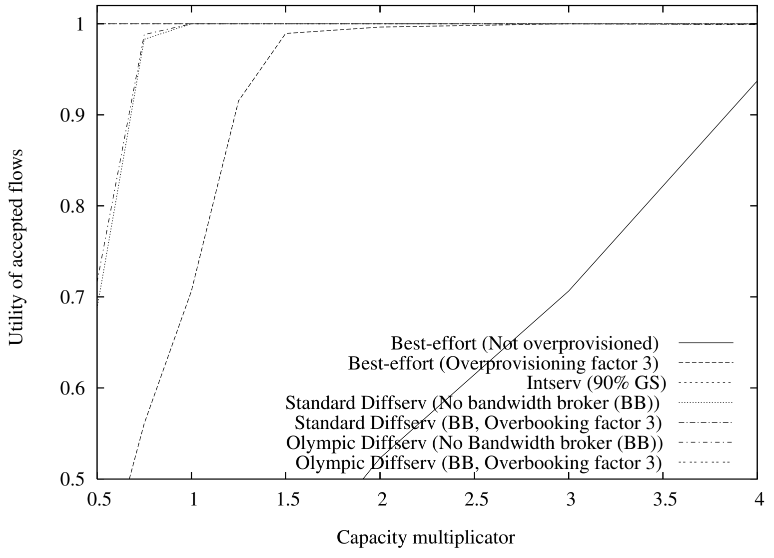

(a) VBR Traffic

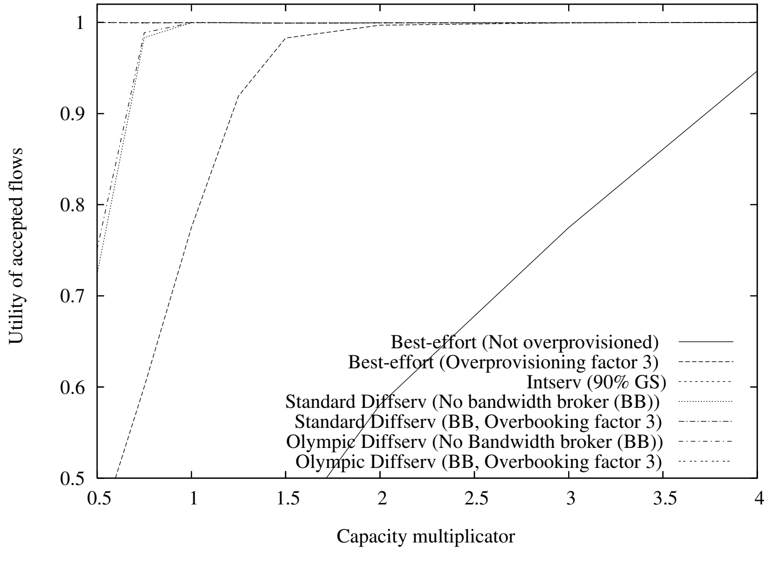

(b) CBR Traffic

Figure 13.2 Utility of the Accepted Inelastic Flows for Different QoS Systems as a Function of the Capacity

The utility of all QoS systems without admission control (Best-effort, Diffserv without Bandwidth Broker) breaks down quickly if the capacity of the network is too low. For the QoS systems with admission control, the utility of the accepted CBR and VBR flows does not break down as these flows are protected by the QoS system. However, the number of rejected customers increases: For Intserv and CBR traffic, the rejection rate is 48.1% for a capacity factor of 1 and 3.4% for one of 4.

For the experiment, the CBR and VBR flows were assumed to be of higher importance than the short-lived TCP flows. The long-lived TCP flows were assumed to resemble peer-to-peer or similar traffic with the lowest importance. The strongest differentiation between the flows is visible for the Olympic Diffserv QoS systems, where the performance of the long-lived TCP flows breaks down long before that of the CBR/VBR flows.

Figures 13.1 and 13.2 also show that in order to support CBR/VBR (multimedia) flows with a plain best-effort architecture, sufficient capacity is even more important than for the other QoS systems.

On the other side, the experiment also shows that if capacity is available in abundance, there is no significant difference between the various QoS systems.

13.2 Capacity Expansion*

Because Internet traffic is continuously increasing, capacity expansion is extremely important to maintain QoS. While QoS systems differ in their ability to maintain a high QoS in the face of scarce capacities, the performance of all systems breaks down if the capacity is too low, as was shown in the previous section. On the other hand, if capacity is expanded too early, the additional capacity remains largely unused for some time and the efficiency of the network suffers. We found that most INSPs use rules of thumb as link capacity expansion strategy in a continuous planning process. The typical rule of thumb is to trigger the expansion of a link once a certain utilisation threshold is exceeded.

In this section, the capacity expansion problem is modelled as an optimisation problem. The mutual influence of capacity expansion and traffic engineering is also considered. Different strategies are compared with the mentioned rule of thumb and some variations of it in a series of experiments in order to analyse the influence of the strategies and to identify the best strategy.

Capacity expansion is based on predictions of future traffic that are typically uncertain – contrary to the traffic engineering experiments in the previous chapter that is based on actual (measurable) traffic. Therefore, and contrary to almost all of the related works (see Section 11.1), we now also consider the uncertainty involved in predicting future traffic demand in our experiments.

13.2.1 Capacity Expansion Process

The typical capacity expansion process is depicted in Figure 13.3. Multiple periods t are investigated; the traffic changes from period to period. In every period, the traffic is routed through the network. If traffic engineering is used, the routing can change from period to period, adapting to changed capacities and flow sizes.

* Reproduced with permission from Oliver Heckmann and Ralf Steinmetz, Capacity Expansion for MPLs Networks, Proceedings of INOC 2005.

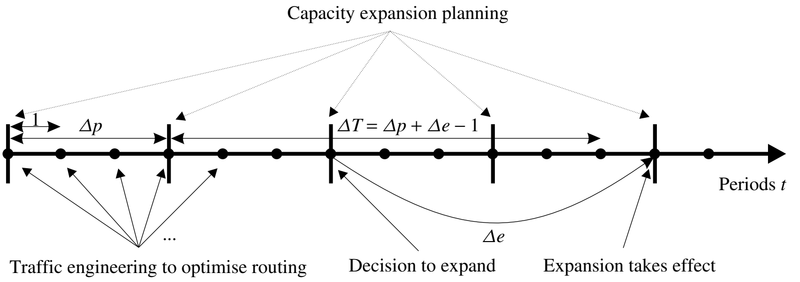

Figure 13.3 Capacity Expansion Process

An INSP decides on capacity expansions every Δp periods (e.g. once per quarter). It takes Δe periods from the decision to expand the capacity of a link until the expansion actually takes effect and the capacity is increased. Δe can be larger or smaller than Δp.

