

Introduction

This thesis deals with mathematical optimization methods for the design of survivable optical networks in telecommunications. The work on this topic has been initiated and motivated by a couple of industrial projects at the Zuse Institute Berlin (ZIB) with Telekom Austria AG and T-Systems International GmbH (formerly T-Systems Nova), continued in ongoing cooperation. Characteristic for research driven by real-world applications, the tasks to solve are typically tackled by a team of experts from different fields. In our case, the team consisted of the practitioners at our partners and a group of colleagues and myself at the optimization department at ZIB, headed by Martin Grötschel. Therefore, this thesis is written from our common perspective including all people I had the opportunity to cooperate with.

The first project with Telekom Austria started at the beginning of 2000 and was supposed to create a software tool for direct support of the network planning process in practice. Much attention was paid to operate on an as accurate as possible model of the optical network on disposition and its functionality in order to ensure that generated designs are in fact realizable in the field. Though not all involved details are incorporated in the following, the study has fundamentally determined our particular perception of optical networks and their working. The developed tool, in a preliminary version presented at CeBit 2001, was finally delivered to our partner in December 2001 and served further as groundwork for our continued research on the issue.

The second project with T-Systems Nova (meanwhile T-Systems International) was accompanied by a parallel project on dynamic call admission and other online tasks in traffic engineering. It was funded by the DFN-Verein, the German research network operator, and started in autumn 2000 with the goal to economically evaluate optical network designs under various settings and architectural prerequisites. To this end, we extended and enhanced our framework, models, and algorithms to integrate variable configuration alternatives. The continuously enhanced methodology has been successfully applied in several computational studies during the project and, perpetuating the fruitful cooperation up to date, it has proven a vital base for further investigations and innovations.

In consequence, we focus on two important aspects for the design of optical networks:

- economic efficiency, i.e., cost minimization accompanied by a mathematically proven quality guarantee, and
- a realistic model of optical networks and their various devices, covering all relevant properties and rules of interaction.

Optical technology is nowadays broadly employed in modern telecommunication networks. After emerging on the commercial markets, optical transmission systems were soon recognized as the future technology superseding electronic transmission in those (sub-) networks where large amounts of data are to be shipped, predominantly in the so-called backbones. Consequently, a rapid penetration of these networks by optical fibers took place, further driven by ongoing progress in extending throughput capabilities. In parallel, engineers and physicists intensified their efforts to overcome bottlenecks resulting from physical limitations in order to enable more and more functionalities being directly applied to the optical signals, like switching or signal regeneration. As a result, such new devices either are nowadays commercially available or are expected to be soon. Currently, first generation optical networks with restricted functionality in the optical domain are already state-of-the-art in practice, and though still high prices for innovative devices slow down the process, a migration to second generation optical networks will, following common opinion, take place in the near future. Our methods, explicitly taking upgrade planning into account, provide helpful support for this purpose, since they are flexible with respect to the architecture and able to cope with the design of networks of either generation.

The great advantage of optical technology, to transport immense data volumes with maximum possible speed, bears a risk, too. An unexpected disruption of working connections can yield tremendous loss of payload (as optical signals cannot be buffered) and breakdown of numerous services. Hence, network providers are in particular interested to design their high capacity optical networks in a survivable manner, taking preventive care for failure situations in order to speed up the recovery process. Therefore, we put special emphasis on the issue of survivability in optical networks. Various protection mechanisms have already been considered, and each one comes with individual properties regarding occupation of additional backup capacity, recovery time, and complexity of network management. We present a new scheme, called Demand-wise Shared Protection, which has been particularly tailored to optical networks and offers a special balance among these criteria.

Mathematical methods for network design in telecommunications have a long standing history, as such problems show a high combinatorial complexity which is hard to master by planners without appropriate support. Besides optical networks, various network architectures and technologies have been and are currently under investigation, like networks based on the Internet Protocol (IP) or by use of the Synchronous Digital Hierarchy (SDH), a widespread data transport technology. The common core task of such design problems typically consists of finding the routing of traffic connections in combination with a corresponding dimensioning of capacities required for their accommodation. For optical technology, this task extends due to a new feature: a spectrum of different wavelengths used for parallel transmissions. As a consequence, carrying out a wavelength assignment for the connections becomes part of the network planning problem. From our perspective, the complete task can be roughly subsumed as follows:

Given is a network topology consisting of nodes and (potential) links between node pairs. For each node and link, existing and installable devices are specified. Communication forecasts define a set of demands for connections to be established between pairs of nodes together with

additional survivability requirements.

