

9

Interconnections Overview

The community believes that the goal is connectivity

RFC 1958 (see Carpenter (1996))

The Internet consists of a variety of interconnected heterogeneous networks (autonomous systems, AS¹), managed by multiple independent INSPs. Despite the competitive INSP market, each INSP must interoperate with its neighbouring Internet networks to provide efficient connectivity and end-to-end service. No INSP can operate in complete isolation from others; therefore, every INSP must not only coexist with other INSPs but also cooperate with them. Contrary to the situation in most telecommunication markets, there is no central authority in the Internet that enforces cooperation.

Both the number of networks and ASes as well as the average number of ASes a given AS is connected to are increasing at a fairly high rate. The number of ASes rose from 909 in 9/95 to 4427 in 12/98, 7563 in 10/00 to over 30,000 in 2004; see CAIDA – Cooperative Association for Internet Data Analysis (2004); Fang and Peterson (1999); IANA (Internet Assigned Numbers Authority). Similarly, the average interconnection degree, that is, the number of providers a certain provider has interconnection agreements with, rose from 2.99 in 9/95 to 4.12 in 12/98. It is also notable that a single provider may interconnect with up to 1000 other providers; see Fang and Peterson (1999).

Considering this, and the fact that the highest cost factors of Internet Service Providers (ISPs) are, typically, the interconnection costs and line costs, it is obviously important to study the effect of the interconnections on the network efficiency and Quality of Service (QoS) of an INSP.

In this book, we define a *network edge* as a connection between two different networks; there are two types of network edges:

- **Homing** describes the connection of an end-user and End-User Network Operator (ENO) to an Access Internet Service Provider (AISP) network.
- An **interconnection** is the connection between the networks of two different AISPs/Backbone Service Providers (BSPs).

¹ An autonomous system (AS) is a group of IP networks operated by one or more Internet Network Service Providers (INSP(s)), which has a single and clearly defined exterior routing policy.

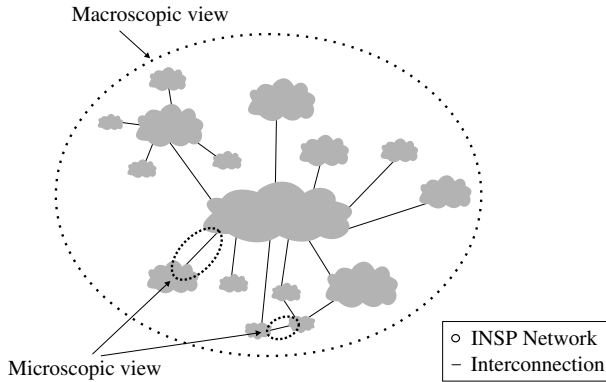


Figure 9.1 Macroscopic and Microscopic View on Interconnections

In this and the next chapter, we focus on the latter type of network edge: interconnections. There are two different ways of looking at interconnections as shown in Figure 9.1. The macroscopic view focuses on the large-scale connection structure of many networks as a whole, while the microscopic view analyses a single interconnection.

We start with a macroscopic view on interconnections in Section 9.1. In Section 9.2, we look at individual interconnections (microscopic view). Peering and transit interconnections are elaborated in that context; they form the basis of the analysis in the next chapter.

One important aspect of an interconnection is the interconnection method. Different methods are discussed in Section 9.3.

Real INSPs almost always use a mix of different (peering and transit) interconnections as discussed in Section 9.4 and further analysed in the following chapter.

9.1 A Macroscopic View on Interconnections

There are many different ways to connect a given set of networks with each other. The two extreme structures are the strictly hierarchical and the fully meshed structures as shown in Figure 9.2. The Internet is a heterogenous network of networks and follows neither of these two structures. However, as aspects of both the structures can be found in real connection structures (see e.g. Huston (1999a)) and as they are often referenced in literature, we investigate them first. Towards the end of this section, we look at empirical results about the real interconnection structure.

9.1.1 Strictly Hierarchical Structure

A strictly hierarchical structure, also called *tier structure*, consists of a small number of global INSPs at the ‘top’ that are referenced as *tier 1 INSPs*. Kende (2000) specifies five tier 1 providers in 2000, also called the *Big Five*: Cable&Wireless, WorldCom, Sprint, AT&T, and Genuity². These few large backbones interconnect solely by peering and do

² Genuity is now a member of Level 3. WorldCom filed the largest bankruptcy in the US history in 2002.

