Chapter 3

Communication traffic engineering

3.1 Statistical nature of communication traffic

In its most elementary form, the basic purpose of any communication network is to satisfy the communication requirements of its customers through the delivery of products and services over limited physical infrastructure. From the physical infrastructure's perspective, these service and product demands are dealt with as communication traffic exhibiting non-uniform statistical behaviour. The fluctuating flows of traffic through a communication network not only influences network management, load balancing, capacity utilisation, and quality of service, but also the way in which the network should have been designed. This results in the familiar chicken-egg dilemma that is best resolved through the use of traffic models and simulation during the network design phase.

Communication networks typically exhibit periodically varying traffic levels that follow daily, weekly and monthly cycles. As the geographical area of large wide-area networks increase and extend over several time zones, the appearance of this kind of periodical behaviour gradually change from a network-wide phenomenon to more localised occurrences. The variation of network traffic as a function of time does however still account for great uncertainty in the utilisation levels of the various links that comprise a network.

3.1.1 Geographical distribution of communication traffic

A strong correlation exists between the add/drop traffic of a network node and its proximity to other network nodes. This is due to the tendency of people to populate metropolitan areas and conduct their economic activities in these metropolitan areas. It is therefore expected for metropolitan areas to contain more network nodes than rural areas, and for these network nodes to have more add/drop traffic than those in rural areas. The add/drop traffic of a network node as well as its proximity to other network nodes should thus be taken into account when designing wide-area networks that provide connectivity between all network nodes.

The concept of communities of interest evolved from the geographical correlation frequently observed between the source and destination in general communication network traffic. To some extent it can even be argued that network designers and operators encourage the seemingly natural occurrence of communities of interest, because of the cost-saving benefits of networking functions such as the proxy and mirror commonly utilised in packet-switched data retrieval applications such as the World Wide Web (WWW).

The clustering of network nodes in such a way as to exploit communities of interest and provide the connectivity associated with wide-area networking, is discussed in section 5.3. The creation of clusters residing on the various levels of the multi-level network model, described in section 4.1, is an intuitive result of the communities of interest concept. It is however important to note that the geographical distribution of communication traffic relates to the network's source-destination pairs' contribution to the total network traffic, as opposed to the geographical proximity of network nodes alone.

3.1.2 Traffic models

The traditional application of traffic models is to assist network operators to estimate the traffic load distribution to be expected if changes in the network should occur. Changes such as network faults or malfunctions can be simulated, allowing for analysis of the resultant impact on the load on protection paths. Another change that can occur in the network might be on the service layer, where traffic models can be used to predict how new services supported over the network would affect its ability to provide required levels of reliability and capacity for delivering existing services.

Theoretical traffic models that attempt to describe actual communication traffic in a network usually have difficulty in obtaining good fit with regards to the marginal probability distribution and autocorrelation function of the empirical time series [28]. Several time series models exist for describing Internet traffic, but are usually limited to specific network topologies and implementations. The late NSFNet was well studied in this regard [29, 30] due to the availability of traffic statistics [31] that assisted in the verification, benchmarking and further development of network traffic models.

Burstiness, described as significant positive autocorrelation in the inter-arrival process, results in increased waiting times without considerably influencing the nett arrival rate.

These are parameters that typically do not behave in such a fashion when generated in theoretical traffic models, hence the motivation for traffic models that can more accurately represent the activity on the channels of modern WDM optical communication networks.

Traffic models are also used when add/drop traffic statistics are unavailable for a wide-area network that is to be designed, or when predictions are required on the influence of planned modifications to the network. Not only the amount of a network node's add/drop traffic is important, but also its geographic distribution. Gravity models using statistics such as regional population and economics can be used to estimate the traffic relationship between all network nodes, from which the estimated add/drop traffic of all network nodes can be calculated.

Gravity models

The simple gravity model, as depicted in figure 3.1, is a popular mathematical tool, inherited from physics where it is used in various of its branches, including the fields of statics, dynamics, and even astronomy. The underlying principle of the gravity model is the weighting of inter-point relationships based on their relative importance in a system of points. In the communication networking context, the points constitute network nodes and the inter-point relationships give an indication of the logical topology as described in section 4.2.2. The relative importance of a network node in a network comprising several network nodes is the key parameter in demand estimation, because it directly determines the logical topology of a network.

