

Part I

Introduction and Basics

1

Introduction

1.1 Motivation

The Internet is a large network formed out of more than 30000 autonomous systems (AS), in which each AS is a collection of IP networks sharing a common routing strategy. These networks are operated by thousands of **Internet service providers** (ISPs). On the one hand, the ISPs compete with each other for customers and traffic, on the other, they have to cooperate and exchange traffic, otherwise the worldwide connectivity would be lost. In contrast with the traditional telecommunication markets, there are almost no central instances in the Internet enforcing cooperation and regulation of the market.

The ISP market is characterised by serious competition and is currently in a phase of consolidation: according to Access ECommerce (University of Minnesota Extension Service) in the three years after the dot.com crash in 2000, at least 962 Internet companies have either shut down or have declared bankruptcy. In 2002 WorldCom, one of the largest tier-1 ISPs, filed for bankruptcy protection, the largest such filing in US history. A year ago, in the US Internet access market, the top 10 providers accounted for less than half of the Internet users; today the market is consolidated more so that the top 10 providers account for almost three out of four Internet users, see Boardwatch (2004) (i). The situation was aggravated by peer-to-peer applications like Napster, Gnutella and Kazaa that led to extreme traffic growths, causing additional costs and challenges for access providers.

Never before had ISPs to be as competitive as they have to be today. The goal of this book is to help ISPs to be more competitive. The focus of this book is on network operations. It is highly important for ISPs to operate their network **efficiently**. In addition, they have to strive for successful business practices. Traditional successful business practices in a competitive environment are cost leadership, market segmentation and differentiation, see Porter (1980). Market segmentation and differentiation depend on measures to offer different products to different markets. A central service of ISPs consists of forwarding IP packets; this service can be differentiated by price and quality. Therefore, besides efficiency it is important to investigate the **quality of service** (QoS) for ISPs too. The latter is important for a second reason: many emerging multimedia applications such as voice-over Internet Protocol (VoIP) and video communication can greatly benefit from QoS support in a network, see Bhatti and Crowcroft (2000); Crowcroft *et al.* (1999); Steinmetz and Nahrstedt (2004). QoS support in a network therefore opens further possibilities

for value-added services with which providers can differentiate themselves and target new markets.

To summarise the motivation of this book and to be competitive, an ISP has to operate its network efficiently and it must also control the level of QoS offered to its customers. Efficiency and QoS are discussed in more detail before we make an overview and a short summary of the different chapters of this book are given.

1.2 Efficiency and Quality of Service

1.2.1 Network Efficiency

Merriam-Webster defines *efficiency* as an *effective operation as measured by a comparison of production with cost (as in energy, time, and money)*. In the context of network services, the “production” of an ISP’s network can be described by the amount of traffic transported by the ISP. Therefore, we define the efficiency of a network as

$$\text{Network Efficiency} = \frac{\text{Transported Traffic}}{\text{Costs}}$$

Depending on the level of abstraction the traffic can be measured by

- the volume of traffic carried through the network, or
- the number of flows or sessions transported through the network, or
- the number of customers served.

The network costs can be monetary or non-monetary. Typical monetary cost factors for an ISP are

- costs for leasing communication lines,
- interconnection fees,
- costs of the network hardware (e.g. routers, switches and line-cards), and
- costs for the technical and administrative staff.

The examples for non-monetary costs are

- the complexity in computation time or the memory of managing and scheduling a packet,
- the amount of state necessary in a network to provide a certain QoS, or
- the technical effort of changing resource allocations.

In this book, many optimisation problems are presented and solved to maximise efficiency. In many circumstances, one can assume that the amount of traffic is given and constant on the timescale of the investigated problem. In that case, the efficiency is maximised if the costs are minimised.

1.2.2 Network Quality of Service

Quality of service (QoS) is defined in Schmitt (2001) as

the well-defined and controllable behaviour of a system with respect to quantitative parameters.

Typical QoS parameters on the network layer are packet loss, packet delay, jitter (delay variation), throughput etc. Different applications have different QoS requirements: real-time applications are more sensitive to delay and jitter, while elastic bulk-transfer applications are relatively insensitive to the delay and jitter of individual packets but are sensitive to the overall achievable throughput.

For this reason, measuring the QoS of a specific flow or session directly with technical parameters like loss or delay can be misleading. Preferably, *utility functions* should be used; they transform the technical QoS parameters and traffic flow experiences, into a utility value depending on the requirements of that flow's application. Examples for this can be found throughout this book.

1.2.3 Trade-off between Efficiency and Quality of Service

If one looks at certain aspects of an ISP (e.g. the interconnection mix or the QoS system), a **trade-off** between the QoS and the network efficiency can be observed. This trade-off is depicted in Figure 1.1 in which the grey area depicts the **solution space** marking the points where feasible solutions exist for an ISP. Consider for example, interconnections – the connections of the network of one ISP with other networks. Typically, many different possibilities for an individual interconnection exist, resulting in a larger number of possible interconnection combinations as a provider typically has interconnections with many other providers. These interconnection combinations differ in their costs, efficiency and in the QoS they support (see Part III of this book for details); the solution space is formed by all feasible interconnection combinations. Offering a very high QoS usually leads to a lower network efficiency because either less traffic can be supported to provide the QoS or the costs for handling the traffic increase. The same holds true the other way around, leading to the shape of the solution space depicted in Figure 1.1.