At the point in time when the decision is made, the link utilisations of the current period are known exactly as they can be measured. We assume that there are predictions available for at least the next ΔT periods. The predictions, however, are subject to uncertainty. ΔT is called the *planning horizon*. ΔT has to be at least $\Delta p + \Delta e - 1$ so that all periods are covered by the capacity expansion process.

13.2.2 Capacity Expansion Strategies

To describe the capacity expansion strategies, the same modelling parameters and simulation environment are used as for the traffic engineering strategies in the previous chapter (see Section 12.2).

The traffic volume is increasing in the long run, so the link capacities have to be expanded sooner or later. We assume that the capacity expansion of a link results in a doubling of the available link bandwidth. This is common practice at INSPs and represents adding either a second line card for one link to a router doubling the available bandwidth or – if two line cards are already present – switching to the next higher SONET/SDH data rate, which also results in effectively doubling the bandwidth (see Table 4.2).

The topology is modelled as a directed graph to be consistent with the models of the previous chapter; however, in a network the connection between two routers typically has the same bandwidth in both directions. Therefore, we assume that two opposing links between the same node pair always have the same capacity.

There are two types of *costs* involved. (a) The costs for the capacity expansion and (b) the increased congestion if capacity is expanded too late.

Assuming that the Internet traffic continues growing in the long run, the *costs for capacity expansion* are not the absolute costs for the equipment, as that equipment has to be bought anyway sooner or later. Also, the question answered by the capacity expansion strategies in the long run is not *whether* to expand but rather *when* to expand. The true costs of the capacity expansion in period t_a are the opportunity costs representing the missed earnings that could be realised if the expansion was delayed until a later period t_b. These opportunity costs consist of the interest for the invested money plus the savings if the prices for the equipment (line cards, leased lines) falls until period t_b[2]. We accumulate all these costs with the interest cost factor p^i and assume that they are proportional to the capacity.

Obviously, a capacity expansion cannot be delayed forever because the congestion would rise to an unbearable level. In the previous chapter, the congestion was modelled with a *congestion cost function* $p^x(u_l)$ that increases exponentially with the utilisation u_l of a link l (see Section 12.1 and Figure 12.1). The same approach is used in this chapter to model the fictive costs resulting from the congestion of the network. These costs result from the decreased QoS the network offers and the risk of, for example, losing profit and customers as a result of that. This cost term can be hard to quantify exactly in reality as it depends on many variables and on market conditions. The more important the network QoS is for a provider, the higher this cost factor will be.

While this second cost factor is influenced by the network QoS, the first cost factor leads directly to monetary expenses and therefore directly influences the overall network efficiency. Solving the capacity expansion problem means finding a compromise between these two goals. Therefore, we introduce a parameter c that measures how these two goals are weighted with each other. c measures the ratio between the interest cost factor and the congestion cost function. c describes where along the optimal performance boundary (see Figure 1.1) a provider wants to operate. Because the congestion costs depend on the utilisation, we arbitrarily define a reference point for a utilisation of 60% to quantify c

$$c = \frac{p^i}{p^x(60\%)} \tag{13.1}$$

In the experiments below, we evaluate the influence of c on the results.

Throughout this section, we assume that the traffic volume rates r_{tf} are influenced by the capacity expansion itself. This is the typical approach in almost all traffic engineering and network design problems (see Chapter 11). In Section 13.3, we drop this assumption and analyse the effect of elastic traffic – for example, TCP – on capacity expansion.

13.2.2.1 Threshold-based Capacity Expansion Strategy (T)

The *threshold-based capacity expansion strategy* (T_{ut}^{la} or short T) is a simple heuristic with two parameters la and ut. la is called the *look-ahead time* and ut the *utilisation threshold*.

The heuristic works as follows: t_0 is the current period; if the utilisation threshold ut of a link is reached or exceeded in period $t_0 + la$, a capacity expansion is triggered.

For $la > 0$, a prediction of the r_{tf} for future periods is necessary. For $la = 0$, the measured utilisation of the current period is used. This heuristic with $la = 0$ resembles the rules of thumb often used by INSPs. The experiments below will show if the performance can be improved by basing the decision on predicted traffic demands.

This strategy without a look-ahead time is also the basic strategy in the paper of Hasslinger and Schnitter (2004).

---

[2] Prices for line cards seem to be relatively stable and therefore their price development should not influence interest costs. Contrary to that, the price of pure transmission rates dropped significantly in the past. d' Halluin *et al.* (2002) list some numbers for OC-48 links between 1999 and 2002. The prices decline between 5 and 43% per year, which corresponds to 0.4 and 3% per period assuming that a period equals a month.

13.2.2.2 Capacity Expansion Strategy (CE)

Strategy The basic capacity expansion strategy (CE) uses the solution of the optimisation problem that is specified in Model 13.1 in mixed integer programming (MIP) form:

The objective function (13.2) consists of the interest costs for capacity expansion and the congestion costs. The capacity doubling is modelled with constraints (13.3) to (13.6). As two opposing links have to have the same capacity constraints, (13.7) and (13.8) are necessary. To account for the congestion costs, constraint (13.9) is necessary. Finally, constraints (13.10) to (13.12) are the nonnegative binary constraints of the variables.

The optimal solution of Model 13.1 can be obtained with standard MIP solving methods (see Section 3.3) or with the faster algorithm that is presented below. It shows the optimal capacity expansion plan, the variables e_{tl} indicate the periods when the expansion of link l should be finished. The expansion of that link has to be triggered Δe periods before that. Please note that if Δe is rather long, it is possible that the optimal solution indicates that the expansion should have already been triggered before the current planning period t_0. In that case, the strategy triggers the expansion immediately in period t_0. Because of the uncertainty involved in the traffic predictions, this situation can be expected to occur for higher Δp.

Model 13.1 uses the predicted link loads v_{tl} as input. If the link loads are predicted correctly, it leads to the optimal capacity expansion plan. In a network with the shortest-path or any static routing, the link loads can be calculated directly from the predicted flow sizes of the predicted traffic matrix In a network that is using traffic engineering to optimise the routing, however, the routing can change from period to period. Flows are more likely to be routed over links that have just been expanded. Therefore, there is a mutual influence of the traffic engineering and the capacity expansion that cannot be accounted for with the above model as the capacity exact routing is not known in advance. The combined traffic engineering and capacity expansion (TMCE) strategy below extends Model 13.1 and takes this mutual influence into account by optimising the routing and the capacity expansion at the same time.

