10

Optimising the Interconnection Mix*

In this chapter, decision theory and mathematical programming methods are used to model the problem of finding the optimal set of peering and transit providers for an Internet Network Service Provider (INSP). We consider costs, reliability issues, Quality of Service (QoS) and the fact that traffic and tariffs are changing over time or that new providers enter the market. Heuristics and exact algorithms are presented and their performance is evaluated in extensive simulations. Related works were discussed in the previous chapter.

For the ease of presentation, in this chapter we assume that only the classical peering and transit interconnections are used (see Section 9.2.2 and 9.2.3). However, the models can easily be extended for nonclassical interconnection types.

We start with the presentation of an optimisation model for minimising the interconnec tion-related costs in Section 10.1. An exact method for solving the problem is presented; it can be used to find the optimal set of peering and transit partners for one INSP. It is compared with some heuristics in a performance evaluation. The evaluation shows that the interconnection mix significantly influences the cost structure and overall efficiency of the network of an INSP.

Minimising the interconnection-related costs, however, is not the only important goal of an INSP with respect to interconnections. The interconnection reliability and the influence of the interconnections on the QoS are important aspects, too. Therefore, we show and analyse how the original model of Section 10.1 can be extended to also consider reliability aspects in Section 10.2 and QoS requirements in Section 10.3. We evaluate the different reliability and QoS policies by simulations.

In Section 10.4, we show how previous strategies can be extended for the dynamic problem situation, which is evaluating whether a given set of peering and transit partners is still optimal considering changes in the traffic mix or the cost structure of the other providers. The administrative costs of changing peering and transit partners or Internet exchange points (IXPs) are also considered. Again, the models are evaluated using simulation.

*Reprinted from Elsevier Computer Networks Journal, 46(1), Oliver Heckmann, Jens Schmitt, Ralf Steinmetz, Optimizing Interconnection Policies, 19-39, Copyright 2004.

The Competitive Internet Service Provider: Network Architecture, Interconnection, Traffic Engineering and Network Design
Oliver Heckmann ISBN: 0-470-01293-5

10.1 Costs

First, a mathematical programming model for finding the optimal set of peering and transit partners for one INSP is presented. It is important to model cost functions for peering and transit partners and for IXPs realistically. We interviewed different INSPs and IXPs and model the cost functions based on these interviews and based on Norton (2002).

We show how the model can be solved exactly and present some heuristics that closely resemble what providers do today. The heuristics are evaluated against the optimal solution that can be obtained by our model.

10.1.1 Description

Finding the optimal transit and peering partners as well as the necessary IXPs for one INSP is modelled by the following optimisation model. We assume that there are R different routes, and that the provider can predict the traffic for each route[1]. There are J transit providers offering transit service for all routes. The transit providers can be connected via a direct line. There are I peering providers offering peering only for some specific routes. Each peering provider is connected to at least one IXP and can be reached only via an IXP it is connected to.

The optimisation model tries to minimise the total costs that consist of the costs for connecting to IXPs, the additional peering costs for peering with a provider at an IXP, and the transit costs, including the direct line costs for connecting to the transit provider's closest POP (point of presence).

The costs for connecting to an IXP are largely fixed costs (leased line and backup line to the IXP, rent for rack space, costs of the exchange router, fixed IXP fees, etc.) that increase if the volume transferred via the IXP exceeds certain thresholds (representing upgrades to the leased lines and the exchange router, additional IXP fees, etc.); see Figure 10.1 (a).

A peering interconnection with provider i can only be made if there is a connection to an IXP where the peering provider i is present. The costs for peering are largely volume-independent fixed costs (transactional costs for the peering agreement, engineering costs). Some low volume dependent costs for peering[2] can also occur as some IXPs, for example LINX, also charge per volume of peering traffic, see Figure 10.1 (b). Giovannetti *et al*. (2003) presents interesting results from an empirical study of peering via the IXP in Milan. The peering decision is influenced, for example, by the proximity of the INSP's headquarters; a factor that can be modelled via the provider-specific fixed peering costs.

The costs for transit consist of fixed costs plus volume-dependent variable costs that decrease when certain volume thresholds are reached (see Figure 10.1 (c)). We assume that each transit and peering provider and each IXP can only accept traffic up to a certain maximum capacity.

[1] Please note that a route in the context of this book can be individual Border Gateway Protocol (BGP) routes. For performance improvement, nonoverlapping BGP routes *can* be aggregated to a single route in the optimisation problem.

[2] The variable costs for peering could also be accounted for in the IXPs cost function but to add them to the peering costs is more efficient from a modelling point of view and more flexible.

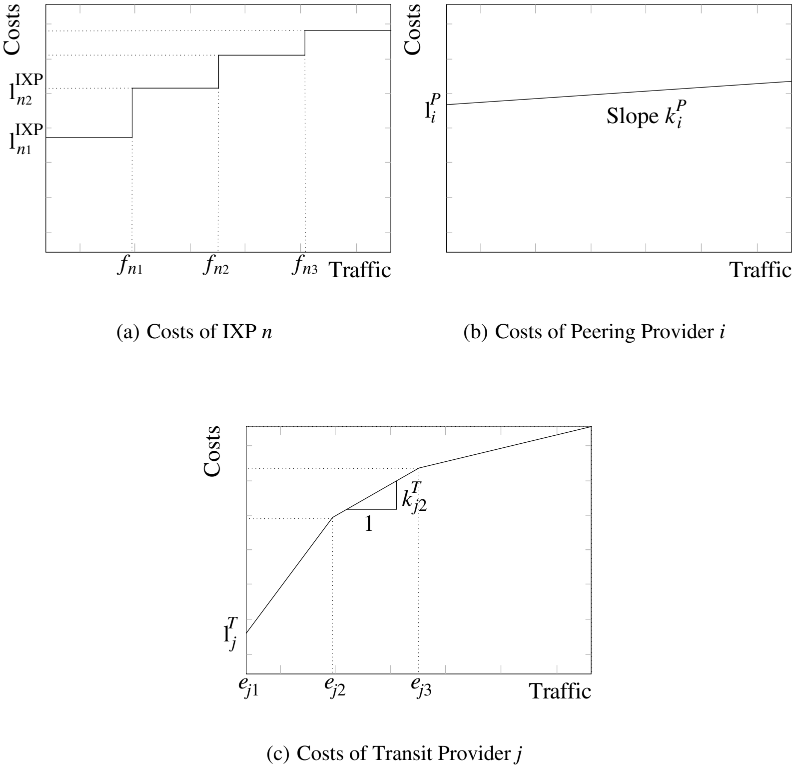

Figure 10.1 Cost Functions

The problem is formally described by the mixed integer programming (MIP) model given with Model 10.1. The target function (10.1) minimises the total costs (IXP, peering and transit). We use the variables $\tilde{x}^{T}_{jm}$ to keep track of how much of the traffic of provider j is in segment m of its cost function. Constraint (10.2) connects the variables $\tilde{x}^{T}_{jm}$ to the routing variables x^{T}_{jr} of the same transit provider j: the total amount of traffic that is divided among all routes has to be equal to the traffic in all segments of the cost function. Constraint (10.3) ensures that the entire traffic demand for one route is satisfied by the combination of peering and transit interconnections.

Constraints (10.4) and (10.5) ensure that the transit cost function segments are 'filled' correctly: (10.4) limits the amount of traffic in one segment to the segment size. For the highest segment, (10.4) forms the capacity constraint of the transit provider. For concave cost functions, the higher segments would be filled first because of their lower volume costs. Therefore (10.5) is necessary; a higher segment of a cost function can only be used once the lower segment is completely full.

Constraint (10.6) is the capacity constraint for the peering providers. The other constraints are the nonnegativity and binary constraints of the variables. Constraints (10.4), (10.5) and (10.6) also connect the binary y variables to the corresponding x variables and make sure that traffic can only be routed (indicated by the x variables) through providers with which an interconnection exists (indicated by the y variables).