It is important to stress that the solution space only contains *feasible* solutions. In the example above, if a specific interconnection mix violates the requirements of the ISP with respect to other criteria – e.g. security – it is not considered feasible and therefore not a part of the solution space. This book takes the position of an ISP and optimises the efficiency and QoS of its network, taking the trade-off between the two goals into account.

The **optimal performance boundary** is marked by the upper right border of the solution space (see Figure 1.1). As long as an ISP does not operate at the optimal performance boundary, it can improve either QoS or efficiency without having to reduce the other goal. It is clear that the goal of a competitive ISP is to operate at the optimal performance boundary. A major contribution of this book is that it investigates how the optimal performance boundary can be found for different aspects from building, operating and managing a network.




Figure 1.1 Trade-off between Network Efficiency and Quality of Service

A different question is where along that boundary should an ISP operate. Internet is a heterogeneous network of networks; the ISPs that control these networks differ from very small regional niche providers to huge multinational backbone providers. Where along the optimal performance boundary an individual ISP should operate, is basically its own decision. And that decision, of course, depends on the requirements of the market and the ISP's customers. The customer requirements and market structures are constantly developing. This book does not make any limiting assumptions in this respect as, for example, only looking at one point of the performance boundary (e.g. providers with high QoS but also high costs) this would make its results only applicable to a small subset of all ISPs. Even for those ISPs, the results would only be valid until the next change occurs in a fast evolving business like the Internet service provisioning business.

Instead, the book strives for generic solutions, if possible. It tries to investigate the entire boundary. For example, a large variety of different QoS systems are evaluated in this book, each with distinct advantages and disadvantages in which none of the systems is clearly better than another. They rather represent different points on the performance boundary. Which of these systems a provider should employ, depends on the specific situation it is in: its customers, its financial situation and many more factors .

1.3 Action Space and Approach

The next important point to discuss is the action space investigated in this book. The book focuses on *technical* operations, not on marketing or other measures. Of all technical measures, the focus is on building, running and interconnecting the *network* of an ISP. In Chapter 2, the term *INSP* for *Internet Network Service Provider* is therefore introduced to describe the entity responsible for these actions. The different actions can be grouped into the following three areas (see also Figure 1.2):

- **Network Architecture**

The network architecture defines the properties of the INSP's network itself. Four sub-architectures can be distinguished:

- the quality of service architecture,




Figure 1.2 The Different Parts of this Book

- the forwarding architecture,
- the signalling architecture, and
- the security architecture.

For the purpose of this book, the most important part of the network architecture is the QoS architecture as it directly determines the QoS and considerably affects the other sub-architectures.

- **Interconnection**

The network of an INSP has to be connected to other INSPs' networks; the connection between two INSPs' networks is called an *interconnection*. There are different types of interconnections and typically many possible interconnection partners. The interconnection mix strongly influences the amount of traffic an INSP can carry, its costs and the achievable QoS, as this book will show.

- **Traffic and Network Engineering**

Apart from operating a network on the basis of a selected network architecture and interconnecting it with other networks, the INSP constantly has to manage its own network: the INSP can use *traffic engineering* methods, e.g. to avoid bottlenecks by rerouting traffic and on a larger timescale it can use *network engineering* to update the topology and upgrade the capacity of the network.

Solutions obtained for one goal or in one category while ignoring the other goals or categories respectively can easily result in an overall suboptimal system: a QoS solution that offers good QoS but would result in unacceptable (monetary or non-monetary) costs or could not cope with the required amount of traffic would be inefficient and probably useless. Also, if the network architecture is highly optimised for efficiency and QoS, that advantage would be easily lost if the interconnections are not optimised for the same goal. Further, it would be lost after a few months if the network engineering process fails to upgrade the network to increasing traffic. To avoid this, a system-oriented view on building, operating and managing a network is necessary.

1.4 Overview

The book is structured along three areas: network architecture, interconnection and traffic and network engineering. Every single problem area has a large influence on the overall QoS and efficiency of the provider as the book will show. From this it follows that considering only one of these aspects – e.g. the network architecture – is not enough because the gain of e.g. QoS by carefully optimising the network architecture is immediately lost if the interconnection mix is not adapted to support this level of QoS. Also, it is lost soon if the capacity expansion process (part of the traffic and network engineering problem area) fails.

Part I presents the introduction and an overview of the ISP market. In addition, different basic methods for network performance analysis are presented. They are employed in the rest of the book to analyse the different facets of networks. The Internet protocol stack and the most important protocols are discussed. Finally, the most important applications for the present and the future of Internet are discussed with focus on the QoS requirements and traffic behaviour.