Minimizing the total cost for installations, find a feasible design of the network that meets the traffic demands. An optical network design comprises to select the required devices at each link and node such that sufficient capacities are provided for the connection establishment, and to determine a suitable routing of all needed optical connections with specified wavelength to use on each traversed fiber link such that any two connections sharing a fiber use different wavelengths.

This brief description outlines a comprehensive and complex problem, especially when focussing on a detailed network representation. A vast variety of alternative equipment configurations and routing possibilities in combination with associated wavelength assignments (or path colorings) has to be evaluated, being compatibly plugged together for construction of a feasible network design at least possible cost. We develop dedicated mathematical optimization methods providing a structured approach to determine suitable solutions. As fundament, we derive integer linear programs to model the problem formally. For increasing computational tractability, we propose a suitable decomposition approach carried out such that limited degree of accuracy is sacrificed in order to obtain good designs with approved quality guarantee. Our way of decomposing grants access to sophisticated methods for the generic core task, and special emphasis is put on the extension feature modeled as generalized coloring problem. Besides its complexity, we study primal and dual approaches, combining fast construction heuristics and lower bounds for the optimum value. For the bounds, particular strength is achieved from the devised path packing formulation solved by a column generation algorithm which is further developed to an exact branch-and-price method. The suitability of both methods and their results is evaluated in an extensive computational study.

Outline. The outline of the thesis follows a natural way when developing mathematical methods and tools for practical applications: first specifying the real-world problem, then translating it into suitable mathematical models, next finding a strategy how to approach their solving, followed by working out details and required algorithms, and finally applying the developed methodology to evaluate its performance and capabilities. This order of matters is reflected by the structuring of chapters. In Chapter 1, we give an introduction into telecommunication networks and optical technology. We explain the practical background and discuss the layout and organization of optical networks, including the issue of network survivability. Since plenty of optional settings are possible, we make also basic assumptions defining the scope of problems which are considered.

In Chapter 2, the focused optical network design problems are formalized and described by mathematical models. A general framework provides the basis for representation of networks and their composition. For three characteristic architectures, we then discuss the modeling by integer linear programs in detail, together with selected variations and extensions. In addition, we introduce the novel survivability concept Demand-wise Shared Protection and explain its integration within the program formulations.

Chapters 3 and 4 are dedicated to the solution methodology. In Chapter 3, we explain our basic approach to solve optical network design problems by decomposition

into two subsequent subproblems and discuss its advantages. For the first of these subproblems, the dimensioning and routing subtask, we then present our solution methodology including alternative ways of application.

In Chapter 4, we investigate the second subproblem: wavelength assignment with converters, forming the characteristic subtask in optical network design. We discuss the complexity of the problem and derive integer linear program formulations. In addition to a combinatorial method, the linear program relaxations provide lower bounds on the number of unavoidable conversions. As constructive approaches, we propose several fast heuristics and provide finally an exact branch-and-price method. In Chapter 5, we report on the results obtained with the developed methods when computing optical network designs for realistic instances. The purpose of this computational study is twofold, serving on the one hand as performance evaluation of our solution approach and, exploiting its flexibility, on the other hand as a demonstration of possible ways of utilization for design analyses and comparisons. The first part, devoted to assessment of entire configurations, is organized in several case studies on different settings in view of survivability schemes, upgrade planning, alternative hardware prerequisites, and architectures. In the second part, we appraise the various algorithms focusing on the wavelength assignment subproblem.

Some concluding remarks on the achieved results and directions for further research close the thesis. The appendix provides an overview on used notation as well as a detailed compilation of the original computational data for both input and results that have been discussed in more significant aggregation in Chapter 5.

Chapter 1

Optical networks in practice

The idea of using light as long distance information carrier arose as early as 1958, when the laser was conceived. A controlled emission of defined light signals builds the basis for photonic communication. Next milestones have been the development of guided wave transmission in the mid 1960s and the fabrication of first low-loss glass fibers around 1970, making the optical transmission practical for telecommunication networks. Although first considered as a curiosity technology, operators soon recognized the potential and the superior properties offered by optics as medium, and the light entered the networks.