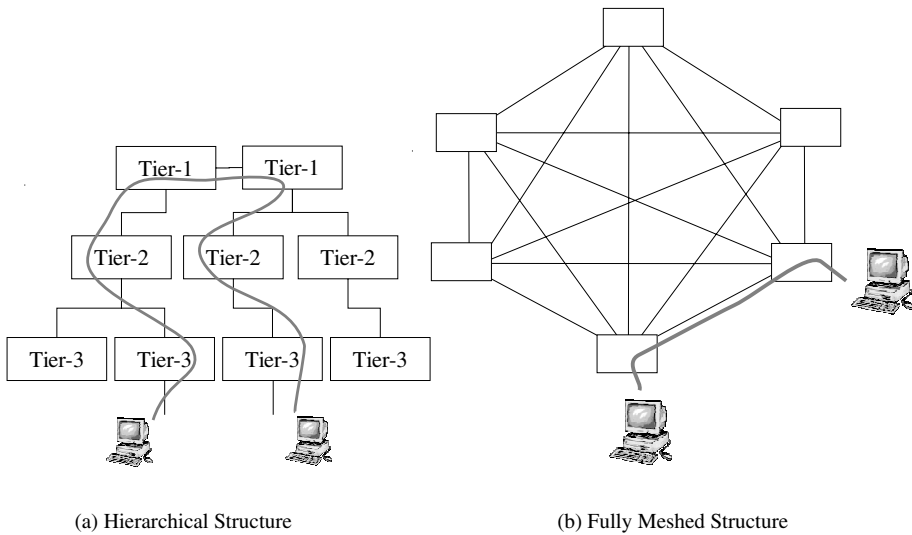


Figure 9.2 Archetypical Structures

not need to purchase transit from any other backbone; they incorporate a pure BSP role. The *tier 2* consists of national INSPs; they have a smaller presence than tier 1 INSPs and may lease part of the network structure of tier 1 INSPs. The INSPs that are considered as tier 2 can be big AISP like America Online (AOL) as well as national BSPs like Deutsches Forschungsnetz (DFN). Local AISP are considered as *tier 3* INSPs. At each tier, the INSPs are clients of the tier above, see Figure 9.2 (a).

The microscopic interconnection relation between two INSPs of different tiers in that structure is typically the *classical transit relation* that we describe and analyse in Section 9.2.

If the hierarchical model is strictly enforced, the traffic between two local INSPs may need to transit all the way through a tier 1 provider. Such extended paths are inefficient, because they generate extended transfer delays and increased costs. In a competitive market like the INSP market, there is strong pressure to reduce costs, which explains why the reality does not match the strictly hierarchical structure. A typical modification is a local interconnection between two neighbouring tier 3 or tier 2 INSPs. However, a benefit of the hierarchical structure is the relatively small number of interconnections needed for each INSP to establish end-to-end connectivity.

9.1.2 Fully Meshed Structure

The other extreme is the fully meshed structure of Figure 9.2 (b) which shortens the path length compared to that of the hierarchical structure. The transmission in such an environment is fast, because the AS-level distance is always one hop³.

However, the fully meshed structure does have obvious scaling issues if the number N of providers interconnected in that way becomes large, because the total number of interconnections is $N \cdot (N - 1)/2$.

³ Of course, one AS system level hop can consist of many IP level hops.

The microscopic type of interconnection in the fully meshed structure is typically a *classical peering relation* as we define and analyse it in Section 9.2.

9.1.3 Realistic Structures

The two structures above are archetypical structures that do not represent the true structure of the Internet as is shown by a number of works discussing the properties of the Internet structure, for example, Aiello *et al.* (2001); Bu and Towsley (2002); CAIDA – Cooperative Association for Internet Data Analysis (2004); Chen *et al.* (2002); Faloutsos *et al.* (1999); Medina *et al.* (2000); Palmer and Steffan (2000); Spring *et al.* (2002); Tangmunarunkit *et al.* (2001); Zegura *et al.* (1997).

In Faloutsos *et al.* (1999), power law relationships are found in three inter-domain (AS-level) topologies of the Internet, which were constructed from Border Gateway Protocol (BGP) data. This paper started a discussion on power law AS-level topologies: Medina *et al.* (2000) investigate, on the basis of the work of Barabasi and Albert (1999), possible origins of these power laws using topology generators to create artificial topologies. Aiello *et al.* (2001); Bu and Towsley (2002); Palmer and Steffan (2000) are based on the power law relationship.