With reference to figure 3.1, the meanings of the employed symbols are as follows. W_i is the relative weight of network node i, while W_j is the relative weight of network node j.

In the denominator of the inter-nodal weight expression, a summation of relative nodal weights is made over all network nodes by using k as an index to N total network nodes. The inter-nodal weight expression $W_i \times W_j/(\sum_{k=1}^{N} W_k)^2$ shown in figure 3.1, is evaluated for all combinations of network nodes, indexed by i and j, as per the logical connectivity requirements to be presented in the network's logical topology.

Elaborations on the simple gravity model exist and are referred to as modified gravity models. These models exhibit various weighting preferences and can be customised to suit the requirements of a given problem. A popular modification is the inclusion of a distance parameter to allow for the phenomenon of communities of interest, that form based on geographical proximity. Caution should however be taken in the modification of gravity models to avoid the introduction of bias that can negatively impact on the dynamic range and resolution of the achievable inter-nodal network weights. Relationships ranging from linear to exponential and even polynomial can be achieved for the mapping from node weighting parameters to inter-nodal network weights. The careful selection of the relationship can ensure more robust capacity provisioning, capable of sustaining fluctuations in network load and unbalanced demand growth.

The population-distance model used in the European Optical Network (EON) project [16], is very similar to the gravity model, modified to include the distance metric. The lack of reliable demand estimates from regulatory bodies required the EON network designers to do a demand estimation for the required capacity between points s and p in the network by evaluating the equation

$$D_{s,t} = K\frac{P_s P_t}{distance_{s,t}}, \tag{3.1}$$

where K was a constant of 5.25 Erlang and the units of population P being in millions

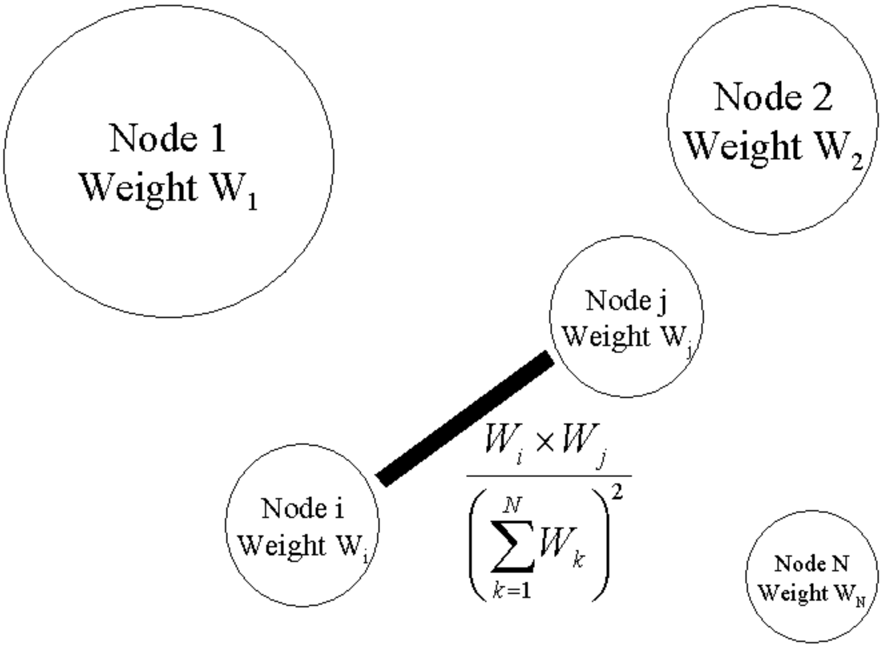

Figure 3.1: The simple gravity model.

and *distance* in kilometres. From the constant K being in Erlang it is easy to deduce the circuit-switched thinking that was still prevalent during the time of the EON project.