Model 10.1 Cost Minimising Interconnection Model

Indices

$i = 1, \ldots, I$	Peering provider i
$j = 1, \ldots, J$	Transit provider j
$r = 1, \ldots, R$	Route r
$m = 1, \ldots, M_j$	Part m of the cost function of transit provider j
$n = 1, \ldots, N$	IXP n
$s = 1, \ldots, S_n$	Step s of the cost function of IXP n

Parameters

$\widehat{x_r}$	Traffic prognosis for route r
l_i^P	Fixed costs for an interconnection with peering provider i
k_i^P	Price per unit of volume for an interconnection with peering provider i
c_i^P	Capacity of peering provider i
$\Re_i$	Set of routes offered by peering provider i
l_j^T	Fixed costs for an interconnection with transit provider j
M_j	Number of steps in the cost function of transit provider j
$c_j^T = e_{j(M_j+1)}$	Capacity of transit provider j
e_{jm}	Lower volume limit of step m of the cost function of transit provider j, see Figure 10.1 (c)
k_{jm}^T	Price per volume in step m of the cost function of transit provider j
S_n	Number of steps in the cost function of IXP n
Ψ_n	Set of peering providers that are connected to IXP n
f_{ns}	Upper volume limit of step s of the cost function of IXP n, see Figure 10.1 (a)
f_{nS_n}	Capacity of IXP n
l_{ns}^{IXP}	Costs for IXP n if it is used and the traffic volume via IXP n is in step s of the cost function of IXP n
Inf	Large number (resembling infinity)

Variables

$x_{ir}^{P} \forall r \in \Re_i$	Amount of traffic for route r passed through peering provider i
y_i^{P}	Binary variable, set to 1 if an interconnection with peering prov. i is made
x_{jr}^{T}	Amount of traffic for route r passed through transit provider j
$\tilde{x}_{jm}^{T}$	Traffic volume in segment m of the cost function of transit provider j
y_{jm}^{T}	Binary variable, set to 1 if cost function segment m of transit provider j is used
y_{ns}^{IXP}	Binary variable, set to 1 if IXP n is used and the traffic volume via IXP n is in step s of the cost function

$$\text{Minimise} \sum_j \sum_{m \in M_j} k_{jm}^T \tilde{x}_{jm}^T + \sum_j l_j^T y_{j1}^T + \sum_i \sum_{r \in \Re_i} k_i^P x_{ir}^P + \sum_i l_i^P y_i^P + \sum_n \sum_s l_{ns}^{IXP} y_{ns}^{IXP} \tag{10.1}$$

subject to

$$\sum_m \tilde{x}_{jm}^T = \sum_r x_{jr}^T \qquad \forall j \tag{10.2}$$

$$\sum_{i \mid r \in \Re_i} x_{ir}^P + \sum_j x_{jr}^T = \widehat{x}_r \qquad \forall r \tag{10.3}$$

$$\tilde{x}_{j\,m}^T \le (e_{j\,m+1} - e_{j\,m}) \cdot y_{j\,m}^T \qquad \forall j\, \forall m \tag{10.4}$$

$$\tilde{x}_{j\,m}^T \ge (e_{j\,m+1} - e_{j\,m}) \cdot y_{j\,m+1}^T \qquad \forall j\, \forall m = 1, \ldots, M_j - 1 \tag{10.5}$$

$$\sum_{r \in \Re_i} x_{ir}^P \le c_i^P \cdot y_i^P \qquad \forall i \tag{10.6}$$

$$\sum_s y_{ns}^{IXP} \le 1 \qquad \forall n \tag{10.7}$$

$$\sum_{i \in \Psi_n} x_i^P \le f_{ns} \cdot y_{ns}^{IXP} + \mathrm{Inf} \cdot \sum_{t=1,\, t \ne s}^{S_n} y_{nt}^{IXP} \qquad \forall n\, \forall s \tag{10.8}$$

$$x_{ir}^P \ge 0 \qquad \forall i\, \forall r \in \Re_i \tag{10.9}$$

$$x_{jr}^T \ge 0 \qquad \forall j\, \forall r \tag{10.10}$$

$$\tilde{x}_{j\,m}^T \ge 0 \qquad \forall j\, \forall m \tag{10.11}$$

$$y_i^P \in \{0, 1\} \qquad \forall i \tag{10.12}$$

$$y_{jm}^T \in \{0, 1\} \qquad \forall j\, \forall m \in M_j \tag{10.13}$$

$$y_{ns}^{IXP} \in \{0, 1\} \qquad \forall n\, \forall s \tag{10.14}$$

Constraints (10.7) and (10.8) are needed for the IXP cost function: if all y_{ns}^{IXP} for one specific n are zero, then IXP n is not used. Otherwise, exactly one variable y_{ns}^{IXP} will be 1 indicating in which step of the cost function the total traffic volume via this IXP n lies. This is ensured by constraints (10.7) and (10.8). The exact solution for this problem can be found using standard MIP solving techniques as discussed in Section 3.3.

10.1.2 Evaluation

Now, the exact solution of our model and the solutions obtained with some heuristics are evaluated in a series of simulations.

For validating the models and for meaningful results, it is important to use realistic data as input. However, providers are very reluctant to reveal information about their cost structure and explicitly do not allow publication of this information. The input data in the following experiments is based on actual but randomised data of a real (medium-sized) INSP. To cover a wider space, specific parameters are varied systematically.

10.1.2.1 Simulation Setup

We evaluate different *scenarios*. A scenario is specified by a given number of peering providers, transit providers, IXPs, routes, and an interval from which traffic and costs for these providers or routes are drawn. A *scenario instance* is created by randomly creating cost functions and traffic demand vectors from the scenario-specific parameter intervals. Per scenario, n=100 instances are created and solved. The averages of all instances as well as the 95% confidence interval are computed and form the basis of the following evaluation.

The parameter intervals for the basic set of scenarios are given in Tables 10.1 and 10.2. For the simulations, we assume that each peering provider offers one route and always has enough capacity for that route. The traffic demand for one route is drawn uniformly distributed from the *traffic demand for a peering provider's route* interval. The BGP routes not covered by the peering providers routes are modelled with one additional larger route. The traffic for that route is determined by the *traffic demand for the rest of the world* parameter.

The fixed costs for the peering providers are calculated as specified in the *fixed peering costs* interval in Table 10.1. We added a small amount per volume to give the fixed costs of larger peering providers the tendency to be higher than those of smaller providers. The variable peering costs are set uniformly low. The transit costs are calculated as specified in the tables; the transit capacity is drawn from the *capacity of a transit provider* interval and divided evenly across the different segments of the cost function.

The different steps of the cost function for an IXP are calculated as specified in Table 10.1. We assume that all steps are of the same size and that the costs for the later steps are significantly less than for the lower ones.

Table 10.2 lists the four scenario-dependent parameter ranges; all 16 possible combinations are evaluated. Each scenario has a number from 0 to 15. In scenario s the first

Table 10.1 Constant Parameter Intervals

Parameter Description	Parameter Value/Interval
Traffic demand for a peering provider's route	[50, 1000]
Fixed transit costs	[10,000, 50,000]
Capacity of a transit provider	[50%, 150%] of total traffic
Fixed peering costs (not including the costs to connect to the IXP)	[3000, 6000] + [0, 4] times traffic demand of peering provider's route
Variable peering costs	4
Number of steps of the transit cost function	5
Number of IXPs	4
Number of steps in the cost function of an IXP	4
Basic costs for connection to an IXP	[25,000, 60,000]
Costs for additional volume transferred to IXP	[9000, 18,000] and reduced by 10% each step
Volume that can be transferred to an IXP in one step of its cost function	Total traffic demand divided by 16

Table 10.2 Scenario-dependent Parameter Intervals

Bit	Parameter Description	Value/Range if Bit = 0	Value/Range if Bit = 1
#1	Number of peering providers	100	200
#2	Number of transit providers	10	20
#3	Traffic demand for rest of the world	$20 \times$ av. traffic demand of peering provider's route	$60 \times$ av. traffic demand of peering provider's route
#4	Variable transit costs	[20, 80], decreasing by [5%, 20%] each step	75% of the costs for bit #4 = 0

parameter from Table 10.2 is used if the corresponding bit in s is not set, otherwise the second parameter is used. For scenario $s = 7$, the second parameter intervals will be used for the number of peering and transit providers and the traffic demand for the rest of the world (bits #1, #2, #3); the first parameter interval will be used for the variable transit costs (bit #4).