However, Chen *et al.* (2002) show that during the process of constructing the topologies of Faloutsos *et al.* (1999) from BGP data, 20% to 50% of the physical links are missed and that more exact topology graphs do not follow the power law relationship found in Faloutsos *et al.* (1999). The authors also show that works based on Barabasi and Albert (1999), for example, Medina *et al.* (2000), are not supported by the more exact topologies.

A nice macroscopic visualisation of the Internet, based on measurements by CAIDA, is shown in CAIDA – Cooperative Association for Internet Data Analysis (2003).

Besides that, there are a number of different topology generators, for example, BRITE (2004), Tiers (2004), Georgia Tech Internetwork Topology Models (2004), INET (2004) that can be used to generate artificial topologies that are deemed realistic by their authors. An evaluation of topology generators with respect to power law AS-level graphs is presented in Tangmunarunkit *et al.* (2001). A node in an AS level graph represents one AS and a link, an interconnection. The AS level graph is thus a graph representing what we call the macroscopic view in interconnections, see Figure 9.1. A similar study but for the topologies of INSP networks (one node representing a POP) is presented in Heckmann *et al.* (2003). Li *et al.* (2004) present an innovative new approach to understanding the structure of INSP networks. Contrary to the previous works that focus mainly on graph theoretic properties of topologies (e.g. the node-degree distribution), Li *et al.* (2004) take into account in their study the basic technological and economical trade-offs⁴ that network designers face. The authors show that topologies that have the same graph theoretic properties (e.g. node-degree distribution) can have very different throughput performance. They further show that high-performance topologies are not likely obtained by any random graph generation method.

Spring *et al.* (2002) present the tool ‘Rocketfuel’ for measuring router-level topologies based on traceroutes, and BGP and DNS data: Using publicly available traceroute servers,

⁴ A technological constraint is for example the bandwidth over degree function of actual switches/routers as it is determined by the cross-connection fabric. Economical considerations show that the costs of wiring can dominate the infrastructure costs, which gives a practical incentive to wiring networks such that they can support traffic using the least number of links.

the topology of a network can be revealed. Rocketfuel uses BGP data to calculate those traceroutes that most likely traverse the target network, at the same time redundant traceroutes are discarded. DNS information is finally used to cluster the IP addresses of the router interfaces to routers.

9.2 A Microscopic View on Interconnections

For our interest of investigating interconnections from the point of view of a single INSP, the microscopic structure of the interconnections is very important because it directly influences the QoS, cost structure, and transmission capacity of an INSP. Also, the microscopic structure of these interconnections – the mix of different interconnection types – is finally the decision of the INSP.

The literature typically distinguishes only two types of interconnections; see for example, Kende (2000); McGarty (2002); Songhurst (2001); Weiss and Shin (2002). This is not enough. The Internet service market is the outcome of business and technology interaction, rather than a planned outcome of some regulatory process. This leads to the appearance of a wide and diverse variety of interconnection types in reality. Therefore, we start by deriving a more detailed definition and classification of interconnections than typically found in related works.

9.2.1 Taxonomy and Classification of Interconnections

Huston (1999a,b) divides interconnections into physical and financial interactions. He describes the different possible connections. The routing entries that are exchanged at interconnections are called ‘the currency of interconnection’. Also, different financial settlement options of the telephony industry (bilateral settlement, sender keeps all and transit fees) and possible options for the Internet industry (e.g. per packet or session accounting) are discussed. Section 9.3 of this chapter is based on Huston’s physical interaction.

A more economical focused point of view can be found in Bailey (1997). It analyses the different economic incentives associated with different types of Internet interconnection arrangements; it does not consider regulatory issues. Friedmann and Mills-Scofield (1997), from a purely economic perspective, examine the optimal settlement pricing strategies for INSPs.

Kende (2000) gives insights into the market development of interconnection and the two most dominant interconnection forms: classical peering and transit. Kende (2000) also examines interconnection policies and regulatory issues as well as international interconnection.