3.1.3 From network node weighting to demand matrices

The determination of relative network node weights is crucial for the utilisation of techniques such as the gravity model, which is discussed in section 3.1.2. Various parameters obtained from statistical analysis of a planned communication network's geographical area can be employed in the development of a relative network node weight. Popular parameters include absolute, growth potential, growth trends, and demographic breakdowns of metrics such as population, economic activity, teledensity, and available technology.

Network node weighting functions can be as simple as relative population density or as complex as a non-linear polynomial interaction of numerous parameters. Experience in signal processing and pattern recognition disciplines suggest that less complicated characteristic functions produce results that are comparable to that of complex functions, with the benefit of providing insight into parameter interaction. For this reason, lower-order network node weighting functions with fewer parameters are recommended, since complex functions do not enhance the power of an equation that ultimately contributes to a process based on non-discrete decision mechanisms and evaluation criteria.

Population alone is not a sufficient parameter to use when weighting a network node. Economical activity has been identified as being key to the communication requirements of network nodes. For this reason a modified gravity model has been employed to create the flow distribution and subsequent demand matrices that describe the capacity requirements between all network nodes. The gravity model has been modified to include consideration of geographical positions, thus allowing for community of interest factors to influence the creation of the demand matrix. The demand matrix is populated through the evaluation of the following modified gravity model, where a simple nodal weighting of economic activity has been employed:

$$D_{i,j} = K_1 \frac{E_i \times E_j}{dist_{i,j}}, \tag{3.2}$$

where E_i and E_j indicate the economic activity of nodes i and j, $dist_{i,j}$ is the distance between nodes i and j, and K_1 is a normalising constant chosen in such a way as to ensure that the network capacity requirement is satisfied as follows:

$$C = \sum_{i=1}^{N} \sum_{j=1}^{N} D_{i,j}, \tag{3.3}$$

with network capacity C being the chosen aggregate network capacity requirement of the network under design. With reference to $dist_{i,j}$ of equation 3.2, it is important to note that the distance between two points on a sphere is given by:

$$\begin{aligned} dist_{i,j} = 1.852 \times 180\frac{60}{\pi} \arccos\Bigg(& \sin\frac{lat_i}{180\pi} \sin\frac{lat_j}{180\pi} + \\ & \cos\frac{lat_i}{180\pi} \cos\frac{lat_j}{180\pi} \cos\left(\frac{long_i}{180\pi} - \frac{long_j}{180\pi}\right)\Bigg), \end{aligned} \tag{3.4}$$

where $dist_{i,j}$ is the distance along the surface of the sphere from point i to point j in kilometres, lat_i is the latitude of point i in degrees and $long_j$ is the longitude of point j in degrees.

3.2 Matrices representing network traffic and flow distribution

A traffic matrix mathematically presents the volume of traffic that a communication network carries between its nodes, through the mapping of traffic sources and destinations on a two dimensional matrix. These nodes could be ingress and egress points on a backbone transport network, gateway routers at the edge of an Internet service provider's (ISP's) point-of-presence (PoP), or interfaces to the subnets of an enterprise network. The backbone network is considered core to this investigation into factors influencing wide-area network design, which is why the discussion of matrices that describe network traffic is limited to this context.

The terms *table* and *matrix* are used interchangeably in literature and industry to refer to the method of presenting the flow, interrelation and distribution of communication network traffic, whether in estimated, provisioned, demanded, capacity or actual embodiments. Aspects such as symmetry of network traffic and the differences between intra and inter-nodal traffic are discussed in this section. Flow distribution matrices are also discussed at length, motivated by their importance to logical topologies, as discussed in section 4.2.2.

The structure of the matrices under discussion consists of row and column labels, indicating the network nodes that constitute the sources and destinations of the connections represented at the intersecting cells of the matrix structure. Nodal totals are computed and indicated at the end of each row and column, with a grand total indicated at the intersection of the end row and column. No set standard exists for whether sources or destination should be mapped to rows or columns, but it is important for a network designer to specify the chosen convention when an actual matrix is constructed. The convention of sources along rows and destinations along columns will be followed throughout this document.