Faster Algorithm Model 13.1 models the capacity expansion problem assuming that the load of individual links can be predicted. In the resulting problem, the links *between different node pairs* are unconnected in the objective function (13.2) and in all constraints from (13.3) to (13.12). Therefore, the problem can be split up into smaller subproblems (one for every connected node pair). They can be solved independent of each other, resulting in the same optimal solution as Model 13.1. The subproblems can be solved efficiently with the following break-even algorithm:

For links l_1 and l_2 with $(l_1, l_2) \in \Omega$, the optimal period for the capacity to be doubled is when the additional congestion costs ΔC, that would be incurred if the capacity is not expanded, exceed the interest costs ΔI that can be saved by further delaying the capacity expansion. With the congestion costs function $p^x(u)$, the additional congestion costs ΔC in period t are

$$\Delta C = p^x\left(\frac{v_{t\,l_1}}{c_{t\,l_1}}\right) + p^x\left(\frac{v_{t\,l_2}}{c_{t\,l_2}}\right) - p^x\left(\frac{v_{t\,l_1}}{2 \cdot c_{t\,l_1}}\right) - p^x\left(\frac{v_{t\,l_2}}{2 \cdot c_{t\,l_2}}\right) \tag{13.13}$$

while the saved costs of delaying the capacity expansion of one period ΔI is given by

$$\Delta I = p^i c_{tl_1} + p^i c_{tl_2} \tag{13.14}$$

Model 13.1 Capacity Expansion (CE)

Indices	
$t = t_0, \ldots, (t_0 + \Delta T)$	Period t
$s = 1, \ldots, S$	Step s of the congestion costs function, see Figure 12.1
$l = 1, \ldots, L$	Link l
Parameters	
t_0	Current period
ΔT	Planning horizon
v_{tl}	Prognosed load of link l in period t
$c_{(t_0-1)\,l}$	Initial capacity of link l
p^i	Interest costs for link capacity
p^x_s	Additional costs in step s of the congestion costs function
q_s	Lower threshold of step s of the congestion costs function
M	Sufficiently large number, $M \geq \max_l(2^{\Delta T - 1} c_{0l})$
Ω	Set of link pairs (l_1, l_2) with opposite directions
Variables	
x_{stl}	Congestion costs variable, denotes by how much traffic the threshold of step s of the congestion cost function has been exceeded on link l
c_{tl}	Capacity of link l in period t
e_{tl}	Binary variable, 1 if the capacity of link l was doubled at the beginning of period t, and 1 otherwise

$$\text{Minimise} \sum_t \sum_l p^i c_{tl} + \sum_t \sum_s \sum_l p^x_s x_{stl} \tag{13.2}$$

subject to

$$c_{tl} \geq c_{t-1\,l} \quad \forall t\, \forall l \tag{13.3}$$

$$c_{tl} \leq 2 \cdot c_{t-1\,l} \quad \forall t\, \forall l \tag{13.4}$$

$$c_{tl} \leq c_{t-1\,l} + M \cdot e_{tl} \quad \forall t\, \forall l \tag{13.5}$$

$$c_{tl} \geq 2 \cdot c_{t-1\,l} + M \cdot (1 - e_{tl}) \quad \forall t\, \forall l \tag{13.6}$$

$$e_{tl_1} = e_{tl_2} \quad \forall t\, \forall (l_1, l_2) \in \Omega \tag{13.7}$$

$$c_{tl_1} = c_{tl_2} \quad \forall t\, \forall (l_1, l_2) \in \Omega \tag{13.8}$$

$$x_{stl} + q_s c_{tl} \geq v_{tl} \quad \forall s\, \forall t\, \forall l \tag{13.9}$$

$$c_{tl} \geq 0 \quad \forall t\, \forall l \tag{13.10}$$

$$x_{stl} \geq 0 \quad \forall s\, \forall t\, \forall l \tag{13.11}$$

$$e_{tl} \in \{0, 1\} \quad \forall t\, \forall l \tag{13.12}$$

Let t^* be the smallest period with $\Delta C > \Delta I$. The capacity should be expanded in that period. As the expansion takes Δe periods, it has to be triggered in period $t^* - \Delta e$.

13.2.2.3 Combined Traffic Engineering and Capacity Expansion (TMCE)

The TMCE strategy is similar to the CE strategy except that it is based on Model 13.2. Routing and the capacity expansion are considered at the same time. The model accounts for the fact that the routing in a subsequent period can be adapted to exploit the increased capacities of the links that were upgraded. The model is a combination of the CC traffic engineering strategy[3] described by Model 12.5 and the capacity expansion strategy of Model 13.1.

The objective function (13.15) consists of the total interest costs for capacity expansion and the total congestion costs. Constraint (13.16) is the routing constraint and constraint (13.17) is used to calculate the true load based on the expanded capacities.

The capacity increase to twice the previous capacity is modelled with constraints (13.18) to (13.21); opposing links are forced to the same capacity by constraints (13.22) and (13.23). The congestion costs are accounted for by constraint (13.24). Finally, constraints (13.25) to (13.29) are the nonnegative binary constraints of the variables.

Please note that Model 13.2 cannot be divided into subproblems as Model 13.1; therefore, the fast algorithm presented for the CE strategy cannot be used here. Instead, the MIP model has to be solved directly.

Model 13.2 Combined Traffic Engineering and Capacity Expansion (TMCE)

Indices

$t = t_0, \ldots, (t_0 + \Delta T)$	Period t
$s = 1, \ldots, S$	Step s of the congestion costs function, see Figure 12.1
$l = 1, \ldots, L$	Link l
$f = 1, \ldots, F$	Flow f
$p \in \rho_f$	Path p

Parameters

ΔT	Planning horizon
r_{tf}	Size of flow f in period t
ρ_f	Set of paths for flow f
ϕ_p	Set of links belonging to path p
$c_{(t_0-1)\,l}$	Initial capacity of link l
p^i	Interest costs for link capacity
p^c	Price for new link capacity

[3] Any of the other strategies could also be easily used, but CC was the best traffic engineering strategy in Chapter 12.

p_s^x	Additional congestion costs in step s of the congestion costs function
q_s	Lower threshold of step s of the congestion costs function
M	Sufficiently large number, $M \geq \max_l(2^{\Delta T-1} c_{0l})$
Ω	Set of link pairs (l_1, l_2) with opposite directions

Variables

v_{tl}	Load of link l in period t
a_{tfp}	Routing variable, flow f is routed via path p by this proportion
x_{stl}	Congestion costs variable, denotes by how much traffic the threshold of step s of the congestion cost function has been exceeded on link l
c_{tl}	Capacity of link l in period t
e_{tl}	Binary variable, 1 if the capacity of link l was doubled at the beginning of period t, and 1 otherwise

$$\text{Minimise} \sum_t \sum_l p^i c_{tl} + \sum_t \sum_s \sum_l p_s^x x_{stl} \tag{13.15}$$

subject to

$$\sum_{p \in \rho_f} a_{tfp} = 1 \qquad \forall t\, \forall f \tag{13.16}$$

$$\sum_f \sum_{p\,|\,l \in \phi_p} r_{tf} a_{tfp} = v_{tl} \qquad \forall t\, \forall l \tag{13.17}$$