Looking at the interconnections observed in the real world and learning from the works cited above, we find that the following three aspects comprehensively classify the variety of the existing interconnections (see Figure 9.3):

- **Route Advertisement**

The route advertisement can be symmetrical or asymmetrical. *Symmetrical* in this context means that both clients exchange their own and their direct customers’ routes. *Asymmetrical* in this context means that one INSP (the upstream provider) offers the other INSP (the downstream provider) access to all destinations in its routing table

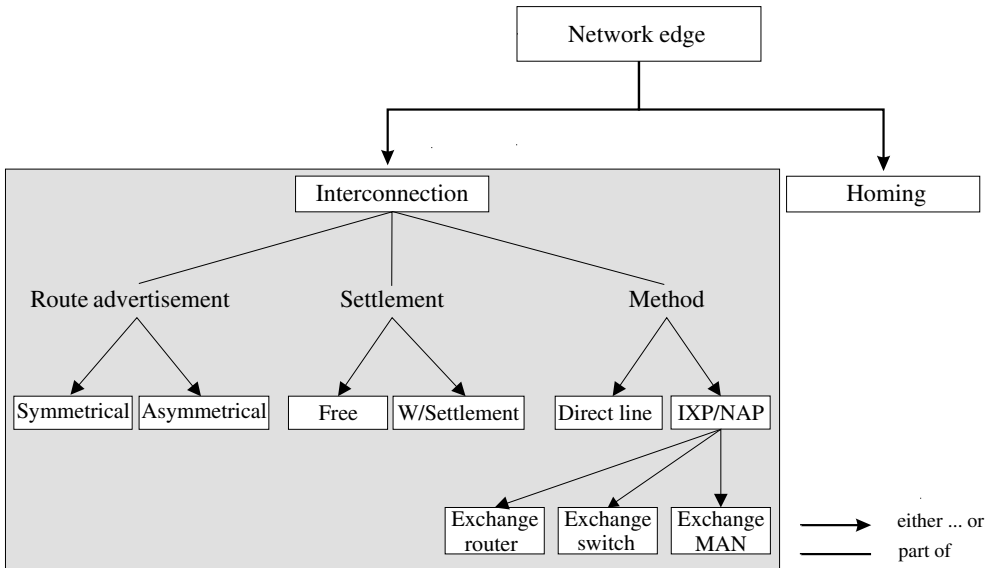


Figure 9.3 Network Edges and Interconnection Types

and advertises the downstream provider’s network entries in its routing table, while the downstream provider only advertises its own and its direct customer’s networks.

• **Settlement**

The settlement aspect of an interconnection is concerned with whether one INSP compensates the other for the exchanged traffic. In *settlement-free* interconnections, the providers might share the costs for the connection method but do not pay each other for the traffic itself, while in interconnections *with settlement*, one provider pays the other for the exchanged traffic. The price usually depends on the volume of the exchanged traffic and typically decreases with the volume (concave cost function), see Norton (2002).

• **Interconnection Method**

The interconnection method describes how the physical interconnection between the two providers is realised: Either through one or more direct connections between the two providers’ networks (*direct-line* method), or through an Internet Exchange Point (*IXP*). An Internet exchange point is typically used by a larger number of INSPs that are connected to

- a central router (exchange router structure),
- a central switch (exchange switch structure) – also called *exchange Local Area Network, LAN* – or
- a Metropolitan Area Network, MAN (exchange MAN structure, distributed exchange).

The exchange switch and exchange MAN (distributed exchange) methods are typically found in large IXPs (like LINX, DE-CIX, Parix). The exchange router method is not very common because the routing between the INSPs is performed by the central router and managed by a central instance, leaving the INSPs with too little influence on the routing and exchange policies.

The settlement, route advertisement, and interconnection methods can be freely combined to $2^3 = 8$ combinations⁵. Two out of these eight combinations make up most of the interconnections currently found in the Internet. We call them *classical peering* and *classical transit* interconnections.

9.2.2 Peering

At the beginning of the commercial Internet, interconnection agreements evolved from the informal interactions that characterised the Internet at the time the National Science Foundation (NSF) was running the backbone. The commercial backbones developed a system of interconnections known as *peering*. Although the term ‘peering’ is used frequently, it rarely has a uniform meaning. There is no set definition for it in the Internet Engineering Task Force (IETF). RFC 1983 (see Malkin (1996)), which provides definitions to important Internet-related terms, has no entry for ‘peering’ or ‘interconnection’.

We use the term *classical peering* for the most common form of an interconnection relationship that treats both INSPs more or less equally:

In *classical peering*, two INSPs use settlement-free symmetrical route advertisement and interconnect at an IXP.