3.2.1 Symmetry in network traffic

In communication networks, the relationship between source and destination is often defined by the difference in the amounts of traffic that flows in both directions. A source is typically defined as the node in a logical connection from where high volumes of traffic originate, whereas a destination is defined as the node to where high volumes of traffic travel. Except for special broadcast and multi-cast scenarios, it is customary for bi-directional communication to exist between all source-destination pairs in a communication network. The temporary nature of most communication connections is one of the main reasons for the overhead and handshaking required in connection establishment. The connection management component, as discussed in section 4.3.2, is also responsible for the bi-directional nature of most communication links.

Since bi-directionality has been established to be a characteristic of typical network connections, the matter of symmetry between the volume of traffic generated in the two directions emerge. In a conventional client-server model such as what has become popular in the WWW, the observation of highly asymmetrical traffic can be made. On the other hand, typical circuit-switched applications such as a public switched telephone network (PSTN) are perfect examples of truly symmetrical communication. Since the first commercial implementation of optical fiber technology and subsequent optical networks in the telecommunication industry, optical network traffic is typically thought of as being symmetrical. *The prevalence of this is such that optical network designers seldom consider asymmetrical traffic models when designing wide-area networks.*

An asymmetrical traffic model offers the advantage of having symmetrical traffic as a special case, thus allowing for seamless coexistence with existing symmetrical traffic paradigm. It is therefore suggested that a network design approach that considers the possibility of asymmetrical traffic possesses great advantages above one that blindly assumes the provisioning of conventional voice services.

With the many developments and nearing maturity of voice-over-IP (VoIP) technology, the field of conventional circuit-switched voice telecommunication will be forced to adapt to a more generic packet-switched way of thinking, including acknowledgment of the possible asymmetrical properties of network traffic. The aforementioned broadcast and multi-cast scenarios applicable to video-on-demand (VoD) and other content streaming services also lend themselves to an asymmetrical traffic model.

3.2.2 Intra-and inter-nodal traffic

It is conventional to only consider inter-nodal traffic and not the intra-nodal traffic that never travels on the network level under consideration. For this reason it is typical to find an empty diagonal in the flow distribution matrix, and subsequent demand, capacity and traffic matrices, resulting from not considering traffic that originates from a specific network node and terminates at the same network node. Traffic of this nature is typically handled on a lower level of the multi-level network model, as described in section 4.1, resulting in additional matrices being created in a hierarchical fashion.

Consideration of intra-nodal traffic is justified in special situations such as described in the clustering approach discussed in section 5.3. This is a special application where recursive traversing of the multi-level network model is used to determine the way in which network nodes should best be grouped for load balancing purposes. In such an application it would thus be justified not to have an empty diagonal in intermediate flow distribution matrices, but still in the subsequent demand, capacity and traffic matrices.

3.2.3 Flow distribution matrices

One of the most important applications of the traffic models and node weighting techniques discussed in section 3.1.2 and section 3.1.3 is the development of flow distribution matrices. The design of logical topologies, as described in section 4.2.2, depend on flow distribution matrices that show the relative contributions of all logical connection permutations to the total network traffic capacity that should be provisioned for.

In a simple network consisting of a low number of network nodes it might be possible to represent all traffic flows in a single flow distribution matrix. Medium-sized to large networks usually require the creation of several hierarchical flow distribution matrices in accordance with the multi-level network model as described in section 4.1. However, where many flow distribution matrices are necessary, the prerequisite is given that the sum of all relative nodal flows over all the flow distribution matrices equate to 100%, thus ensuring that all anticipated traffic flows are provided for in a relative yet context-aware manner. Figure 3.2 shows a hierarchical collection of flow distribution matrices where symmetrical traffic demand was assumed.

In order to achieve an aggregate traffic flow of 100%, it is necessary to employ techniques for scaling the individual relative traffic flows that do not introduce unplanned non-linearity into the results of an already potentially non-linear node weighting process. An innovative methodology for creating flow distributions from weighted network nodes through the use of statistical clustering is presented in section 5.3.