$$c_{tl} \geq c_{t-1\,l} \qquad \forall t\, \forall l \tag{13.18}$$

$$c_{tl} \leq 2 \cdot c_{t-1\,l} \qquad \forall t\, \forall l \tag{13.19}$$

$$c_{tl} \leq c_{t-1\,l} + M \cdot e_{tl} \qquad \forall t\, \forall l \tag{13.20}$$

$$c_{tl} \geq 2 \cdot c_{t-1\,l} + M \cdot (1 - e_{tl}) \qquad \forall t\, \forall l \tag{13.21}$$

$$e_{tl_1} = e_{tl_2} \qquad \forall t\, \forall (l_1, l_2) \in \Omega \tag{13.22}$$

$$c_{tl_1} = c_{tl_2} \qquad \forall t\, \forall (l_1, l_2) \in \Omega \tag{13.23}$$

$$x_{stl} + q_s c_{tl} \geq v_{tl} \qquad \forall s\, \forall t\, \forall l \tag{13.24}$$

$$c_{tl} \geq 0 \qquad \forall t\, \forall l \tag{13.25}$$

$$x_{stl} \geq 0 \qquad \forall s\, \forall t\, \forall l \tag{13.26}$$

$$v_{tl} \geq 0 \qquad \forall t\, \forall l \tag{13.27}$$

$$a_{tfp} \in \{0, 1\} \qquad \forall t\, \forall f\, \forall p \in \rho_f \tag{13.28}$$

$$e_{tl} \in \{0, 1\} \qquad \forall t\, \forall l \tag{13.29}$$

Hasslinger and Schnitter (2004) present a heuristic for capacity expansion that takes into account the fact that traffic engineering can exploit the expanded capacity. They assume a traffic engineering strategy that minimises the maximum link utilisation and aim at maximising the average utilisation of the network. On the basis of these goals, their heuristic preferably upgrades links on a cut through the network. Their approach does not consider cost terms and the traffic engineering objectives are different from those in this section, TMCE: The TMCE strategy works with any of the path selection traffic engineering strategies discussed in the previous chapter and explicitly considers the trade-off between capacity costs and QoS. In Model 13.2, the path selection algorithm that minimises congestion costs was selected because it showed the best performance in the previous chapter. It explicitly showed better performance than strategies that minimise the maximum utilisation. In addition, TMCE is not a heuristic; it calculates the optimal capacity expansion plan and leads to the optimal solution in the absence of uncertain demands. It might be of higher computational complexity[4] but that should be relatively unimportant for a problem that only has to be solved once a month or once every three months. Because of the different goals and assumptions, it does not make sense to include that heuristic in this evaluation.

13.2.3 Performance Evaluation

13.2.3.1 Experiment Set-up

The same simulation environment and problem generation method as in Chapter 12 are used to evaluate the performance of the different capacity expansion strategies. Contrary to the single period evaluation of Chapter 12, 24 periods are considered here with one period representing one month. The size of the traffic flows r_{tf} is increased with a certain growth rate; the growth rate of the first period is drawn randomly from the interval [4%, 8%] and changed randomly by [–2%, 2%] points per period. The average growth rate of 6% leads to an average increase of roughly 100% per 12 periods; this expected increase is consistent with Odlyzko (2003) and Hasslinger and Schnitter (2004).

In a period t_0, the size of the traffic flows r_{tf} can be predicted with a maximal error $\pm 10\%$ for the following period; the maximal error increases by 3% per period $t > t_0 + 1$.

The expansion time Δe is set to $\Delta e = 3$ in the beginning, it will also be varied below. The decision that links to upgrade is made every $\Delta p = 3$ periods; that means we analyse a situation where the INSP is making the decision of when to expand its network every three periods.

As traffic engineering strategy, the CC strategy is used a maximum number of $n = 5$ paths between each node pair and maximal $\Delta l = 2$ additional hops. This strategy showed very good performance in the previous chapter.

The default congestion cost function from the previous chapter is used here (Function (1) from Figure 12.1). For evaluating the strategies, the absolute interest and congestion

[4] On a 2 GHz Pentium III with 512 MB RAM the TMCE strategy rarely needed more than one hour for the problems presented in this section.

costs are irrelevant as the results only depend on the relationship between those costs. The relationship between the interest for the network equipment and the congestion costs $c = \frac{p^i}{p^c(60\%)}$ is set to 1 in the beginning, it is varied later.

The threshold strategy T_{th}^{la} is evaluated with look-ahead la values of 0, 3 and 6 as well as various thresholds th that are depicted on the x axis of the graphs in this section. The absolute cost-minimal capacity expansion plan for the network can be calculated with the TMCE strategy if the uncertainty is switched off and Δp is set to 1; that is, if capacity expansion is planned every period based on the real future traffic. We call this the reference strategy REF[5].

13.2.3.2 Basic Results

The average congestion costs, the interest costs and the sum of both are shown in Figure 13.4. The 95% confidence intervals are also shown. Each experiment was repeated 20 times for different problem instances; all strategies and all different experiment set-ups solved the same 20 problems so the results are directly comparable. Because of the computational complexity, the experiment was restricted to the Telekom topology (see Appendix A). Selected experiments were repeated for the DFN topology and lead to very similar results. All costs are normalised relative to the costs of the REF strategy.

- The TMCE strategy that is executed only every $\Delta p = 3$ periods on the uncertain traffic predictions leads to only less than 1% higher total costs than when it is executed every period without uncertainty (REF). This strategy is obviously robust against the uncertainty and performs very well even if run only every third period.
- Comparing the CE with the TMCE strategy, there is a significant difference in costs. CE leads to more than 6.5% higher total costs than TMCE. With respect to the individual cost terms, CE leads to only slightly higher congestion costs than TMCE but to much higher interest costs. This results from CE not accounting for the fact that the traffic engineering algorithm can use the additional capacity of an expanded link to decrease the overall congestion in the subsequent periods. Therefore, CE overestimates the true congestion and invests too much in capacity leading to the relatively high interest costs and relative low congestion.
- Looking at the T strategies, one can first notice that all of these strategies reach the performance of the CE strategy if the threshold value th is set correctly. If it is set too high, the congestion costs explode and ruin the performance because capacity is expanded too late. This explosion becomes smaller for high look-ahead periods.

  If the threshold is set too low, too much capacity is bought and the interest costs increase. At the same time, the congestion costs decrease but that decrease becomes smaller and smaller because of the convex shape of the congestion cost function (see Figure 12.1). For decreasing values of th the congestion costs in Figure 13.4 approach a linear function with a small steepness corresponding to the lowest segment of the congestion cost function (1) in Figure 12.1.

[5] The optimal expansion plan can also be calculated by running TMCE once with ΔT encompassing all 24 periods. This, however, leads to a much higher overall computational complexity than solving TMCE with smaller ΔT every period.