Therefore, classical peering has the following distinctive characteristics:

1. Peering INSPs exchange traffic on a settlement-free basis; they do not charge each other for the transfer volume between them as in a transit relationship.
2. Peering INSPs use symmetrical route advertisement. They exchange traffic that originates with the customer of one INSP and terminates with the customer of the other peered INSP. To enable this, they exchange their own and their direct customer’s routes. As part of the classical peering arrangement, an INSP would not, however, act as an intermediary and accept the traffic of a peering INSP and transit that to another connected INSP.
3. The classical peering interconnection exchanges traffic via an IXP. The peering INSPs have to own or lease lines to the access point of the IXP. The connection to the IXP gives the connecting INSP also access to a wide number of other possible peering and transit partners.

Initially, most peering traffic took place at IXPs as it was efficient for each INSP to interconnect with as many INSPs as possible at the same location. The rapid growth in the Internet traffic caused the IXPs to eventually become congested; see for example, Robertson (1997). This led to the situation that some INSPs avoided IXPs and peered directly with each other. This kind of peering differs from classical peering by using a direct-line interconnection method instead of an IXP and is known as *private peering*.

Badasyan and Chakrabarti (2003) describe a game-theoretic model in which INSPs decide on private versus classical peering agreements as a multistage game; their result is that a mixed approach of connection via private peering and classical peering has the most advantages.

⁵ If we ignore for the moment that there are different realisations of IXPs. They are discussed in Section 9.3 and Appendix C.

Nowadays, the congestion problem at IXPs seems to be solved. There are recent indications that a large proportion of the Internet traffic is again exchanged via classical peering interconnections. For example, Boardwatch (2003, viii) state that 90% of UK Internet traffic is routed through the IXP LINX in London and that LINX provides access through its memberships to around 50% of the world's Internet networks.

Another atypical peering interconnection type is *peering with settlement*. It has all the properties of the classical peering arrangement as described above with the sole exception that one INSP receives a financial compensation from the other INSP, usually because the peered traffic is unbalanced in favour of the second INSP.

9.2.3 Transit

Because each peering arrangement only allows INSPs to exchange traffic destined for each other's customers, INSPs would need to peer to a significant number of other INSPs to gain access to the full Internet. One alternative to classical peering is the classical transit interconnection:

In a *classical transit* interconnection, the two INSPs can be clearly distinguished as a *customer INSP* and a *transit INSP*. They are sometimes also called *downstream and upstream INSPs*. Asymmetrical route advertisement is used, the customer INSP pays the transit INSP for the exchanged traffic (settlement) and a direct-line connection is used.

The main differences between classical transit and classical peering are thus:

- In a transit interconnection, one INSP pays another INSP for the exchanged traffic; the amount of settlement typically depends on the exchanged traffic volume.
- The transit INSP advertises the customer INSP's routing table entries and routes its traffic to all its peering and transit partners, thus connecting the customer INSP to 'the rest of the world'.
- The customer INSP, on the other hand, only advertises its own routes and thus only receives traffic from the transit INSP that ends in its own network.
- Transit agreements often include Service Level Agreements (SLAs); see below.

There are several *non-classical transit-like* interconnections imaginable. Sometimes, for example, an IXP could be used instead of a direct-line. However, not all IXPs allow these type of agreements over their infrastructure.

9.2.4 Service Level Agreements

Service Level Agreements (SLAs) are bilateral contracts at a network edge between an INSP and a customer that can be either another INSP or an end-user. RFC 3198 defines a SLA as *the documented result of a negotiation between a customer/consumer and a provider of a service, that specifies the levels of availability, serviceability, performance, operation, or other attributes of the service*; see Westerinen *et al.* (2001). SLAs are typically used in transit-like interconnections agreements. They contain a *Service*

Level Specification (SLS). A SLS is a set of parameters and their values, which together define the service offered to the customer. Besides the SLS, a SLA can contain pricing, contractual and other information.