The rows and columns of flow distribution matrices represent the sources and destinations of the traffic flows. In symmetrical traffic matrices, as described in section 3.2.1, the column and row totals of the flow distribution matrix separately add up to the same 100% total network capacity. In the case of asymmetrical network traffic a slightly different approach is used where individual cell values add up to the 100% total network capacity due to the inequality of row and column totals.

Table 3.1 shows an elementary flow distribution matrix with asymmetrical traffic as opposed to symmetrical traffic as shown in table 3.2. Each individual cell entry indi-

		Destination			
	Nodes	**1**	**2**	**3**	
Source	**1**		19%	8%	**27%**
	2	12%		21%	**33%**
	3	22%	18%		**40%**
		34%	**37%**	**29%**	**100%**

Table 3.1: An elementary flow distribution matrix with asymmetrical traffic.

		Destination			
	Nodes	**1**	**2**	**3**	
Source	**1**		10%	15%	**25%**
	2	10%		25%	**35%**
	3	15%	25%		**40%**
		25%	**35%**	**40%**	**100%**

Table 3.2: An elementary flow distribution matrix with symmetrical traffic.

cates the amount of traffic flow from a specific source node to a specific destination node relative to the total network traffic, where the source nodes can be chosen to be the rows and the destination nodes the columns of the flow distribution matrix. The mapping of source and destination to row and column should always be specified in a flow distribution matrix since a wrong assumption can result in the development of an incorrect logical topology.

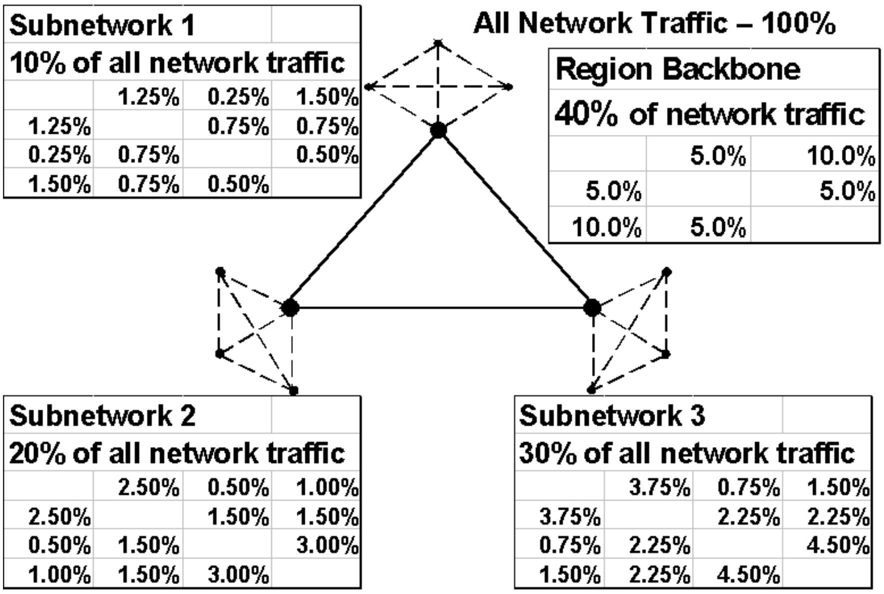

Subnetwork 1
10% of all network traffic

	1.25%	0.25%	1.50%
1.25%		0.75%	0.75%
0.25%	0.75%		0.50%
1.50%	0.75%	0.50%	

Region Backbone
40% of network traffic

	5.0%	10.0%
5.0%		5.0%
10.0%	5.0%	

Subnetwork 2
20% of all network traffic

	2.50%	0.50%	1.00%
2.50%		1.50%	1.50%
0.50%	1.50%		3.00%
1.00%	1.50%	3.00%	

Subnetwork 3
30% of all network traffic

	3.75%	0.75%	1.50%
3.75%		2.25%	2.25%
0.75%	2.25%		4.50%
1.50%	2.25%	4.50%	

Figure 3.2: Hierarchical collection of flow distribution matrices assuming symmetrical traffic demand.