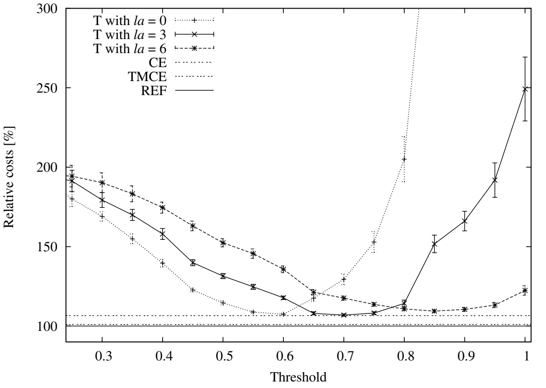

(a) Total Costs

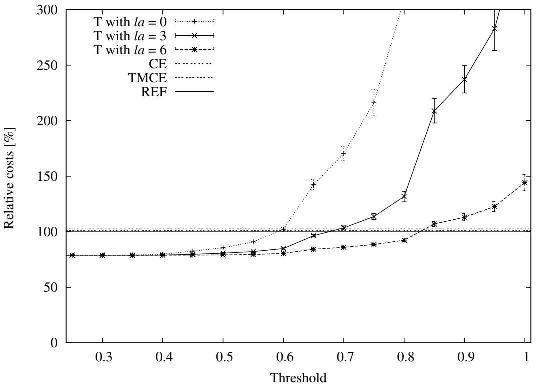

(b) Congestion Costs

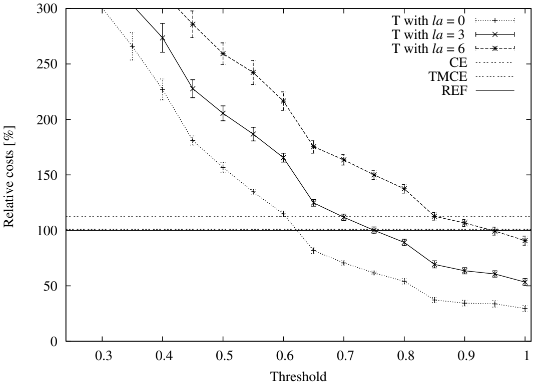

(c) Interest Costs

Figure 13.4 Costs of the Different Capacity Expansion Strategies for $c = 1$, $\Delta e = 3$, $\Delta p = 3$

- Comparing the different look-ahead values la for the T strategies, the lower the look-ahead value, the lower the optimal capacity expansion threshold. For $la = 0$ – that is, if the capacity expansion is based purely on measurements of the current period and no traffic predictions – the optimal threshold is around 60%, while it is close to 70% for $la = 3$ and 80–85% for $la = 6$. Obviously, as the traffic volume is generally increasing from period to period, higher look-ahead values lead to higher predicted utilisations and therefore higher optimal thresholds ceteris paribus.
  For a given threshold, the threshold strategy with the highest look-ahead time la leads to the lowest congestion costs and the highest interest costs because it triggers expansions significantly earlier because of the higher la value. As a result of that, this strategy leads to the highest total costs for low thresholds because the interest costs dominate in that region and to the lowest total costs for high thresholds because the congestion costs dominate in that region.

13.2.3.3 Variation of the Cost Ratio

Next, the effect of changing the cost ratio c is analysed. c measures the ratio between the interest costs for the equipment and the congestion costs. The interest costs are determined by the prices for the network hardware and the interest rate of the financial market. The congestion costs, however, are largely determined by the provider itself depending on how important QoS (low congestion) is for its network, its business model, and its customers. In Figure 13.4, the results for a cost ratio of $c = 1$ are depicted, Figure 13.5 shows the results for lower and higher cost ratios:

- The general shapes of the congestion and interest cost functions remain the same but as they are added in different ratios to the total cost function now, the total cost function is distorted compared to the original one in Figure 13.4.
- If c is set to 0.2, the congestion costs are judged five times higher than before. This resembles a provider for which QoS is highly important. The congestion costs dominate the overall performance and the total costs more closely resemble the congestion cost function. The optimal threshold for the T strategies is significantly lower now as can be expected. The TMCE strategy offers a 3% cost advantage compared to the best T strategies and the CE strategy; it leads to only 0.35% higher costs than the optimum.
- If c is increased to 5, the influence of the congestion costs is five times smaller than before. The general shape of the total cost function is now strongly influenced by the shape of the interest cost function. The optimal expansion threshold of the T strategies is higher than before. The TMCE strategy offers a 16% cost advantage compared to CE and a 12% advantage compared to the best T strategies. It comes as close as 2.2% to the optimum.

13.2.3.4 On the Capacity Expansion Process

Next, the parameters of the capacity planning process are changed. So far, for every $\Delta p = 3$ periods the capacity planning strategies were run and a single expansion took