The commercial MIP solver CPLEX (see ILOG CPLEX (2004)) is used to calculate the exact solution for the optimal interconnection Model 10.1 (**OPT**). We compare the solution we obtain with several heuristics.

10.1.2.2 Description of the Heuristics

Our comparison includes heuristics that can mimic the behaviour of some real-world providers:

- **Transit Heuristic (H TR)**
  The transit heuristic uses the cheapest transit provider or, if the capacity of the cheapest is not sufficient, the cheapest set of transit providers. It does not use peering.
- **Peer-With-All Heuristic (H PA)**
  The peer-with-all heuristic connects to all available IXPs and peers with every peering provider available. For the remaining traffic, it chooses the cheapest transit provider (or set of transit providers).
- **Peer-at-Selected-IXPs Heuristic (H PS)**
  The peer-at-selected-IXPs heuristic is similar but more careful in the selection of IXPs. If connected to an IXP, it peers with all available peering providers[3] at that IXP. In order to decide which IXPs to choose, it starts with the cheapest transit provider (or set of transit providers).
  It then evaluates all IXPs, starting with the one that has the highest ratio of traffic volume of the peering providers at that IXP to the costs for connecting to that IXP. Connecting to a new IXP will reduce transit costs; if the saved transit costs are greater than the additional costs required to connect to the IXP and peer with the peering providers, it connects to that IXP, otherwise it does not.
- **Evolution Heuristic (H EV)**
  The evolution heuristic describes an evolutionary approach that could describe how a real INSP has found its interconnection partners over its lifetime:
  Go with the cheapest transit provider (or set of transit providers) first and connect to the best IXP – the one with the highest volume to cost ratio (see H PS). Then, successively evaluate the peering providers available at that IXP. Peer with a new peering provider if the saved transit costs are higher than the additional peering costs.
  In the second part of the heuristic, the other IXPs are evaluated with the following method: the INSP assumes it is connected to the new IXP and chooses its peering partners at that IXP in the same way as in the first part of the algorithm. Then, the INSP compares the total costs after connecting to this IXP with the total costs when not connecting and connects only if that reduces the costs.

10.1.2.3 Performance Evaluation

We first compare the solution obtained by our model (called *OPT*) with the solution obtained by the heuristics. The results are summarised in Figures 10.2 and 10.3 for $n = 100$ instances per scenario. Each of the algorithms solved the same 100 instances per scenario. The averages over the instances and the corresponding 95% confidence intervals are shown in the figures.

The costs of the heuristics (Figure 10.2 (a)) are measured relative to the total interconnection costs of the OPT algorithm – that is, the optimal costs. The H TR and H PA

[3] This heuristic mimics a behaviour that is used by some INSPs. These INSPs are sometimes called 'peering-sluts' because they have peering relationships with a lot of other providers at selected IXPs.

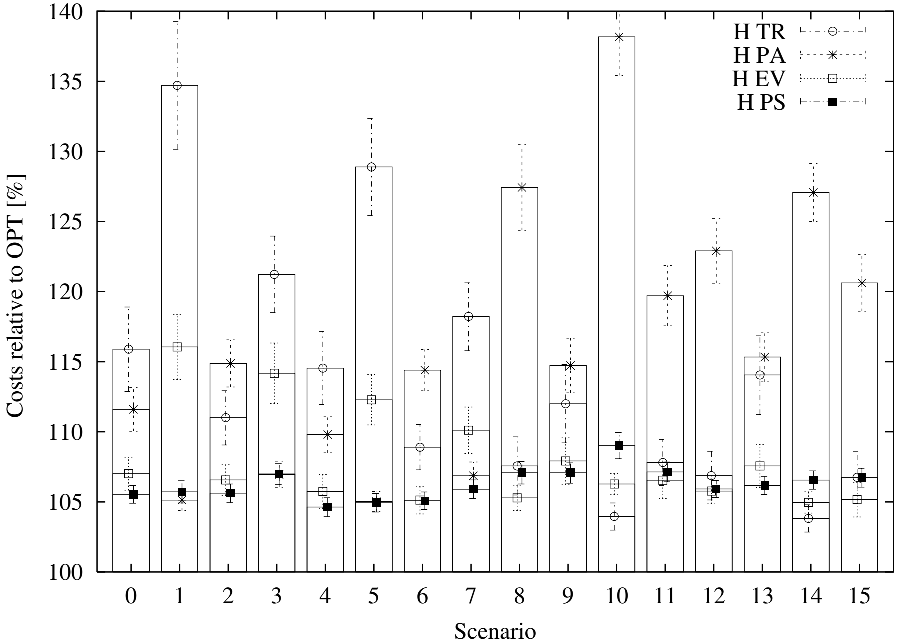

(a) Costs Relative to OPT

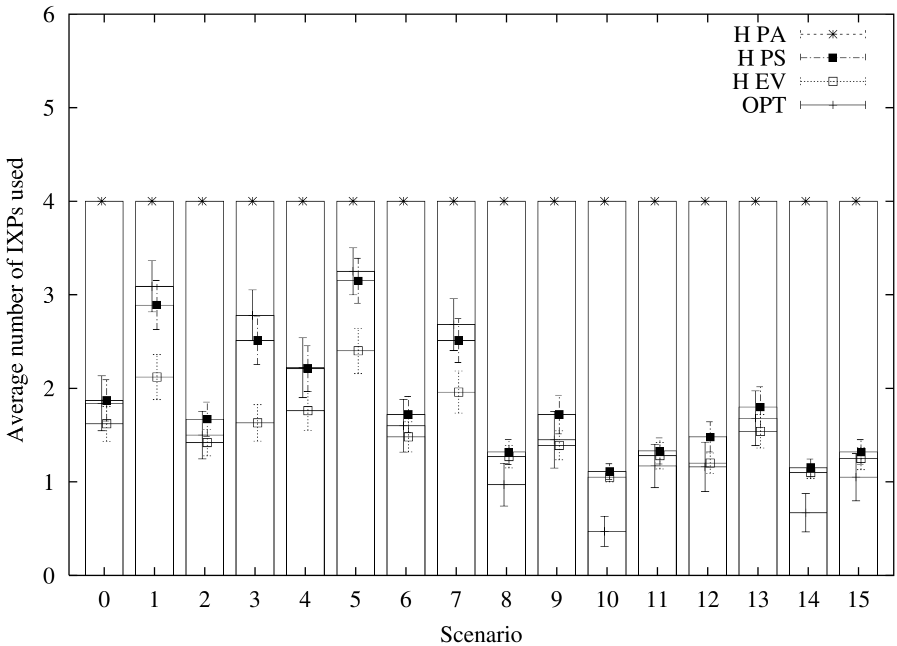

(b) Number of IXPs Used

Figure 10.2 Performance Evaluation

heuristics clearly perform worst and lead to up to 35% higher costs than the OPT algorithm. Choosing none or too many peering providers is obviously not a good idea: H PA performs systematically better than H TR for the higher transit costs (scenario 0–7) and vice versa. The best heuristics are H EV and H PS, with H PS performing better than H EV most of the time, but they still lead to 5 to 9% higher costs than the OPT algorithm. The explanation is shown in Figure 10.2 (b)[4] of Figure 10.2. In many cases, H PS chooses nearly the same number of IXPs as OPT. H EV systematically chooses too few for the higher transit costs (scenario 0–7) and thus misses some chances of saving costs by peering. It chooses too many for the lower transit costs (scenario 8–15) and thus misses out cheap transit opportunities in that case. When looking at the ratio of peering to transit providers (Figure 10.3 (a)) and peering to transit traffic (Figure 10.3 (b)), H EV systematically has too few and H PS has too many peering relationships.

The results show that the OPT algorithm presented in this chapter can save significant amounts of interconnection costs for all the different scenarios when compared with heuristics actually used in the real world. Remember that the interconnection costs are typically the largest cost factor for an INSP. The best real-world heuristic is the so-called peer-at-selected-IXPs heuristic (H PS); it peers with every possible partner at a number of carefully selected IXPs.