The installation of the first fiber connections in the early 1970s marks the beginning of the evolution of optical networks. Due to physical enhancements, the new technology did not only compete with established systems, but soon outperformed them in effectivity. In fact, the bitrate-distance product offered by optical transmission roughly grew exponentially over the years. As a result, more and more copper cables were replaced by fibers. The 1980s brought further major breakthroughs which enlarged the application field for optics and supported to its penetration of the telecommunication infrastructures. The first fiber trans-Atlantic cable laid in 1988 provides a striking example. Consequently, fiber systems became the dominating long-distance connection technology from the late 1980s on.

Until the end of the 1990s, optics was used in telecommunication networks as pure transmission technology, whereas electronics accessed and handled the data flows in the network nodes. In the meantime, the rapidly increasing optical bandwidth capacities overstretch the capabilities of electronic equipment which is subject to fundamental limitations. As a consequence, more and more functionality “goes optic”, from switching optical channels directly up to refreshing the signals, exchanging wavelengths, and more.

Nowadays, modern optical telecommunication networks form complex infrastructures composed of a plenty of various devices. Operating such a network is a challenging task. In particular, it requires an accurate planning, taking into account all technological “needs” and “offers” of the optics, in order to end up with a working network which provides the best return on the spent investments.

In this opening chapter, we give a brief introduction to optical telecommunication networks. The demanding task to accurately plan such complex infrastructures re-

quires a basic understanding of the applied technology and their mechanisms as well as general networking aspects. For the scope of this thesis, we confine the description on those facts that are relevant for the design of optical networks, taking the perspective of the network planners, not of the engineers. The presentation aggregates various aspects in our appropriate perception and is guided by experiences we learned from the practitioners at our industrial cooperation partners. Although many further issues are indispensable for the practical operation of optical networks, we leave out detailed physical and technical explanations. The books of Mukherjee [123], Ramaswami and Sivarajan [147], and Stern and Bala [154] provide a comprehensive introduction to optical networking with an engineering-oriented emphasis.

We begin with the main ingredients of optical networks, the involved hardware devices and their operational properties. Due to the rapid progress in research and product development, the attempt to give a full account of the state-of-the-art technology is futile. Moreover, network operators do not upgrade their hardware with any innovation arising, but make only major reformations from time to time. The network planning thus has to deal with combinations of new and old (if not ancient) technologies, while foreseen advancements of the next years should also be prepared within the designs to provide the future. Therefore, the description is not restricted to hardware that currently is available, but also includes those components that will be available soon, or are under development due to a demanding market. Occasionally remarked technical quantities serve just to illustrate the concerned magnitudes without claiming to represent current technological capabilities.

After having discussed the possible devices and their functionalities, we consider fundamentals of operating optical networks. The physical network and its layout build the groundwork for managing transmissions. To transport the data traffic, connections between network nodes have to be set up. We describe optical connections and discuss their operational issues. Moreover, we introduce the concept of a lightpath as a characteristic connection type for optical networks.

Since the network operation shall finally be robust against failures, survivability is a further important issue. We discuss some basic aspects and present known concepts to realize survivable networks. The most relevant schemes are categorized and compared with respect to several practical criteria.

The migration of real-world networks has brought up several constitutive network architectures, subsuming combinations of the used technology and applied concepts. To illustrate the evolution, we briefly sketch the progression of optical networks in practice and present main types of architectures that consolidated over time. For design tasks, we focus on these basic network types and their relevant properties.

The chapter is completed with a description of the planning task(s) for optical networks. A frictionless operation relies on a proper network configuration. We specify this comprehensive term and identify configuration subproblems which compile highly related issues. Moreover, planners in practice are confronted with various questions to answer. We discuss some major aspects and present typical planning types and variants. These descriptions serve as orientation for the final specification of those practice-relevant optical network design problems that are studied in this thesis and for which we develop appropriate solution methods in the subsequent chapters.