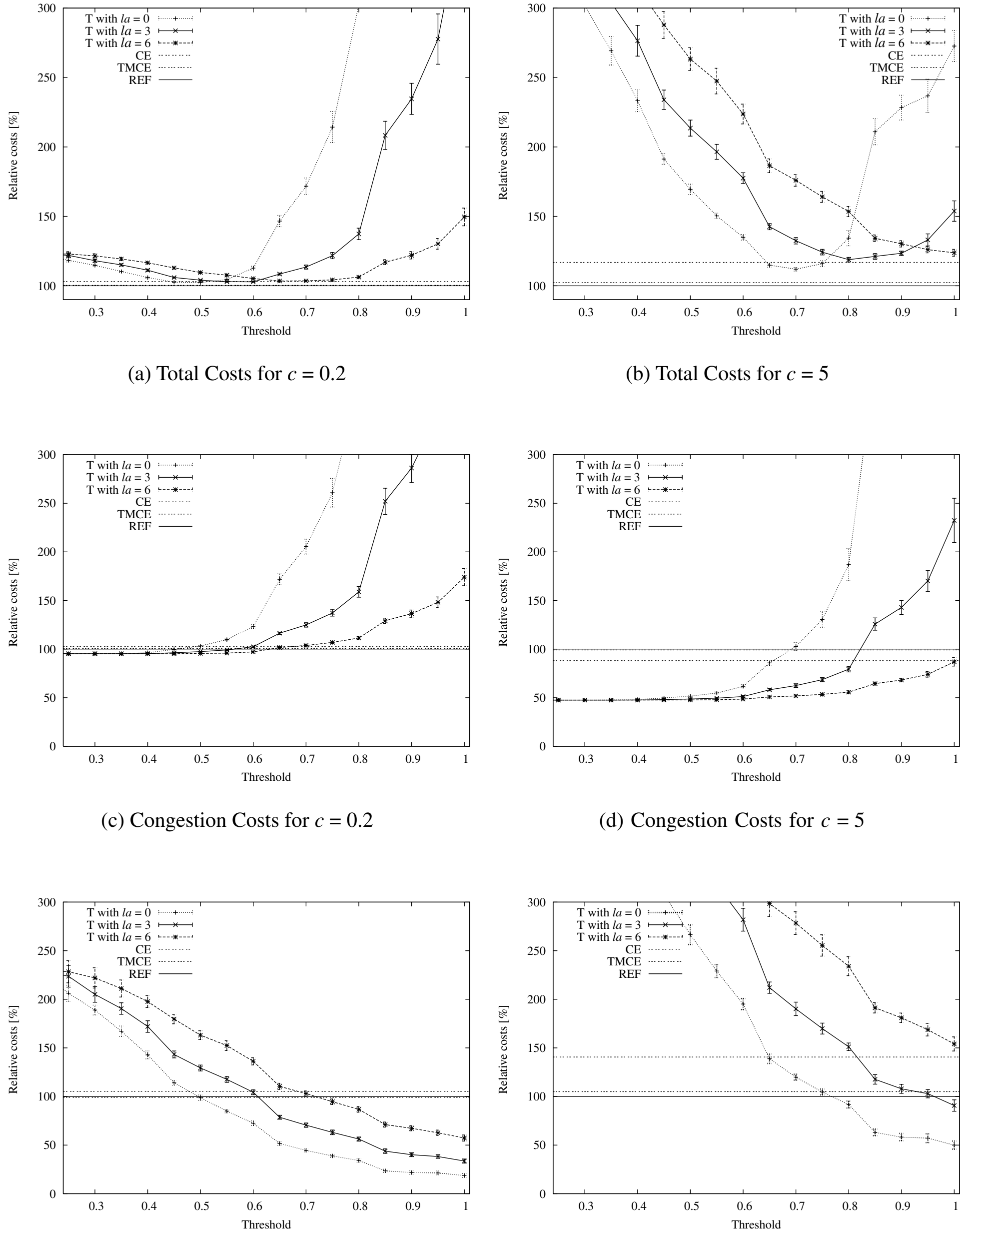

(a) Total Costs for $c = 0.2$

(b) Total Costs for $c = 5$

(c) Congestion Costs for $c = 0.2$

(d) Congestion Costs for $c = 5$

(e) Interest Costs for $c = 0.2$

(f) Interest Costs for $c = 5$

Figure 13.5 Costs of the Different Capacity Expansion Strategies for Different Cost Ratios c; $\Delta e = 3,\ \Delta p = 3$

$\Delta e = 3$ periods to take effect. Figure 13.6 shows the resulting total costs for different values of Δe and for different values of Δp.

- If the expansion time Δe increases, the thresholds when an expansion should be triggered obviously decrease as visible on the left-hand side of Figure 13.6. The performance of CE and TMCE is not influenced significantly. The same holds true for the respective optimal values of the T strategies.
- On the right-hand side of Figure 13.6, the effect of an increasing time between two planning periods Δp is visible. An increase in Δp leads to a higher planning uncertainty that should be countered by decreasing the expansion threshold of the T strategies. The overall performance of all strategies decreases with an increasing Δp. TMCE for $\Delta p = 1$ leads to optimal performance in almost all cases uninfluenced by the uncertainty, while for $\Delta p = 6$ it loses 5% performance. CE loses 13% while the T strategies lose 6%. For high Δp, the T strategies perform significantly better than CE.

13.2.4 Recommendations

In the face of traffic volumes that are growing in the medium and long run, the capacity expansion decision is not about *whether* to upgrade capacity but rather *when* to upgrade capacity. This decision is directly influenced by the trade-off between the costs of the network (therefore the network efficiency) and the QoS. This trade-off was modelled by the price ratio c.

We evaluated different capacity expansion strategies with respect to their total costs. The total costs are the interest costs for the networking equipment and the congestion costs, a fictive cost term describing the ill-effects of a congested network. We now summarise the conclusions for the different strategies:

- The CE strategy bases its decision on the solution of an optimisation problem that assumes fixed routing for the network. This strategy leads to significantly worse performance than the TMCE strategy and in some cases worse than the T strategies. It cannot be recommended for networks that use traffic engineering. For networks with a fixed routing (e.g. plain shortest-path routing), this strategy is equivalent to the TMCE strategy and can be recommended.
- The TMCE strategy takes the mutual influence of the capacity expansion and the traffic engineering strategy into account. It led to the best performance in all experiments. Depending on the settings, this comes as close as 0 to 5% to the optimal solution. This strategy can be clearly recommended. In the absence of uncertainty and for $\Delta p = 1$, it yields the optimal solution.
- The threshold strategies (T) are simple rules of thumb used by today's INSPs that expand a link once a certain utilisation threshold is reached in the current period or predicted to be reached in a certain future period. These strategies can lead to good performance if the threshold parameter is set to the correct value. The performance degrades rapidly if it is set too high, especially when using current and not predicted future link utilisations. These strategies can be recommended only if the threshold value is set correctly.