An example for a SLA is the one MCI/UUNET is offering for Internet services, see MCI (2004):

- A 100% *network availability* is promised. For each cumulative hour of network unavailability or fraction thereof, the customer is credited one day of charges.
- Within the United States, *latency* guarantees of 55 ms between MCI's inter-regional transit backbone routers (hubs) are guaranteed; transatlantic latency guarantees of 95 ms are guaranteed. Also specified in the SLA is how the latency is measured: by averaging sample measurements taken during a calendar month.
- Packet *delivery* guarantees of at least 99.5% are given between hub routers. Again, this is a monthly average. The credit is one day of the MCI monthly fee.
- MCI notifies customers by email or pager within 15 minutes after it is determined that their *service* is *unavailable*. Unavailability is assumed if the edge router does not respond after two consecutive 5-minute ping cycles.
- A scheduled *maintenance* notification is specified to reach the customer 48 hours in advance; maintenance is performed during a standard maintenance window as also specified by the SLA.
- A response time of maximal 15 minutes is guaranteed for *denial of service* attacks.
- If acts of God, embargoes, terrorism, fires, sabotage, and so on, are responsible for the breach of the SLA no credit has to be given.

Several commercial SLA management solutions exist and are deployed, e.g. CiscoWorks (2004) and Lucent (2004).

SLAs are also an integral part of the Differentiated Services QoS architecture (see Black *et al.* (1998)). In that context, they are discussed in Section 6.2.4.2.

9.3 Interconnection Method

9.3.1 Internet Exchange Points

Since the introduction of the four original Network Access Points (NAPs), aka IXPs, in the NSF-proposed post-NSFNET architecture in 1995, the IXP market has developed significantly. More IXPs have emerged all over the world, enabling local as well as global interconnection.

Europe's leading IXPs have set up the European Internet Exchange (EURO-IX) Association with LINX in London, United Kingdom and AMS-IX in Amsterdam, Netherlands as the biggest members (see EURO-IX (2004)). As mentioned above, more than 90% of UK's Internet traffic is routed through the LINX exchange, Europe's largest IXP, which provides access through its memberships to around 50% of the world's Internet networks, see Boardwatch (2003, viii).

INSPs (especially AISPs) may connect to more than one IXP to ensure better connectivity and gain access to more peering partners. 47% of the INSPs connected to an IXP in the EURO-IX Association are also connected to at least one other IXP within




Figure 9.4 IXP Network Structure

the EURO-IX Association. Besides connecting via IXPs, AISPs typically also have direct interconnections with other INSPs. The resulting network structure evolving around IXPs is depicted in Figure 9.4.

A problem with respect to IXPs is that small regional providers that join an IXP will have the same reach as large INSPs that have invested in the national backbones. However, small regional networks will have smaller sunk costs that might lead to prices lower than that of the national providers. Cost recovery may become a significant issue for the larger providers. However, if the national backbones cannot recover cost and go out of business, the small regional INSPs lose the connectivity their customers want.

Because of this problem, membership to IXPs is sometimes restricted to national backbones, or small regional INSPs are required to have settlements with the larger INSPs. Both approaches can be found in the IXP market. All IXPs in the EURO-IX Association have the prerequisite that candidates for membership shall have an ASN (Autonomous System Number), see e.g. AMS-IX Membership (2004); German Internet Exchange DE-CIX (2004); LINX (2003). The DE-CIX IXP in Frankfurt, Germany, requires its members to peer with at least 10 other DE-CIX members after 6 months of membership, see German Internet Exchange DE-CIX (2004).

9.3.2 Evaluation

If we analyse the connection structures that are dominated by peering-like interconnections, it is quite intuitive that using IXPs can save costs compared to a connection structure that purely relies on a larger number of direct connections as shown in Figure 9.2 (b). A simple analytical model in Appendix C sheds some light on this intuition and shows that already for a very few number of providers within a city, a central exchange point is more cost efficient than direct connections.

There are different types of IXP structures. Our classification of interconnections shown in Figure 9.3 and Huston (1998) distinguish between the exchange router, exchange LAN, and exchange MAN (distributed exchange) structures. The first IXPs were built using the exchange router structure. Nowadays, this structure is not used anymore owing to its technical limitations. A simple analytical comparison of the exchange MAN and LAN structures is also given in Appendix C and suggests that the LAN structure is more cost efficient with a relatively small number of connected INSPs. The MAN structure gains cost advantages with a rising number of INSPs. However, these are just theoretical structures and in reality the distinction between MAN and LAN structures is fluid. The most dominant structure for bigger IXPs like LINX and DE-CIX is the combination of both structures; they use multiple colocation facilities containing a LAN (exchange switch) that are connected via a MAN.

9.4 Interconnection Mix

Almost all providers use a combination of different interconnection types. Most INSPs make use of an interconnection mix as they peer with other local INSPs and transit with at least one BSP to ensure global connectivity.