3.3 Traffic grooming

When evaluating the cost of contemporary WDM optical networks, it is found that the SONET/SDH multiplexing equipment found in the digital cross-connect systems (DCSs) of the network nodes contribute substantially to the total equipment cost of the network. Figure 3.3 shows a high-level diagram of three network nodes with emphasis on the DCS that interfaces lower data rate traffic streams to the unlabeled optical add-drop multiplexer (OADM) in the middle network node. A number of transceivers are located at the interface between the DCS and OADM and constitutes a dominant cost to be minimised through efficient traffic grooming.

The compilation of a data stream for transmission on a wavelength has become as challenging as the wavelength-division multiplexing function itself. The data rates of the independent traffic streams are substantially lower than that of the optical data rates achieved on a wavelength channel. The term *traffic grooming* refers to the techniques used to combine lower data rate traffic streams onto available wavelength channels to achieve design goals such as cost minimisation and restorability.

Grooming can be seen as the time-domain equivalent of wavelength division multiplexing, with the only exception being that individual traffic streams of various, often different, data rates are combined as opposed to traffic streams of the same data rate. The problem of assigning shared wavelength channels to several individual traffic streams is complex due to the different source and destination combinations of the various individual traffic streams. Figure 3.4 shows how the number of SONET/SDH add-drop multiplexers (SADM) required in the DCS, shown left, can be greatly reduced by implementing wavelength add-drop multiplexer (WADM) functionality, shown right, in the OADM of a network node.

The combination of individual traffic streams into shared wavelength channels requires consideration of several factors including channel capacity, time-domain multiplexing and demultiplexing resolution, near-minimum hop routing, reliability through protection and restoration, and billing complexity. Since traffic grooming operates on the SONET/SDH level, it is not surprising that most research [32, 33, 34] on the topic have focused on traffic grooming in ring topologies, the most popular implementation of SONET, as shown in figure 3.5. Besides the conventional techniques for doing traffic grooming in SONET rings, some novel approaches employing genetic algorithms [35] and simulated annealing [36] have also been proposed.

The creation of concepts such as local and express traffic routes, results from the grooming of communication traffic. The combination of various individual traffic streams onto shared wavelength channels is done in such a way as to minimise the standard deviation of add-drop multiplexing required per wavelength channel. Some channels will be used for short routes, and then re-used in other parts of the network, whereas other chan-

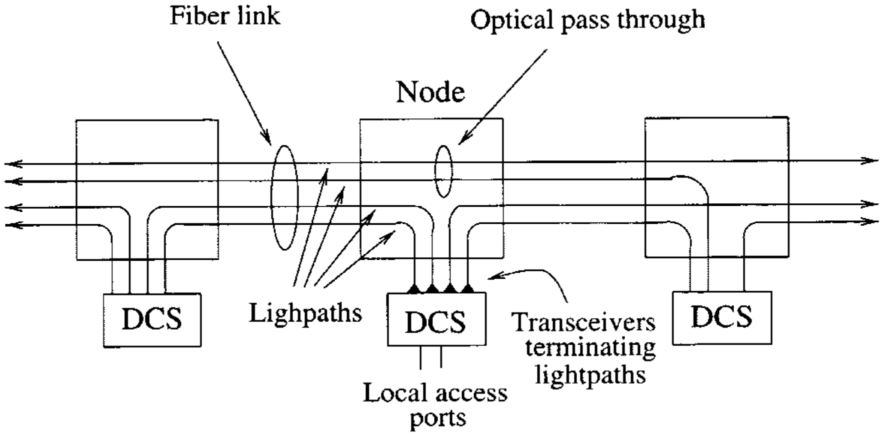

Figure 3.3: Three optical network nodes with emphasis on the interface between DCS and OADM of the middle node [33].