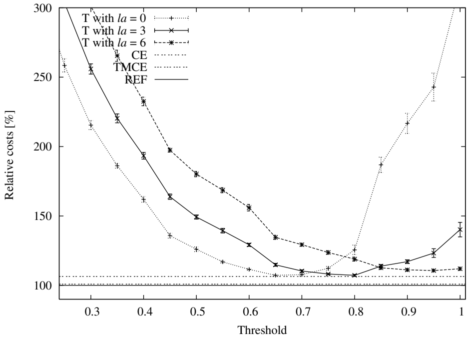

(a) $\Delta e = 1$, $\Delta p = 3$

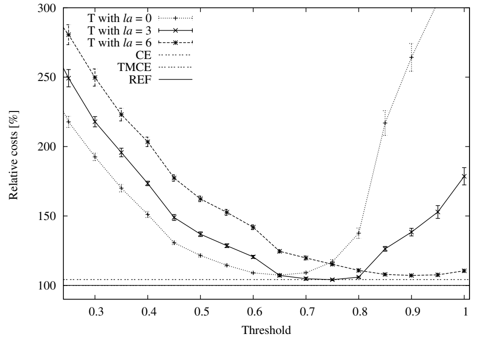

(b) $\Delta e = 3$, $\Delta p = 1$

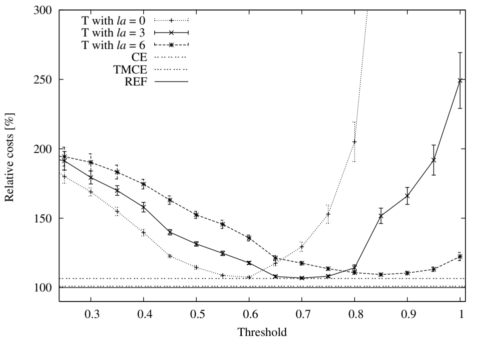

(c) $\Delta e = 3$, $\Delta p = 3$

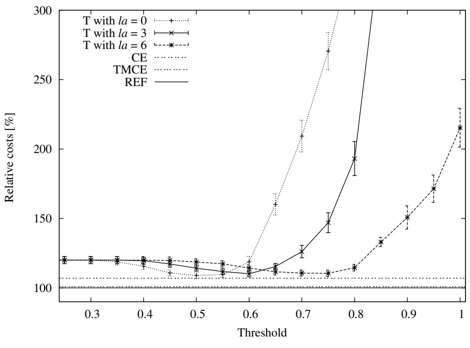

(d) $\Delta e = 6$, $\Delta p = 3$

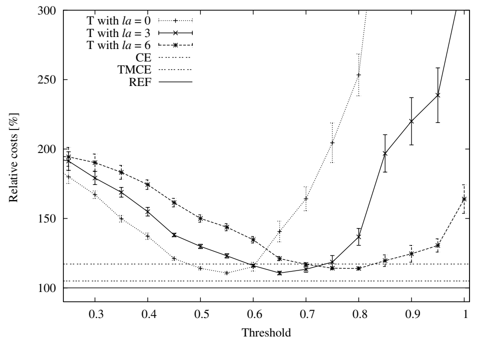

(e) $\Delta e = 3$, $\Delta p = 6$

Figure 13.6 Total Costs for Variation of the Parameters Δe and Δp of the Capacity Expansion Process; $c = 1$

If the T strategy is used with predicted demands, the overall performance does not increase significantly, therefore it is probably not worth the effort for predicting the future demands. However, if a provider is unsure about the correct setting of the threshold parameter, it is worth considering a higher look-ahead time because it can significantly reduce the ill-effects of a too high threshold value.

With respect to the overall capacity expansion process, the expansion time Δe for a link has no massive influence on the overall performance but the time interval between two planning periods Δp (when capacity expansions are considered) has. For the given parameter settings, a capacity expansion planning every three months yielded satisfactory results that were improved by only 1% if reduced to every month.

13.3 On the Influence of Elastic Traffic*

As argued before, traffic matrices are fundamental for network design and traffic engineering problems. Normally, the traffic matrix entry r_{ij} is expressed statically as a scalar – we call a traffic matrix with static predictions r_{ij} a *static traffic matrix*. However, Internet traffic is dominantly TCP traffic that adapts to changing network conditions like routing or the link capacity. This effect is systematically neglected when using static traffic matrices. The effect of capacity changes was probably negligible in times when the Internet was dominated by web traffic that consisted of huge numbers of short-lived TCP connections dominated by the slow start and not the elastic congestion avoidance phase. Traffic matrix entries at these times mainly increased if the customer base or browsing behaviour changed.

Nowadays, however, most of the traffic is generated by peer-to-peer (P2P) applications, see Chapter 5. As discussed in Section 5.2, these applications use mainly long-lived TCP connections for file transfers. This supports the assumption of this chapter that long-lived reactive TCP connections start dominating the Internet traffic.

Besides P2P traffic, future multimedia Internet traffic like streaming videos can also be expected to be TCP friendly and therefore show similar reactive effects as long-lived TCP connections that we are looking at in this section, see Handley *et al.* (2003).

Because of this, it is important to investigate the effect of the elasticity of long-lived TCP connections in their congestion avoidance phase on traffic matrices used as input for *network design* and *network engineering problems*. Normally, these problems are based on a static traffic matrix and ignore the effect that the new capacity (or capacity change) has on the amount of traffic matrix itself. We use the term *elastic traffic matrix* for a traffic matrix M with entries $r_{ij} = f(\ldots)$ that capture the elasticity of the TCP traffic and investigate the use of these elastic matrices in this section.

We developed three different network models to analyse this effect. They consist of a combination of the TCP formula and queueing theory. They are presented in Appendix D and form the analytical foundation for further analysis. We first generally analyse the elasticity of traffic matrices and then determine the impact on capacity expansion.

13.3.1 Elasticity of Traffic Matrices

The influence of the elasticity of a traffic matrix when the capacity of the network changes while all other conditions remain the same (ceteris paribus) is being analysed in this section. The effects described here are neglected when static traffic matrices are used.

* Reproduced by permission of VDE Verlag GMBH.

We base our analysis on the different network models derived and described in Appendix D.

13.3.1.1 Single Link Experiments

We start our analysis with an extensive series of experiments on a single link. Figure 13.7 shows the rate increase $(r_{ij}^{\text{new}} - r_{ij}^{\text{old}})/r_{ij}^{\text{old}}$ of the symmetrical macroflows over the single link topology for different queue lengths B (measured in packets) and different values for the external loss $\tilde{p}$ and delay $\tilde{q}$ when the link capacity μ_l is doubled $\mu_l^{\text{new}} = 2 \cdot \mu_l^{\text{old}}$. Figure 13.7(a) lists the results for the basic model of Section D.1, Figure 13.7(b) shows the results for the model with discrete service times of Section D.2. We used two different service time distributions. Distribution A consists of 50% packets with a size of 40 bytes and 50% packets with a size of 1500 bytes. Distribution B consists of packets of size 1000 bytes only. We assumed a line rate of 1 Mbps and had to use a rather low queue length of $B = 10$ packets because the loss probability formula gets too complicated for larger values of B to be handled analytically.

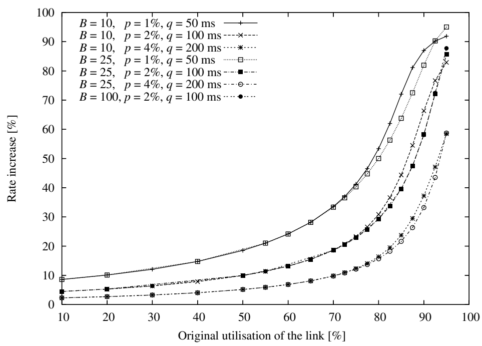

(a) Results for the Basic Network Model

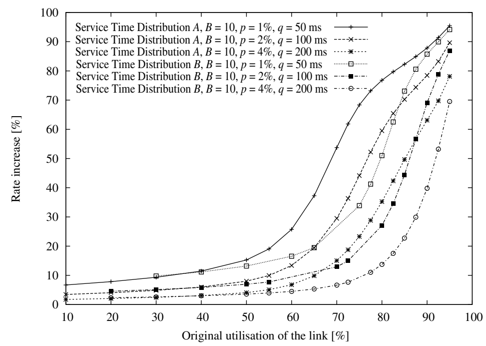

(b) Results for Discrete Service Times

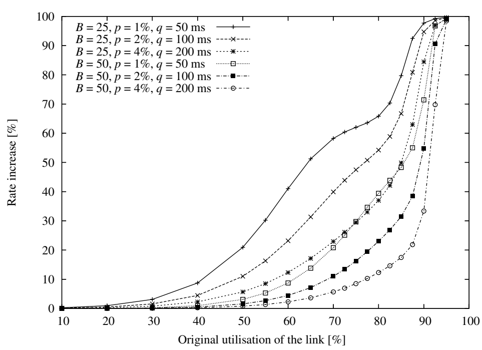

(c) Results for Self-Similar Traffic

Figure 13.7 Single Link Experiment Results. (Reproduced by permission of VDE Verlag GMBH)

Figure 13.7(c) shows the results obtained if we apply the model for self-similar traffic from Section D.3. A Hurst parameter of $H = 0.75$, a line rate of 1 Mbps, an average service packets size of 1000 bytes and the corresponding average service time were used.