The next question we investigate is whether the computational complexity of the OPT algorithm might be an obstacle for using it rather than the heuristics.

10.1.2.4 Evaluation of Computational Complexity

If we define $M = \frac{1}{J} \cdot \sum_j M_j$, $R = \frac{1}{I} \cdot \sum_i \text{Size}(\Re_i)$, and $S = \frac{1}{N} \cdot \sum_n S_n$, then Model 10.1 needs $I(R+1) + J(2M+R) + NS$ variables and $I + J + 2JM + R + N(1+S)$ constraints.

The time it took to solve one problem instance of scenario 0 on a machine with a 700 MHz Pentium 3 and 256 MB RAM is depicted in Figure 10.4. The numbers of peering providers I and transit providers J were increased (x-axis) to increase the complexity of the problem. As Figure 10.4 shows, OPT can be solved in less than 10 minutes for large problems with 1100 providers. Given the fact that in the real world the problem has to be solved only rarely, the computational complexity is not an obstacle for using OPT.

A further advantage of OPT is that it is based upon a MIP model that can be further extended in different ways, as shown in the next sections. Some of these changes would be anything but straightforward to incorporate into the heuristics.

10.2 Reliability

Reliability is an important issue for INSPs. Model 10.1 can be extended in several ways to include reliability. By reliability, in this context we mean protection against the failure of one or more interconnections. For example, in all the 100 solved problem instances for the OPT algorithm, if the biggest provider selected from that strategy fails, there is not enough free capacity available from the other interconnected transit providers to compensate the

[4] The number of IXPs and the peering to transit ratio for H TR are zero; therefore H TR is not included in Figure 10.2 (b) and in Figure 10.3.

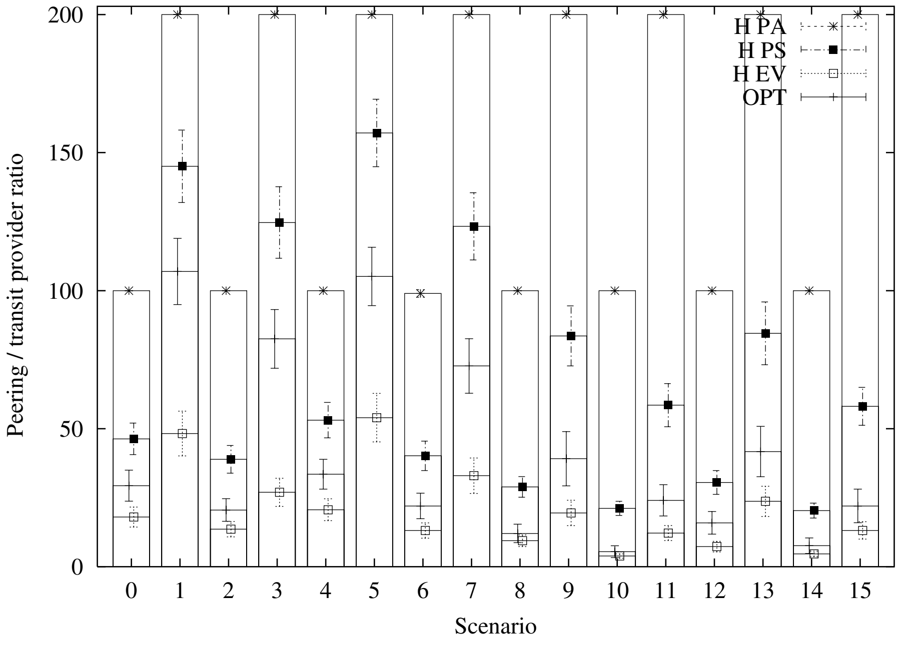

(a) Peering/Transit Provider Ratio

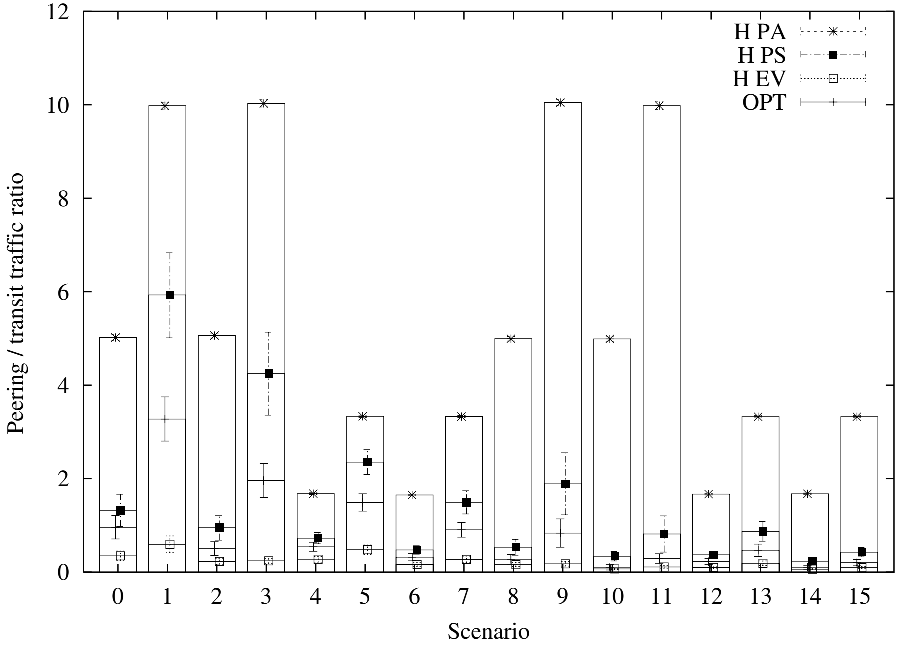

(b) Peering/Transit Traffic Ratio

Figure 10.3 Performance Evaluation

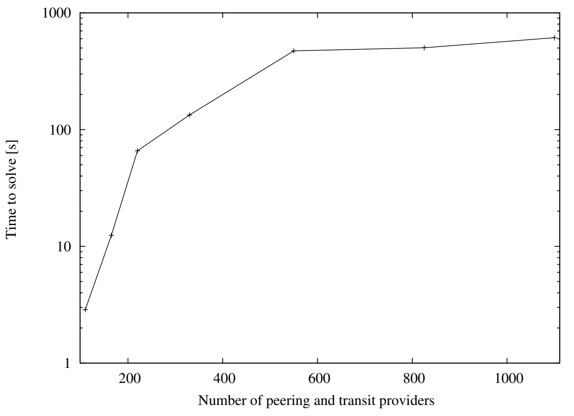

Figure 10.4 Evaluation of the Computational Complexity, Time to Solve for OPT

failure by rerouting the traffic destined for the failed provider. We therefore suggest and discuss several policies for extending the OPT strategy.

10.2.1 Policies

10.2.1.1 Minimum Number of Transit Providers Policy (MT)

One easy reliability policy is to interconnect with a minimum number of transit providers $\underline{Y}$ to reduce the dependency on each of them. This policy can be easily incorporated into the basic model (see Model 10.2). The advantage of this policy is its ease of use; the disadvantage is that it does not provide any guarantees or fine-grained control.

Model 10.2 Minimum Number of Transit Providers Policy (MT)

The following parameter and constraint are added to the otherwise unchanged Model 10.1:

Parameter

$\underline{Y}$ Minimum number of providers

Constraint

$$\sum_j y_{j1}^T \geq \underline{Y} \tag{10.15}$$

10.2.1.2 Minimum Free Capacity Policy (MC)

Another reliability policy is to make sure that there is a minimum amount of free transit capacity available, for example, a percentage of the total traffic. The free transit capacity is the sum of all capacities of the transit providers minus the capacities of the providers that are already being used; see Model 10.3.

Model 10.3 Minimum Free Capacity Policy (MC)

The following new parameter, variables and constraints are added to Model 10.1. Also, we now explicitly have to assume positive fixed costs for transit providers: $l_j^T > 0$.