9.4.1 Negotiation Process

An interconnection arrangement is based on a negotiation process that results in the general framework for the interconnection. Usually all parties want to act as transit INSP, as it is preferable to be paid for an interconnection. The decision on which party takes on the role of a transit INSP and which one the role of a customer INSP is not easy to make. Generally, the closer the INSPs are, considering their size, their customer base, and their infrastructure, the more difficult is the negotiation process and the more likely is a peering arrangement. The negotiation over being the transit/customer INSP is often based on the greater geographical coverage criterion. However, this factor is not the only possible criterion, as one INSP may host valuable content and argue that the access to this content adds value to the other INSPs network. Also, an INSP with a very large client population with a limited geographical coverage may argue that this large client base offsets other possible criteria.

Huston (1999b) describes this negotiation as *two animals meeting in the jungle. Each animal sees only the eyes of the other, and from this limited input they must determine which animal should attempt to eat the other!*. After deciding upon which one will be the upstream INSP, the remaining parts of the contract have to be defined. The fees have to be decided upon as well as the location and the number of exchange points.

If the INSPs cannot solve the problem, the INSP may settle on a peering arrangement. Peering has some appeal to the INSPs as they do not need to track the exchanged traffic volume constantly, like in a transit relationship where the payment is typically based on the exchanged traffic volume. The tracking generates cost and the parties have to consider if transit is worth the effort. In conclusion, it can be said that peering is sustainable under the assumption of costly, unnecessary traffic measuring and mutual benefits.

There are several approaches in literature to model and solve the interconnection decision. For example, Giovannetti *et al.* (2003) identified some criteria to peer in an empirical study of the INSPs connected to the Milan IXP:

1. **Size.** The established peering points, either an IXP or private peering, entail fixed and variable technological costs. The result is that a sufficiently intense traffic flow between the customers of the two INSPs is needed for peering to be economically viable. The larger the two networks are, the more intense the traffic flow will be. WorldCom⁶, for instance, published their criteria for peering in 2001. One criterion is that the traffic volume at the peering points is at least 150 Mbps.
2. **Symmetry.** Since the cost for the peering points are usually shared equally by the two peering INSPs, unbalanced traffic implies unbalanced gain from peering against a balanced distribution of costs. Such unbalanced situations have led to discontinuation of peering agreements and to its replacement with a transit interconnection. One of the criteria published by WorldCom is that the peering network has the geographical scope of at least 50% of its own. Another criterion is that the exchanged traffic volume at peering points does not exceed 1:1.5.
3. **Quality of Service.** The quality of a connection between two end-users depends crucially on the most congested network on the connection path. To ensure a certain degree of quality and to curb a potential free-riding⁷ on infrastructure investments, the last criteria WorldCom sets forth is that most of the peering network has a capacity of 622 Mbps.

The study also shows that the peering decision is influenced, for example, by the proximity of the INSP's headquarters and their distance to the IXP. Other types of works for determining the interconnection mix are discussed next.

9.4.2 Determining the Interconnection Mix

There are two basic types of work that model the decision of an INSP on which interconnection type or mix to choose: game-theoretic and decision theoretic works. In decision theoretic works, the optimal decision of one INSP is analysed under a *ceteris paribus* constraint, which effectively means that possible reactions of the other parties involved are not anticipated. Game-theoretic works focus on the anticipation of possible reactions of competing INSPs and typically model the optimisation problem itself in much less detail.

Game-theoretic works are, for example Baake and Wichmann (1998); Badasyan and Chakrabarti (2003); Dewan *et al.* (1999, 2000); Giovannetti (2002); Norton (2004).

⁶ Since 2003 WorldCom is known as MCI (www.mci.com).

⁷ There is a potential free-riding on infrastructure investments as the quality of a connection between two end-users depends crucially on the most congested network in the path. When two networks peer with each other and one of them is congested, the quality of the connection does not improve when the non-congested network upgrades its infrastructure. If the congested network chooses not to upgrade its infrastructure, it would have the full cost savings, and would share the reduced performance with all the networks it peers with; see Giovannetti *et al.* (2003).

The rationales behind peering decisions for commercial INSPs and for academic research networks are analysed in Baake and Wichmann (1998); the focus lies on analysing competition and business stealing effects.