Figure 1.1: *Guided wave propagation through an optical fiber.*

1.1 Optical technology

For a long time, telecommunications was a pure electronic domain. Initially, networks of copper cables grew rapidly and brought phone services into everyone's life. As next step, data entered the networks. The extension of computer utilization and the digitalization—an ongoing trend—turned communication networks into a platform for improved services and applications, for instance in e-commerce. The continuous increase of data traffic demanded for more and more bandwidth to be provided by the networks. For this, the electronic equipment was at some point found to be constricted by physical limits. To achieve further advancement, there was need for a new medium.

1.1.1 Optical transmission

Using light as data transmission medium is based on the principle of guided wave propagation. An emitted wave propagates through a carrier material whose shape guides the wave through space to reach the intended target. For optics, the carrier material usually consists of the highly transparent glass (silica) core of an *optical fiber*. Light waves entering the fiber propagate along the glass core, being reflected whenever they hit upon the side borders of the glass.¹ Forwarded this way, the waves can only abandon the glass core at the end of the fiber (see Figure 1.1). This enables the directed transmission of light between any two points connected by optical fibers.

Optical channel. To transport data optically, the digital information delivered by electronic devices passes a *transmitter* which converts the bits into short light pulses by a laser. The light then propagates along fibers until it reaches a *receiver* where the signals are reconverted by photodiodes or photodetectors and handed back to the electronic equipment. Transmitters and receivers build the interfaces between electronics and optics. We call the optical part of such a connection on a fiber an *optical channel*, as illustrated in Figure 1.2. The shorter the emitted light pulses are, the more bits can be transmitted within some time period. The *bitrate* of a transmission system is defined as number of bits per second (b/s) that can be transmitted/received. The achievable bitrate is given by the physical skills of transmitter, receiver, and fiber. In today's practice, systems with a bitrate of

¹ Note that 'reflection' is just a simplifying description. There are different physical principles applied to direct the light waves along a fiber core, depending on the properties of the carrier material (step- vs. graded-indexed, single- vs. multi-mode etc.) as well as the specific type of light waves (mode, polarization etc.) at hand.




Figure 1.2: *Optical data transmission on fibers.*

10 Gigabit per second (Gb/s) are widely used. Systems with 40 Gb/s are also commercially available, although rarely applied (see Freeman [51]). News from the laboratories indicate that manufacturers already head for 160 Gb/s systems and beyond.

Signal degradation. Under ideal conditions, the received optical signals would have exactly the same shape as emitted. Unfortunately, real systems do not provide ideal conditions. Depending on the material of the glass core and its quality, there are also undesirable effects (accidental such as material impairments, but also systematical such as dispersion, attenuation, cross-talk, and others). These effects disturb the light propagation on optical fibers. Hence, the signal quality suffers a loss, limiting the distance over which an optical signal can be carried without losing or falsifying the information. More precisely, the emitted light wave carries a rectangularly shaped amplitude coding the 0-1 bit information stream as light off/light on. On its way, it blurs to a wave with fuzzy shaped amplitude which only resembles at the original form. The receiving photodetector reinterprets the light information as bits by applying a threshold to the light amplitudes, as depicted in Figure 1.3. The longer the distance of the optical transmission is, the more the original amplitude shape is tampered. When exceeding a certain length, the light wave does not allow anymore for a correct reconversion of the original information, and bit failures occur. In practice, the *bit error rate* is used as a measure for the signal quality, defined as the ratio of incorrectly interpreted bits over all transmitted bits. The lower the bit error rate is, the less coding effort and redundant information is required for error detection and correction. Currently, fibers available at the market typically set 70 to 120 km as transmission bound without additional equipment and a bit error rate ranging between 10^{-7} and 10^{-9} .

Signal regeneration. To extend the maximum optical transmission distance on fibers, so-called *amplifiers* are used to refresh the signals optically. As the name indicates, an amplifier simply scales up the amplitude of the light wave to allow for a correct signal reversion after an increased total transmission length. The




Figure 1.3: *Transformation of emitted light pulses along fibers and their restoration by a photodetector.*




Figure 1.4: *Optical signal regeneration stages.*

maximum distance can be further elongated by placing amplifiers consecutively in regular distances. Such a *reamplification* cannot be repeated arbitrarily often, because the light wave blurring does not only influence the amplitude heights but also spreads the widths of the original signals (as wave crests and troughs). At some point, wave equalization effects dominate and make a correct signal interpretation impossible. Amplifiers are usually placed each 60 to 80 km at a fiber, enabling to transmit optical signals over up to around 1000 km, with special ultra long-haul equipment even several thousand km.