Parameter

Γ Required amount of free capacity as fraction of the total traffic

Variables

f_j^T Free capacity of transit provider j

Constraints

$$f_j^T \leq c_j^T - \sum_{m \in M_j} \tilde{x}_{jm}^T \qquad \forall j \tag{10.16}$$

$$f_j^T \leq c_j^T \cdot y_{j1}^T \qquad \forall j \tag{10.17}$$

$$\sum_j f_j^T \geq \Gamma \cdot \sum \hat{x}_r \tag{10.18}$$

$$f_j^T \geq 0 \qquad \forall j \tag{10.19}$$

Constraint (10.16) limits variable f_j^T to the free capacity of transit provider j, (10.17) forces f_j^T to zero if there is no interconnection with transit provider j, (10.18) enforces the minimum amount of free capacity, and (10.19) is the nonnegativity constraint for the new variables.

This policy gives the decision maker fine-grained control over the free capacity. Its drawback is that if one interconnected provider who carries more than the fraction Γ of the traffic fails, there is not enough spare capacity. This is avoided by the next policy.

10.2.1.3 Anticipating Failure Policy (AF)

Another approach would be to make sure that there is enough spare transit capacity if a single transit or peering provider fails completely. This policy is described in Model 10.4.

Constraint (10.20) anticipates the failure of each transit provider j, (10.21) does the same for each peering provider i.

10.2.1.4 Combined MC and AF Policy (MCAF)

The MC and the AF policy can be combined as they use the same variables f_j^T; we name this approach MCAF.

Model 10.4 Anticipating Failure Policy (AF)

The following new variables and constraints are added to Model 10.1:

Variables

f_j^T Free capacity of transit provider j

Constraints (10.16), (10.17), (10.19) and

$$\sum_{k\,|\,k\neq j} f_k^T \geq \sum_m \tilde{x}_{jm}^T \qquad \forall j \qquad (10.20)$$

$$\sum_j f_j^T \geq \sum_{r\in\Re_i} x_{ir}^P \qquad \forall i \qquad (10.21)$$

10.2.2 Evaluation

In order to evaluate the reliability policies given above we use simulations again. The results presented here are based on scenario zero of Section 10.1 but there was no fundamental difference observed for the other scenarios.

In order to evaluate the reliability performance, we calculate the free transit capacity of the solutions obtained by the different policies as a percentage of the total traffic. The higher the free capacity, the more safety buffer remains if – for example – one provider fails. For each solution, we also determine whether there would be enough free capacity to carry the traffic of the biggest (peering or transit) provider, if it fails; we call this measure *robustness*. The average results and the 95% confidence intervals are depicted in Figure 10.5, as are the average costs of the solutions obtained by the different policies. The graphs also contain the reference reliability and cost measures of the solutions obtained for the same problems by the unmodified OPT algorithm above (0% robustness, 32% free capacity).

Again, we generated n=100 instances that were solved by the Minimum Number of Transit Providers Policy (MT), the Minimum Free Capacity Policy (MC), and by the combination of the Minimum Free Capacity Policy and the Anticipating Failure Policy (MCAF). The results for the AF Policy (AF) alone are included in the results for MCAF when the minimum free capacity is 0%.

If we look at MT, which has a parameter that can only be increased in integer steps, it can be seen that the costs increase very quickly if the minimum number of transit providers is increased. It is important to remember that the reference costs (100%) refer to the total interconnection costs, and not only the transit costs; they are therefore very high in absolute terms. The cost increases of the MC and MCAF policies are much smoother and more controlled.

If we analyse the reliability measures, the robustness increases quickly for MT and more slowly for MC; MCAF automatically leads to full robustness because of the AF constraints. The free capacity explodes for the MT policy while it is obviously more controlled with the MC policy, because the minimal free capacity is a parameter of that

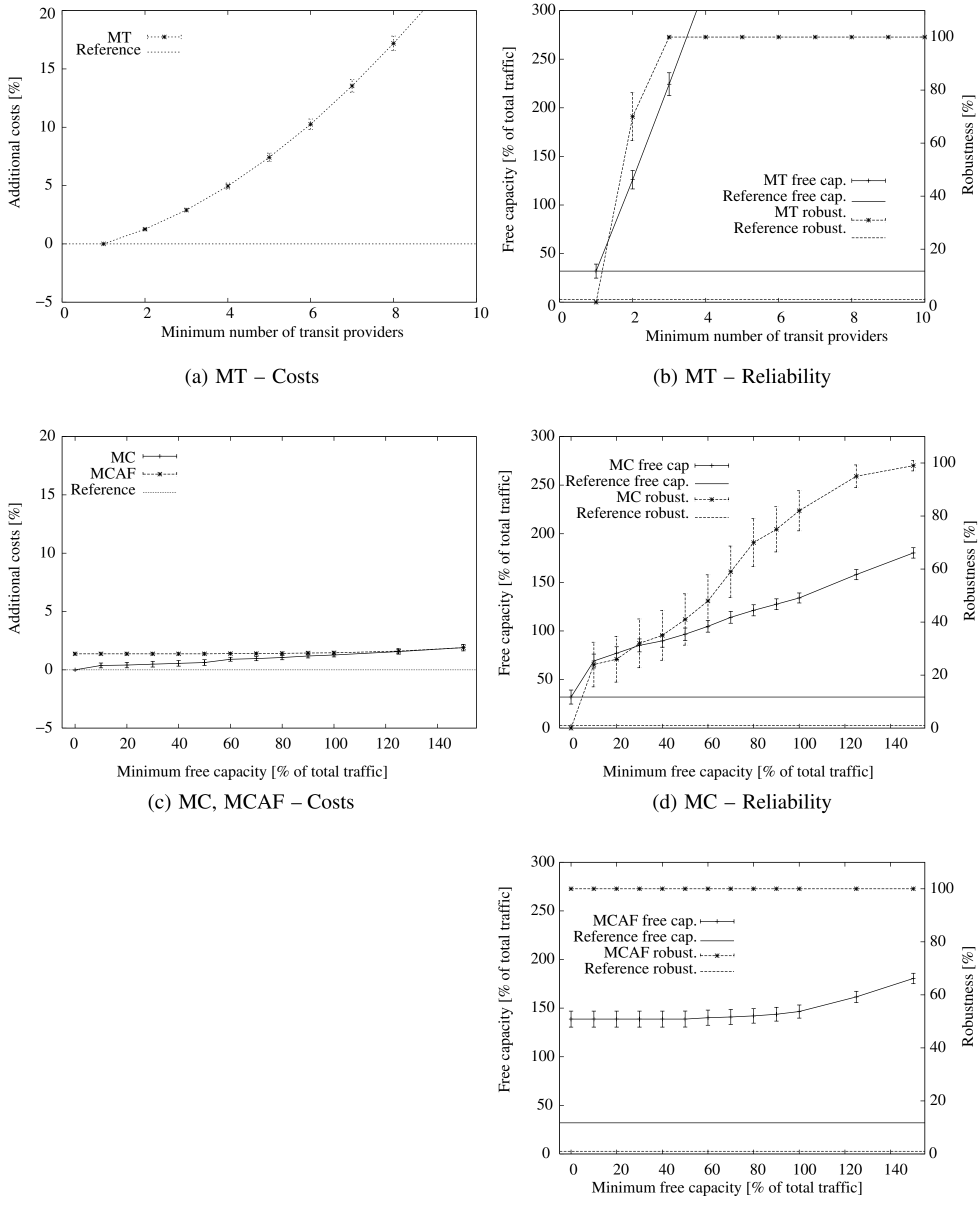

(a) MT – Costs

(b) MT – Reliability

(c) MC, MCAF – Costs

(d) MC – Reliability

(e) MCAF – Reliability

Figure 10.5 Evaluation of the Reliability Policies

policy. Because of the AF constraints in MCAF, the free capacity does not decrease for lower values of the minimum free capacity parameter.

The MT policy represents an easy rule of thumb (using a minimum number of transit providers) as it can be expected to be used by some INSPs. It can be used to increase the reliability. However, costs can explode, especially as the parameter (the number of transit providers) cannot be increased smoothly but only in steps of one. Since the policy parameter only indirectly influences reliability metrics such as free capacity and robustness, the MT policy cannot be recommended. The more sophisticated approaches developed here (MC and MCAF) perform better compared to MT. Contrary to MT, which increases free capacity by adding new transit providers, MC and MCAF can also increase the free capacity by choosing more peering providers and *bigger* providers instead of just *more* transit providers.