Looking at the results, one notices that for all three different network models and most parameters, the general behaviour of the traffic is the same. Up to a certain utilisation threshold of the analysed link, the traffic is affected by the increase in capacity only slightly. Then, the traffic increases very quickly. If the initial utilisation of the link is high enough, the analysed link forms a strong bottleneck and all additional capacity is used up completely by a rate increase of 100%.

The step is steeper for the $M/M/1/B$ network model than for the other two models that can be deemed more realistic.

13.3.1.2 Different Topologies

We now analyse the elasticity in the form of the rate increase for more complex topologies than the single link topology of the previous experiments. Figure 13.8 summarises the results for three different topologies, the backbone of the Deutsche Telekom, a dumb-bell topology with a single bottleneck link and three nodes on each side of the bottleneck and a simple star-shaped topology with one internal and four external nodes. The value of t_{ij} is varied between 10^{-1} and 10^{2}. The network capacity for each t_{ij} is doubled and the rate increase recorded. As one can see, the different topologies lead to similar results. While most of the rate increases are very small (more than 50% of the times the rate increase was below 10%), there are a significant number of times where the rate increase was very high. Because of the different paths the different flows take through the topology, the rate increase can be higher than 100% if a series of links is doubled in capacity for a flow.

If a traffic matrix is used in the context of network design or capacity expansion, the elasticity of the traffic can be neglected up to a certain utilisation of a link. Once that threshold is passed, the error can be significant.

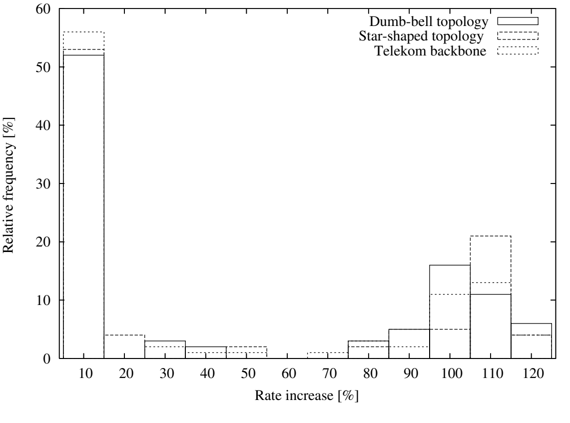

Figure 13.8 Rate Increase for Different Topologies. (Reproduced by permission of VDE Verlag GMBH)

13.3.2 Impact on Capacity Expansion

We now address the question of how the elasticity of the traffic matrix affects capacity expansion and how the capacity expansion strategies of Section 13.2 can be adapted.

If the network models of Appendix D are combined with the MIP model of the CE or TMCE strategy of Section 13.2, the resulting optimisation problem becomes nonlinear and can no longer be solved easily. For these strategies, an iterative approach could be used to take the elasticity of the traffic matrix into account. The threshold heuristic T of Section 13.2, however, can be combined directly with the network models of Appendix D. We do so exemplarily for the threshold heuristic T with a look-ahead value of $la = 0$. Using that heuristic, we can evaluate the impact of the elastic traffic matrices on capacity expansion: If the utilisation ρ_l exceeds a certain threshold th on a link l, the capacity expansion for that link is triggered. For this analysis, we assume that the link capacity is effectively doubled to the beginning of the next period after the one that triggered the expansion.

Traffic is given in form of the parameter t_{ij} of equation D.1. The actual traffic volume passed through the network is elastic and thus reacts to changes in capacity.

In 'classical' network design and capacity expansion algorithms, the elasticity of the traffic is ignored. The problem is that by increasing the capacity of a link, the traffic flows through that link will increase their rate and therefore also the utilisation of the other links they are flowing through. This can lead to the situation (a) that immediately after the expansion the threshold th on other links is exceeded and not predicted by the classical model with static traffic matrices. It will take an additional period until these links too can be expanded. Furthermore, if a link is an extreme bottleneck for some flows, it is possible that the utilisation will not significantly decrease if the link is doubled. This effect (b) can also not be predicted with static matrices. This effect was, for example, observed when the UK ISP Rednet quadrupled their DSL access link capacity, as reported to the author.

Using the models of Appendix D, we can predict the traffic increase and utilisation change of a planned network expansion and avoid the effects (a) and (b). We use the following simulation as a proof of concept:

Using the backbone topology of the Deutsche Telekom again, we generate a traffic matrix with random entries r_{ij} between 1.0 and 5.0. We use this for the initial parameters t_{ij}. A starting line rate of 1 Mbps is used for all links; it is doubled for each link before the actual simulation until all link utilisations are below 70%. We then simulate 10 periods; at the beginning of each period each traffic matrix entry is increased randomly between 5 and 20%. The basic model of Appendix D is used to calculate the link utilisations – we assume that the result of these initial calculations represents the SNMP (simple network management protocol) data collected by the provider. An external loss of 2% and delay of 100 ms is assumed; this results in a not too aggressive behaviour of TCP. In the experiment, the expansion of a link l is triggered if it has a utilisation of $\rho_l \geq th = 0.75$.

In order to capture the elasticity of the traffic matrix, we can again use our basic model to predict the effect of the triggered capacity expansions in order to avoid the effects (a) and (b) described above. We do so and measure how often these effects occur.

Effect (b) was not observed. Because we increase the rates only in moderate steps and allow the capacity to increase in each period effect, (b) does not occur in our simulations and can therefore not be avoided by the model. Effect (a), however, occurred 12 times (in 23% of all expansions) in the experiment and can be avoided by using our prediction of the elastic traffic. This example demonstrates that our concept works and helps in capacity expansion decisions.

13.4 Summary and Conclusions

Capacity expansion is an important and frequent task in today's IP networks because the traffic volume is increasing steadily. In this chapter, the influence of capacity expansion on the performance of the different QoS architectures of Chapter 8 was analysed first. If capacity is abundant, the differences between the QoS architectures vanish. However, if capacity is scarce, the systems with a strict admission control manage to maintain QoS while the other systems suffer to different extents.

Different capacity expansion algorithms were presented and evaluated. One of the introduced algorithms considers the effect of traffic engineering and capacity expansion at the same time. It leads to the best performance and is very robust against uncertain demand predictions. The simple heuristics that are often used by actual INSPs also show good performance – but only if their parameters are set correctly. The effects of several parameters on these parameters were also studied in this chapter.

Finally, the effects of elastic TCP traffic on traffic matrices and capacity expansion were discussed with some analytical models. It influences the capacity expansion measures if the network is highly utilised before the expansion. It was shown how this effect can be predicted and reacted upon accordingly.