Reamplification of light waves is just the simplest kind of signal refreshing. After having passed an amplifier, the light wave shows heightened amplitudes, but still has a fuzzy waveform. A more extensive signal regeneration consists in restoring the original rectangular shape of the wave by use of a threshold similar to a receiver processing. This *reshaping* process reproduces a rectangular waveform, but the rectangle sizes may differ from the original ones due to the distorting effects. A full restoration of the original signal thus requires an additional step to unify the bit lengths: The emitting has to undergo a *retiming*. This step finally results in a rectangular waveform with steady sized rectangle lengths. It equals the original light emission as long as the received signals have been interpreted correctly, i.e., no bit errors occurred. Figure 1.4 depicts the three stages of optical signal regeneration: reamplification, reshaping, and retiming. Since each stage also involves the former stages in this order, they are also denoted as 1R, 2R, and 3R regeneration (cf. the ITU standardization recommendation in G.872 [160]).

Regenerator. The differentiation between the regeneration stages results from their distinct difficulties for optical realization. As indicated above, 1R regeneration can be optically performed by rather simple amplifiers. Using a pump laser, power is injected to increase the light wave's energy and thereby its amplitudes. Reshaping and retiming, however, are not that easy to realize. In practice, a full signal regeneration currently requires to transform the signal back to electronic form by a receiver and then pass it back to a laser reemitting light pulses. Due to the signal transformation from *optics* to *electronics* and back to *optics* again, this process is abbreviated as *o-e-o conversion*, and the executing hardware device called

transponder is built of a receiver and a transmitter back to back. Recently, progress in optical 3R regeneration has been reported, see Nolting [128]. The integration of *optical 3R regenerators* would overcome the need to interrupt optical transmissions by o-e-o conversions and lift the ban of a maximum distance. To simplify the further description, we call any device that performs full signal regeneration a *regenerator* independent of the applied technique.

Further extensions of the maximum optical transmission length have been attempted by experiments with various fiber materials and exploiting several light wave properties, such as polarization, modulation, and others. While some compositions have transcended certain disturbing effects, other distorting influences become more dominant for them. Up to date, no optical transmission system without limitation on the maximum distance to bridge is available. Research, however, yielded a variety of different kinds of devices, including fibers, transmitters, receivers, and amplifiers, which are now available at the market. Each device type comes along with specific properties that make it the better choice for specific situations.

Compared to electronic signals on copper cables, optical transmission on fibers provides many obvious advantages: higher bitrates, lower loss of signal quality, immunity against electromagnetic influences and noise, lighter and more supple carrier material, corrosion resistance, and advanced security (since optical signals are more difficult to tap). From an operator's point of view, however, the most striking argument is given by the incurred cost per transmission unit. At first sight, optical systems have a much higher price than electronic devices, but they also offer a much higher transport capacity. Electronic transmission speed has reached its limits imposed by physics, with a maximum bitrate around 10 Gb/s. So, further capacity expansion is only possible by massive parallel use of devices. This correlates with a cost explosion when demanding a moderate capacity increase, say by a factor of two or four. Optical transmission already enables higher bitrates than electronics, and further advancements in ultra-short light pulse emission give rise to expect further bitrate boosts. Moreover, optics provides further unique opportunities for capacity expansions with low effort: multiplexing.

1.1.2 Wavelength Division Multiplexing (WDM)

In general, the term *multiplexing* refers to techniques which allow for a shared use of a scarce resource. The basic idea of this principle is to divide a common resource appropriately into 'smaller parts' which can be used in parallel to fulfill a certain task. In this way, the resource utilization is multiplied without enlarging the resource itself. In most cases, the possibility to multiplex a resource is based on advanced technology which allows to use the smaller parts.

Multiplexing techniques. In telecommunication networks, multiplexing has been applied successfully in several manners. The most simple case consists in the parallel use of a conduit by multiple cables, which is known as Space Division Multiplexing. It is also applied in optical networks by bundling many fibers into one cable. A more prominent example of the principle is given by Time Division Multiplexing (TDM). In order to benefit from an increased transmission speed without