The MCAF policy is clearly the best choice: it offers full robustness and full control over the free capacity. Its parameter is the minimal free capacity, which can be easily estimated by the decision maker. If the failure of the biggest provider is unlikely, MC can also be used. It can save some costs compared to MCAF, but only for low values of free capacity.

10.3 Quality of Service

One of the typical parameters that an INSP would like to optimise is the QoS achievable with its interconnections. In the context of interconnections and with the information available at that abstraction level[5], the QoS can be mainly influenced by selecting interconnections so that the length of routes in terms of AS (autonomous systems) hops is kept low. In addition, peering or transit providers could be rated in some fashion with respect to the QoS they usually offer, and the solution could take those ratings into account. We will focus on the more objective measure of route lengths and show several possibilities for extending the basic model (Model 10.1) to include the QoS that is achieved by the interconnection policy chosen. These extensions can be directly combined with those discussed in Section 10.2.

10.3.1 Policies

The typical QoS metric used on the timescale of interconnections is the average number of AS hops for a route from the provider's network to the end-point. A lower number of hops correlate with lower delay and a lower loss probability for the packets, and thus a higher utility for the end-user. This is especially important for routes carrying traffic from real-time multimedia applications and network games. Peering interconnections usually offer a lower hop count than transit interconnections, because the traffic ends in the peering network. This is, in fact, the main reason why some larger INSPs accept peering with significantly smaller INSPs.

10.3.1.1 Peering Bonus (PB)

The easiest way of taking the lower hop count of peering providers into account is giving peering providers with QoS sensitive routes a bonus b_i that reduces their fixed peering costs and thus makes peering with them more attractive; see Model 10.5.

[5] e.g. obtained from BGP data.

Model 10.5 Peering Bonus (PB)

The parameter l_i^P of the basic Model 10.1 is replaced with the new parameter $\tilde{l}_i^P = l_i^P - b_i$.

The advantage of this approach is its ease of use; the disadvantage is that the parameter b_i can be difficult to estimate, and only indirectly influences the QoS.

10.3.1.2 Hop Constraint (HC)

Another approach that gives the decision maker more control of the QoS parameter hop count is to add an additional constraint for the average hop count of the traffic; see Model 10.6.

Parameter q_r is used as a weight when determining the average hop count of the traffic. Routes known to carry delay-sensitive traffic (e.g. to gaming sites) should obtain a higher than average q_r. Using the parameters of Model 10.6, the weighted average AS hop count is

$$H = \frac{1}{\sum_r q_r \hat{x}_r} \left(\sum_i h_i^P \cdot \sum_{r \in \Re_i} q_r x_{ir}^P + \sum_j h_j^T \cdot \sum_r q_r x_{jr}^T \right) \tag{10.22}$$

It is limited to $\overline{H}$ by Constraint (10.23) of Model 10.6. Not only the AS hop count, but also any other QoS metric can be modelled with this approach. Instead of looking at the entire traffic, this approach can be easily modified to take into account only a subset of the routes. For a more fine-grained prediction of the hop count for transit providers, h_j^T could be replaced by a route-dependent prediction h_{jr}^T for route r through the network of transit provider j.

The advantage of the Hop Constraint (HC) approach is that it gives the decision maker a finer control and, with the maximum hop count, an easy to understand design parameter. The disadvantages are the higher number of parameters and the slightly higher complexity of the optimisation model with the additional constraint.

10.3.1.3 Hop Count Penalty Costs Policy (HP)

Decreasing the hop count can quickly lead to increasing costs (as shown below). The HC policy enforces a maximal hop count without constraining the cost-increase. The hop count penalty costs policy (HP, Model 10.7) is similar but does not enforce a maximum hop count with a constraint. Instead, it adds the hop count, weighted with some Penalty Costs (PC), to the target function. This allows a trade-off between decreasing the hop count (which typically leads to increasing costs, as we will see in the evaluation below) and decreasing the costs.

10.3.2 Evaluation

In order to evaluate the QoS approaches, we use simulations based on scenario 0 again; the results observed for the other scenarios were not fundamentally different. The hop

Model 10.6 Hop Constraint (HC)

The following parameters and constraint are added to Model 10.1:

Parameters

h_i^P Average hop count for traffic through peering provider i (typical value is one)

h_j^T Estimation of the expected hop count for traffic through transit provider j

q_r Delay sensitivity of the traffic on route r used as weight for the average hop count

$\overline{H}$ Allowed maximal average hop count allowed

Constraints

$$\sum_i h_i^P \cdot \sum_{r \in \Re_i} q_r x_{ir}^P + \sum_j h_j^T \cdot \sum_r q_r x_{jr}^T \leq \overline{H} \cdot \sum_r q_r \hat{x}_r \tag{10.23}$$

Model 10.7 Hop Count Penalty Costs Policy (HP)

Using the parameters h_i^P, h_j^T, q_r from Model 10.6, (10.22) is added to the target function (10.1) of the otherwise unchanged Model 10.1 weighted with penalty costs ψ:

$$\text{Minimise (10.1)} \; + \frac{\psi}{\sum_r q_r \hat{x}_r} \left(\sum_i h_i^P \cdot \sum_{r \in \Re_i} q_r x_{ir}^P + \sum_j h_j^T \cdot \sum_r q_r x_{jr}^T \right) \tag{10.24}$$

count for peering providers is set to 1, and for the transit providers, it is drawn uniformly distributed from the interval [3.0, 6.0].

The averages of $n = 100$ problem instances and the 95% confidence intervals are shown for the Peering Bonus (PB), HC, and HP policies in Figure 10.6. For reference purposes, the costs and the hop count from the plain Model 10.1, without any QoS features, are also depicted (labelled ‘reference’).

Figure 10.6 shows that all policies can decrease the hop count. At the same time, the costs increase. The costs for a low hop count are higher when using the PB policy than either of the other two. This occurs because the PB policy distorts the costs of the providers with the PB and minimises the distorted rather than the real costs. Also, the decision maker cannot easily guess the optimal parameter of the PB policy, therefore, this policy cannot be recommended.

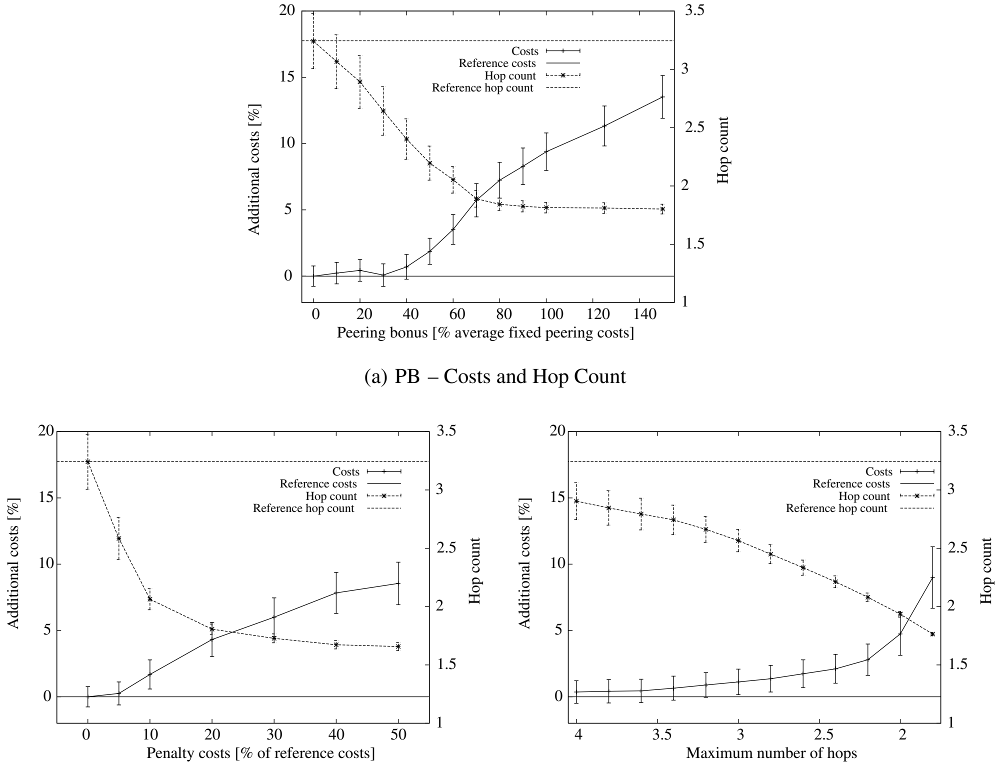

(a) PB – Costs and Hop Count

(b) HP – Costs and Hop Count

(c) HC – Costs and Hop Count

Figure 10.6 Evaluation of the QoS Strategies

The HC offers direct control over the hop count, which the other policies do not. Its parameter is therefore easy to set – if the decision maker has an indication of what the hop count should be. The costs increase more quickly for lower hop counts. The HP policy does not enforce a certain hop count but, instead, evaluates the value of the decreased hop count (expressed by the PC) against the hop count. This avoids the danger of exploding costs in case the maximum hop count of the HC has been set too low.