Dewan *et al.* (2000) and Dewan *et al.* (1999) concentrate on the economics of direct line interconnections, assuming that IXPs are congested and there are thus incentives to move away from them. INSPs differ on the basis of connected content providers. Dewan *et al.* (2000) discuss direct-line interconnection agreements between INSPs that compete for customers in the same area, while Dewan *et al.* (1999) discuss the same approach for INSPs that do not compete for customers in the same area. However, as the congestion problem at most IXPs seems to be solved, the results are of less interest today.

Giovannetti (2002) presents a game-theoretic analysis of the effect that offering transit for other providers (including direct competitors) has on a provider, who monopolistically controls a bottleneck, and on its competitors.

The work of Badasyan and Chakrabarti (2003) in which INSPs decide on private peering is also relevant. INSPs compete by setting capacities for their networks, capacities on the private peering links, if they choose to peer privately, and access prices. The model is formulated as a multistage game and examined from two alternative modelling perspectives – a purely non-cooperative game, where the subgame perfect Nash equilibrium is solved through backward induction, and a network theoretic perspective, where pair-wise stable and efficient networks are examined. The INSPs in this model compare the benefits of private peering relative to being connected through an IXP. The result of Badasyan and Chakrabarti (2003) was that a mixed approach of connection via private peering and IXP has the most advantages.

An interesting work related to the game-theoretic works is the Peering Simulation Game described in Norton (2004), where the participants play providers and negotiate interconnections.

Decision theoretic works are, for example Awduche *et al.* (1998); Heckmann *et al.* (2001); Hwang and Weiss (2000); Liu *et al.* (1998); Weiss and Shin (2002).

In Heckmann *et al.* (2001) a part of MPRASE (Multi-Period Resource Allocation at System Edges, see Heckmann *et al.* (2002)) is presented. It is a mathematical framework that describes and solves all kinds of resource allocation problems at the edge between two networks. Heckmann *et al.* (2001) discuss (among other things) the selection of the cheapest provider or the cheapest combination of providers from the point of view of a customer of an INSPs (which could be another INSP). A dynamic problem with multiple periods is investigated; Heckmann *et al.* (2001) make a decision in the first period about the combination of providers used for the rest of the planning horizon. The models of Heckmann *et al.* (2001) contain far less complex cost functions and do not include reliability and QoS issues.

Hwang and Weiss (2000) present an interconnection problem for a future QoS-supporting Internet, where Diffserv is used as the QoS architecture. The authors investigate how the cost of quality for different QoS networks characterises the optimal resource allocation strategies of the Diffserv bandwidth broker.

Awduche *et al.* (1998) present a mixed integer programming model for finding the cost-minimal placement of a given number of interconnection points within the topology of an INSP, once the decision to interconnect is made. Liu *et al.* (1998) take a similar approach, but additionally consider the switch/router placement (network design problem).

Weiss and Shin (2002) model the decision of two AISPs on whether to use a classical peering or a transit interconnection. They model one BSP as the upstream (transit) provider and two AISPs as the downstream providers with different market shares. The traffic is a function of the market shares of the two AISPs. The potential settlement for the transit interconnection is calculated as a function of the maximum inbound or outbound traffic volume. End-users pay a certain price for their traffic. Weiss and Shin determine the break-even price depending on the market share of the AISPs.

However, their study is fundamentally flawed in two points. First, they assume that as long as the AISPs in the model can make any profit with the transit agreement, they prefer it over peering – ignoring that peering might be more cost efficient. For an economical study this is not convincing; more convincing would be to assume that the AISPs prefer the interconnection type that offers them most profit. Second, in Weiss and Shin (2002) service requests and traffic volume are assumed to flow in the same direction, which they do not in reality for the dominant applications (WWW, P2P file transfer).

9.5 Summary and Conclusions

ISPs must connect their networks with its neighbouring networks to provide end-to-end connectivity. The connections between two networks are called *interconnections* and were discussed in this chapter. We looked at the big picture, the macroscopic topology of the Internet, and then at the individual interconnections. The most common interconnection types are classical peering and transit. How the interconnections are technically realised is also important; we call this the interconnection method. There are two basic interconnection methods, using an Internet exchange point or a direct connection. Most providers have a mix of different interconnection types with different other providers. At the end of this chapter, works related to determining the interconnection mix were discussed.

In the next chapter, the influence of the interconnection mix on efficiency and QoS is studied. Several strategies are derived for optimising the interconnection mix with respect to different goals. The next chapter of this book can be classified as a decision theoretic work.