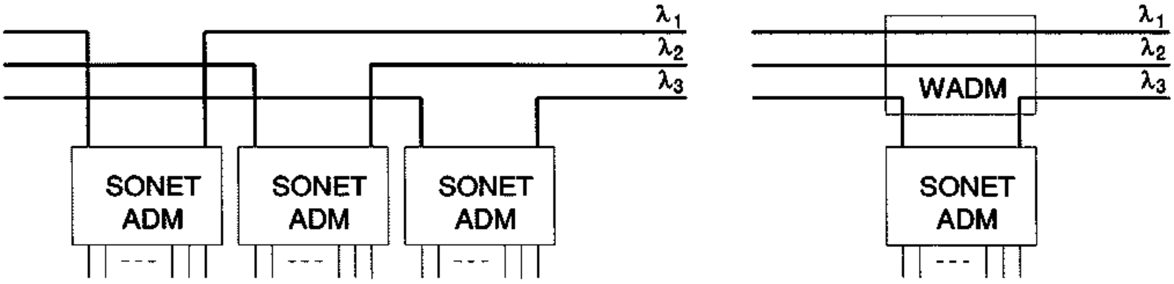

Figure 3.4: Reducing the number of SADMs, shown left, required at a network node through the addition of WADM functionality, shown right [34].

nels will function like express lanes on a highway, where vast geographical distances are covered without allowing for individual traffic streams to exit or join the shared wavelength channel.

3.3.1 The non-trivial nature of the grooming problem

Figure 3.6 shows a simple six-node point-to-point physically connected SONET-over-WDM ring utilising three wavelengths on all the optical links. Note that this network does not contain any true optical nodes, since no traffic can traverse a node without

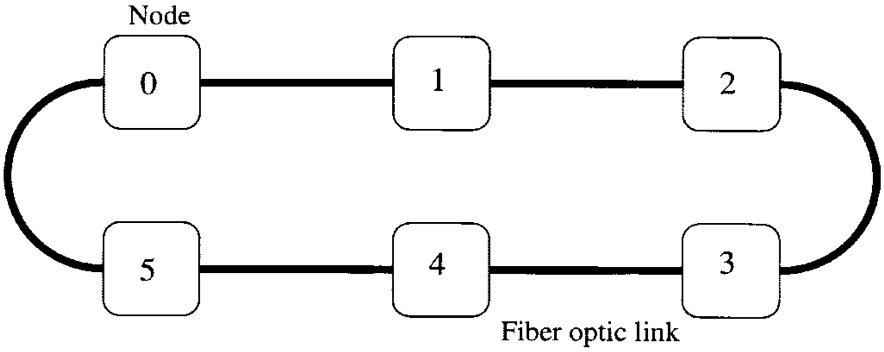

Figure 3.5: Simplified physical topology of a six-node optical SONET-over-WDM ring network [33].

being converted to and from the electronic domain where all processing decisions are made.

The non-trivial nature of the grooming problem can be explained with reference to the network in figure 3.6. Assume the following scenario: each of the three wavelengths in the network carries a SONET OC-48 data stream and it is necessary to extract an OC-3 stream from each of the three wavelengths at a specific network node. This would require all three wavelengths to be received and processed at the network node to obtain three OC-3 data streams that could have easily fitted into the same OC-48 data stream, allowing the other wavelengths to pass through the network node unhindered. In a relatively simple example like this it might seem easy to ensure more appropriate grooming, but complex networks with high numbers of network nodes and varying traffic conditions make this a highly non-trivial problem.

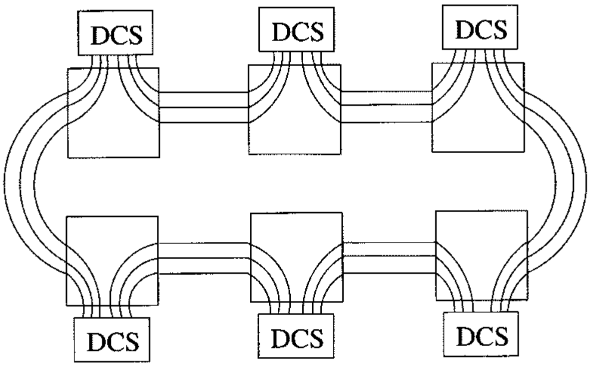

Figure 3.6: Elementary six-node SONET-over-WDM ring with three wavelengths per point-to-point physical link [33].