If a certain maximum hop count is absolutely necessary, the HC policy has to be used. Otherwise, if there is flexibility on the hop count, HP offers the best way of modelling this trade-off. HC and HP can also be combined: HC could be used to ensure that a certain (higher) hop count is not exceeded, while HP could be used to further decrease the hop count, without ignoring the cost increase.

10.4 Environment Changes

The models of Sections 10.1 to 10.3 can be used to calculate the optimal set of peering and transit providers for one INSP at one point in time. This is useful for a new INSP entering the market. An INSP that already has interconnections with a number of peering and transit providers faces a slightly different problem: *is the current set of peering and transit providers still optimal or is it worth changing interconnections, considering the technical and administrative costs for establishing a new or cancelling an existing interconnection?*

We call this the *dynamic* problem and now show that the previous models can be easily extended for the dynamic case. Again, the models are evaluated by simulations.

10.4.1 Adjusting the Basic Models

For the dynamic case, we now assume that there are interconnections to a set Θ of the I peering providers, to a set θ of the J transit providers and to a set ϑ of the N IXPs. As the traffic requirements and the cost functions of the providers change, the dynamic problem is solved every period in order to find the new optimal set of providers.

Typically some technical and administrative effort is necessary to establish a new interconnection that can be expressed by a cost term (transaction costs). Cancelling an existing interconnection also typically involves some effort that can be expressed by a cost term.

10.4.1.1 Penalty Costs Policy (PC)

The costs for establishing a new interconnection can be expressed as penalty costs per period by dividing them by the number of periods an interconnection is expected to last, or by a typical amortisation or planning horizon. These penalty costs can be added to the fixed costs of the providers that are not in set Θ or θ. Similarly, the costs for cancelling an existing interconnection can be transformed into bonus costs per period that are subtracted from the fixed costs for the providers in set Θ and θ respectively. The same can be done for the IXPs. This provides an incentive to stick with the current set of providers and IXPs; we call this the PC policy; see Model 10.8.

The advantage of this policy is that the basic models are easily extended this way and the cost terms involved can typically be estimated quickly and easily.

10.4.1.2 Limiting Change Policy (LC)

Another policy for dealing with the dynamic problem would be to limit the amount of change (new interconnections and cancelled interconnections) per period, reflecting the limited technical capacities for these changes in a period, or the risk of change the provider is ready to take. We call this policy Limiting Change (LC) policy; see Model 10.9.

Constraint (10.28) limits the permissible number of changes. The left-hand side of constraint (10.28) counts the binary y-variables that are 1 if an interconnection to provider *i/j* is made for all providers *i/j* with which no previous interconnection agreement existed. It adds all cancellations of interconnection agreements by counting the zeroes in the binary y-variables of the providers *i/j* with which an interconnection agreement existed during the last period.

10.4.2 Evaluation

For the simulative evaluation, we create $n = 25$ problem instances. To simulate the dynamic environment we simulate p periods per instance. At the beginning of each period,

Model 10.8 Penalty Costs Policy (PC)

Model 10.1 is extended as follows:

Parameters

Θ	Set of peering providers that an interconnection exists with at the beginning of the current period
$s_i^P \ \forall i \notin \Theta$	(Per period) penalty costs for establishing a new interconnection with peering provider i
$b_i^P \ \forall i \in \Theta$	(Per period) bonus for not cancelling an existing interconnection with peering provider i
θ	Set of transit providers that an interconnection exists with at the beginning of the current period
$s_j^T \ \forall j \notin \theta$	(Per period) penalty costs for establishing a new interconnection with transit provider j
$b_j^T \ \forall j \in \theta$	(Per period) bonus for not cancelling an existing interconnection with transit provider j
ϑ	Set of IXPs connected with at the beginning of the current period
$s_n^{IXP} \ \forall n \notin \vartheta$	(Per period) penalty costs for connecting to the new IXP n
$b_n^{IXP} \ \forall n \in \vartheta$	(Per period) bonus for disconnecting from IXP n

Parameters l_i^P, l_j^T and l_{ns}^{IXP} are replaced by $\tilde{l}_i^P$, $\tilde{l}_j^T$ and $\tilde{l}_{ns}^{IXP}$ which are defined as follows:

$$\tilde{l}_i^P = l_i^P + s_i^P \ \forall i \notin \Theta \text{ and } \tilde{l}_i^P = l_i^P - b_i^P \ \forall i \in \Theta \tag{10.25}$$

$$\tilde{l}_j^T = l_j^T + s_j^T \ \forall j \notin \theta \text{ and } \tilde{l}_j^T = l_j^T - b_j^T \ \forall j \in \theta \tag{10.26}$$

$$\tilde{l}_{ns}^{IXP} = l_{ns}^{IXP} + s_n^{IXP} \ \forall s \ \forall n \notin \vartheta \text{ and } \tilde{l}_{ns}^{IXP} = l_{ns}^{IXP} - b_n^{IXP} \ \forall s \ \forall n \in \Theta \tag{10.27}$$

the amount of traffic, the capacity of the providers and the fixed and variable costs vary. The range of the changes is shown in Tables 10.3 and 10.4. As in Section 10.1, we analyse different scenarios where either the first or second option from Table 10.4 is used. If option *All Providers Available at Beginning* is used, all the providers are available for an interconnection agreement at period 0; the only change in this simulation is the

Model 10.9 Limiting Change Policy (LC)

The following parameters and constraints are added to Model 10.1:

Parameters

Θ, θ See above

W Maximum allowed number of new and cancelled interconnections per period

Additional Constraint

$$\sum_{j\in\theta}(1-y^T_{j\,1})+\sum_{i\in\Theta}(1-y^P_i)+\sum_{j\notin\theta}y^T_{j\,1}+\sum_{i\notin\Theta}y^P_i \leq W \tag{10.28}$$

Table 10.3 Constant Parameters

Parameter Description	Interval
Growth of traffic per route per period	[15%, 25%]
Growth of capacity per period	[15%, 25%]

Table 10.4 Scenario-dependent Parameters

Bit	Parameter Description	Bit = 0	Bit = 1
1	Number of periods p	20	10
2	Change of each of the following cost terms per period: Fixed peering costs, fixed transit costs, variable costs, IXP costs	[−20%, +5%]	[−10%, 0%]
3	All providers available at beginning	Yes	No

traffic, capacity and cost change. If this option is not chosen, 25% of the providers are not available in period 0 and become available in a random period of the simulation (each period having the same probability). We now first evaluate the dependency of the results of each policy on the parameters of the policy for scenario 7 and then compare all of the policies for each scenario.