Part II investigates the network architecture, starting with a discussion of the state-of-the-art in Chapter 6. In this part the focus lies on the QoS architecture as it is the foundation for the QoS achievable in a network and also because it largely influences the other aspects of the network architecture and thus indirectly the architectural costs and efficiency. Using analytical methods, two aspects of QoS architecture (admission control and service differentiation) are investigated in Chapter 7. For both the aspects, the over-provisioning factors are derived with different analytical methods. An over-provisioning factor is the relation of capacity (mainly bandwidth) between a plain best-effort system and a QoS system at the point in which both systems offer the same QoS. It captures the benefit of the QoS system.

The benefit of admission control largely depends on the adaptivity of the application and the load distribution. In a well-dimensioned network, the over-provisioning factor is usually less than 300%; for adaptive applications it is significantly even smaller than 150%.

We derive a novel network model that – contrary to the existing approaches – allows us to analyse service differentiation. The over-provisioning factors resulting from service differentiation are significantly higher, typically between 200% and 500%, depending on the traffic assumptions.

Different QoS systems based on the QoS architectures that are in the standardisation process of the Internet Engineering Task Force (IETF) are analysed in a simulative study (Chapter 8). The study sheds light on the quantitative trade-offs of the different approaches to QoS, e.g. per-flow versus per-class scheduling and central versus decentral admission control. One of the conclusions of this chapter is that Diffserv networks can be over booked by at least a factor of three to increase efficiency. The over-provisioning factors determined in this chapter are similar to those determined with our novel analytical models for the service differentiation in the previous chapter. The book also shows that contrary to common belief, the utilisation of a network with Expedited Forwarding (EF) traffic can be higher than a few per cent. For the basic experiment in Chapter 8, the Charny bound predicts a maximal utilisation of 7.98%. A bandwidth broker described in this book can raise the utilisation to over 27%.

In **Part III**, interconnections are investigated starting with an overview and the discussion of the state-of-the-art in Chapter 9. Chapter 10 shows that the interconnection mix has significant impact on efficiency and QoS. Reliability is important in this context too and therefore is also discussed. Different strategies for optimising the interconnection mix with respect to efficiency are presented and evaluated by a series of simulations.

Interconnection related costs are one of the highest cost factors of an INSP. The results show that from 5% to more than 30% interconnection related costs could be saved. The analysed strategies can be easily extended to control reliability and QoS too. Chapter 10 presents strategies that allow to explicitly adjust the desired trade-off between QoS and efficiency. The result is the optimal interconnection mix for an INSP.

Traffic and network engineering is discussed in **Part IV** of this book, starting with the discussion of the state-of-the-art in Chapter 11. Several traffic-engineering strategies are presented in Chapter 12. The chapter shows that the maximum utilisation criterion, which is typically used in related work, is not a good objective function for traffic engineering. The chapter derives a congestion function that should be used instead as an objective function. The impact of traffic engineering on the network efficiency and QoS is evaluated in a series of experiments. The results show that traffic engineering can decrease the congestion in a network during times of high load. A traffic-engineered network can therefore offer higher QoS and/or higher efficiency (because more traffic could be carried than a network without traffic engineering). The absolute benefit of traffic engineering, however, strongly depends on the traffic-engineering strategy. On the basis of the experiments, the book gives recommendations of which strategies to use and also how to use them. However, for certain topologies and traffic distributions, the benefit of traffic engineering can be rather small. In this case, it is very doubtful whether a traffic-engineering solution can amortise its costs.

Network engineering with focus on capacity expansion is finally discussed in Chapter 13. Capacity expansion is an important and a frequent task in today's IP networks because the traffic volume is increasing steadily. First, the influence of capacity expansion on the performance of the different QoS architectures of Chapter 8 is discussed. If capacity is abundant, the differences between the QoS architectures diminish. If capacity is scarce, the systems with a strict admission control manage to maintain QoS while the other systems suffer to different extents. The best effort systems are the ones most sensible to in-time capacity expansions.

Besides this, different capacity expansion algorithms are presented and evaluated in Chapter 13. The best strategy takes the mutual influences of traffic engineering and capacity expansion into account. The rule of thumb often used by today's INSPs shows acceptable performance only if their parameters are set correctly; our strategy performs significantly better and is robust against uncertain traffic predictions. Finally in that chapter, the effects of elastic TCP traffic-on-traffic matrices and on capacity expansion are discussed with some analytical models. The elasticity of the traffic influences the capacity expansion measures if the network is highly utilised before the expansion. The presented models can be used to predict this effect and react accordingly.

To understand the terminology used in the book, we recommend every reader to have a look at Chapter 2. Readers who are technical experts in IP network technology and

performance analysis can probably skip the introductory Chapters 3 and 4, possibly also 5. Readers who are interested in the basics should focus on Chapters 2 to 5 and the first chapters of each part (6, 9 and 11). Moreover, in Chapters 6, 9 and 11 most of the related scientific works are summarised. Finally, readers who are already very familiar with the topics and who look for results might want to focus on the chapters describing the experiments (7, 8, 10, 12 and 13).