10.4.2.1 Dependency on Policy Parameters

We start by analysing the average number of changed interconnections and the probability of a period without any changes. These change metrics are depicted in Figure 10.7 (b)

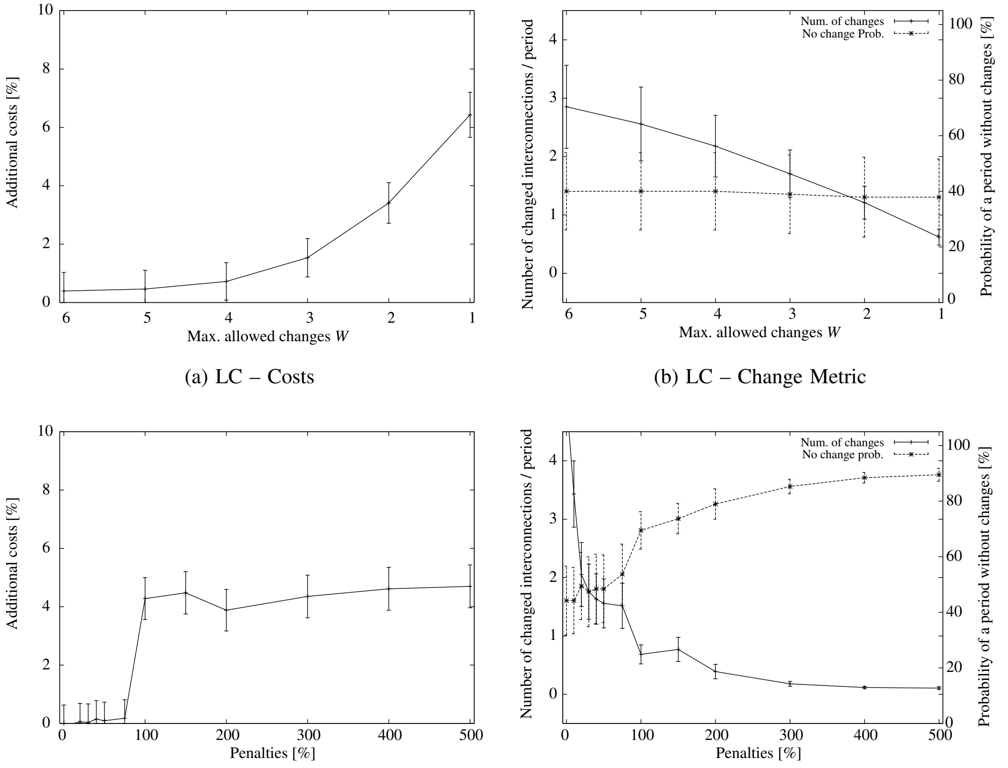

(a) LC – Costs

(b) LC – Change Metric

(c) PC – Costs

(d) PC – Change Metric

Figure 10.7 Evaluation of the Dynamic Strategies, Dependency on the Policy Parameters

for different parameters W that limit the number of permissible changes per period for the *LC policy*. The average of $n = 25$ problem instances and the 95% confidence interval are shown. We can see that the probability that no change occurs in a period remains low, independent of W. The LC policy allows a number of changes in each period and thus uniformly distributes the amount of change over all the periods. This leads to the low probability evident in the figure. The number of changed interconnections per period obviously decreases with W. The additional costs of LC relative to the results for $W = \infty$ are shown in Figure 10.7 (a). They increase by only 6% if W is decreased from 6 to 1.

For the *PC policy*, the PC were calculated as a constant percentage of the fixed peering: transit costs for establishing a new or cancelling an existing interconnection. The same procedure was followed for the IXPs. For PC of up to 100%, the probability that no change occurs increases while, at the same time, the number of changes per period decreases. At the same time, the costs increase slightly. This is a nice result; the PC policy can influence the amount of change better than the LC. However, for very high PC above 100% the amount of change is only slightly decreased. This is reasonable since the amount of change seen for very high PC is the change that is absolutely necessary, such as choosing a new transit provider because traffic demand exceeds the capacity of the existing interconnections. The PC approach is reasonable and practical. If there are

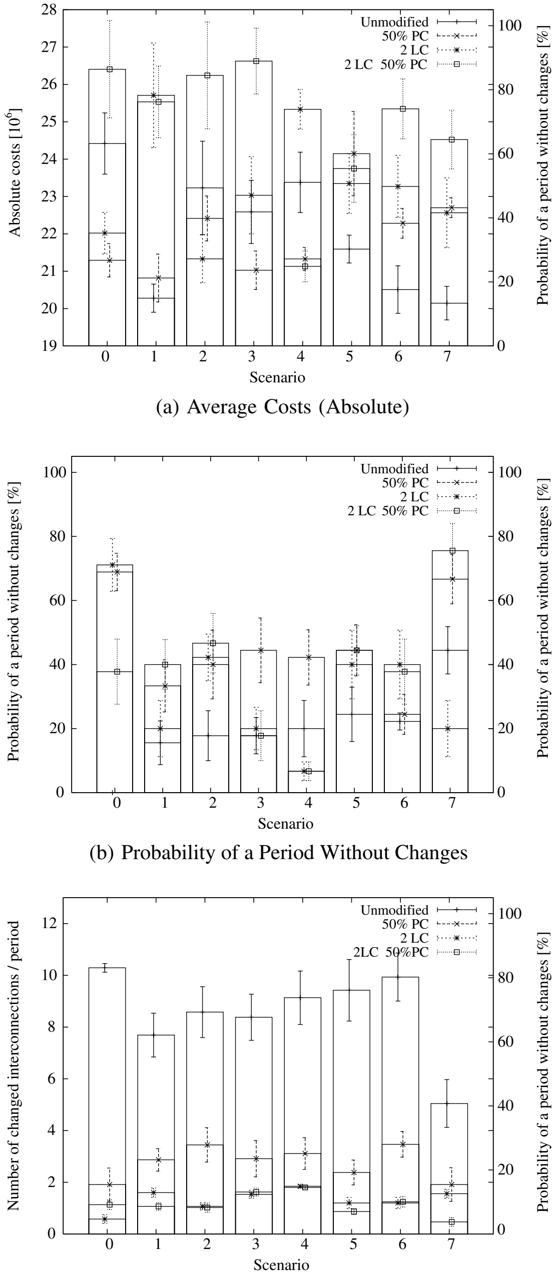

(a) Average Costs (Absolute)

(b) Probability of a Period Without Changes

(c) Number of Changed Interconnections

Figure 10.8 Evaluation of the Dynamic Strategies, Dependency on the Scenario technical limitations restricting the amount of change, LC can be used in addition to PC. We now compare the approaches for the range of scenarios.

10.4.2.2 Evaluation of the Different Scenarios

The results for the different scenarios are shown in Figure 10.8 for the unmodified OPT Model 10.1, the PC policy with 50% penalty costs, the LC policy with $W = 2$ and the combination of PC and LC. The costs can differ by up to 32% between the policies. No one policy leads to clearly lower or higher costs than any other in all scenarios. The combined policy and LC lead to the fewest number of changes. The unmodified algorithm does not control change and thus leads to the highest change rate. PC and the combined policies lead to the highest probability of not having to change an interconnection in one period. To conclude, we can recommend using the combination of PC and LC, as it provides the most robust policy. However, the results also show that the policy and the parameter of the chosen policy strongly depend on the scenario.

10.5 Summary and Conclusions

The interconnection-related costs form one of the highest cost factors for an INSP and are therefore highly important for the efficiency of a network. In this chapter, the optimisation potential with respect to interconnection-related costs was evaluated. Several optimisation models for interconnections between providers were presented. We have shown how to find the most efficient set of peering and transit partners for a provider. Simulations show that our approach is superior to typical real-world heuristic approaches. They also show that, of the heuristics, the one performs best that connects to all possible peering partners at an IXP and chooses the optimal set of IXPs. However, the presented exact algorithm can still save 5% more total interconnection costs. Furthermore, it can be easily extended to take other issues like reliability and QoS into account.

Besides efficiency, the interconnection mix can significantly influence the achievable QoS, for example, via the AS hop count. We derived and analysed strategies for optimising both QoS and efficiency. We have also presented and discussed several ways of extending the basic strategies to take reliability issues into account. In the last part of the chapter, we have shown how to extend the basic models to the dynamic problem situation by evaluating whether a given set of peering and transit partners is still optimal, considering changes in the traffic mix or cost structure of the involved providers. We have also considered the administrative costs of changing peering and transit partners and evaluated different approaches in simulations.