Part VI: Advanced Topics (Bonus Material on CD-ROM)

This part includes additional material that are related to Part IV and Part V; it consists of two sub-parts.

In the first sub-part, three chapters (Chapter 21, Chapter 22, and Chapter 23) cover functions and components of a router in further detail as a continuation of Part IV. First, different approaches to architect the switch fabric of a router are presented in Chapter 21. Second, packet queueing and scheduling approaches are discussed along with their strengths and limitations in Chapter 22. Third, traffic conditioning, an important function of a router, especially to meet service level agreements, is presented in Chapter 23.

In the second sub-part, we include two chapters (Chapter 24 and Chapter 25). Transport network routing is presented first in its general framework, followed by a formal treatment of the transport network route engineering problem over multiple time periods, in Chapter 24. The final chapter (Chapter 25) covers two different dimensions: optical network routing and multi-layer network routing. In optical network routing, we discuss both SONET and WDM in a transport network framework; more importantly, we also point out the circumstances under which a WDM on-demand network differs from a basic transport network paradigm. Furthermore, we discuss routing in multiple layers from the service network to multiple views of the transport networks; this is done by appropriately considering the unit of information on which routing decision is made and the time granularity of making such a decision. We conclude by presenting overlay network routing and its relation to multilayer routing.

24

Transport Network Routing

If you ask any filmmaker how they got into it, everyone came a different route.

Alan Parker

Reading Guideline

The discussion of basic drivers and the need for transport network routing can be read without much dependence on other chapters. However, to understand the optimization models for network engineering, a good understanding of network flow modeling as presented earlier in Chapter 4 is necessary. Transport network routing models are useful in solving relevant traffic engineering problems that arise in MPLS, GMPLS, and optical networking.

D. Medhi and K. Ramasamy, *Network Routing: Algorithms, Protocols, and Architectures.*

Transport networks play important roles in networking as they serve as the bearers of services at a physical level. Note that the term *transport* as used in transport networks has no relation to transport as used in referring to the transport layer protocol in the TCP/IP protocol stack. Traditionally, transport networks were referred as *telecommunication facilities networks* [596], and in the past decade, the term *transport network* has caught on. Along with transport network, the term *transport service* is often used; this means that a customer needs a transport network service from a transport network provider. This chapter covers transport network routing, especially for multiple time periods, to accommodate transport service requests.

24.1 Why Transport Network/Service

Transport, as in transport networks/services, requires some explanation. What does it mean and what is it transporting? Consider the following three examples:

- *IP network links/trunks:* At the view of IP networks, an IP link or trunk is between two routers of certain capacity, for example, a T3 data rate or OC-3 data rate. This view is actually only a logical view. Such links are needed to be carried over actual physical facilities. Thus, the actual physical route taken by this link is different from the logical view.
- *Trunkgroups in the telephone network:* Much like IP network links, trunkgroups that connect two TDM switches in the telephone network are also logical entities. They also need to be physically routed on actual physical facilities. Similarly, SS7 circuits between SSPs, SCPs, and STPs must be physically connected.
- *Private network services:* An entity/customer with offices located in two different cities might want to have the sites connected at a certain bandwidth, say T1, for internal voice or data traffic. In this case, such connectivity can be provided by a transport network provider that then has the responsibility of ensuring that a physical path and the bandwidth are both available to the customer.

Clearly, in each case, a transport service is needed. A key point to note is that such *links* are usually established on a semi-permanent basis, whereas bandwidth remains fixed for a long time. In many cases, such physical paths that use a link may not change for a number of years; for example, in many instances, this length is 3 years because of signed contract agreements.

Example 24.1 *Illustration of logical versus transport.*

We will consider a simple illustration to show the difference between logical and transport views (Figure 24.1). Consider switches in three locations: San Francisco (SF), Kansas City (KC), and New York (NY). We use *switch* as a generic term here for illustration; it can be routers or customer edge equipments for private transport services. They are connected by logical trunks. To provide these logical trunks, ingress and egress transport nodes must be located in close proximity of end switch sites, which are connected by a transport-routed path that can further go through one or more intermediate transport nodes. In general, transport network nodes are referred as *cross-connects*. Note that the cross-connects are not visible to switches.

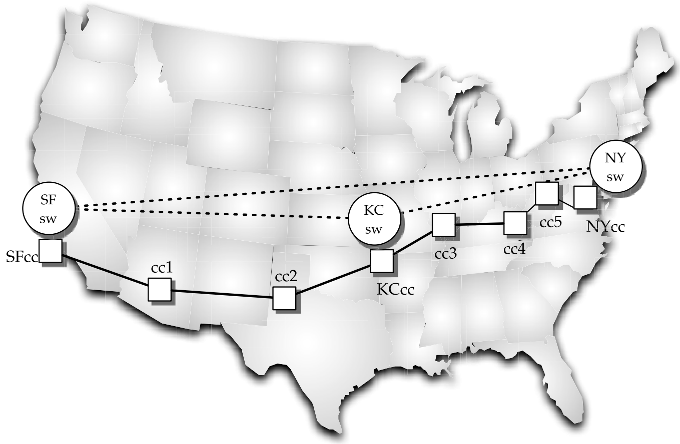

FIGURE 24.1 Transport path for logical links through cross-connects: SF-KC-NY example.

As an example, the SFsw-KCsw logical link is connected as SFsw-SFcc-cc1-cc2-KCcc-KCsw; similarly, the KCsw-NYsw logical link is connected as KCsw-KCcc-cc3-cc4-cc5-NYtn-NYsw (Figure 24.1). However, the SFsw-NYsw link is connected as SFsw-SFcc-cc1-cc2-KCcc-cc3-cc4-cc5-NYcc-NYsw; note that this link does not enter KCsw at all. Thus, we can think of the switching points for a logical link as the end points of a transport route between them.

In what sense, then, does transport network routing come into the picture? The paths shown above that are taken in the transport network are based on certain decisions; these decisions can depend on transport node functionality, port availability at nodes, and other services that need to be accommodated. ▲

It may be noted that in many instances, customers who require transport service do not want the transport service provider to change the paths since such links are carrying mission-critical traffic for which rearrangement of the path can have significant financial implications to the customer. In some instances, some customers are willing to accept it if such a rearrangement is done during a maintenance window, say 2:00 AM–4:00 AM late at night. However, with a global economy and the role communication plays, such a maintenance window, especially for international transports, is hard to find; while it is night in one part of the world, it is daytime in another part. Second, while it may be night in one part of the world, facilities and logical links in this part of the world might be of significance to users in another part where it is daytime due to corporate Intranet-based services.

Another important point is that setting up a logical link over a transport network may require a certain amount of setup time. Depending on the level of transport service requests as well as the actual bandwidth requested, such requests can take from a few days to a few months. It should be noted that a new signaling protocol suite such as GMPLS (see Section 18.4) has been developed to reduce this setup time; this is partly also necessitated by the problem faced in finding a maintenance window. However, one of the issues in setting up such a transport capability is that it involves having agreements in place between customers and providers that may include service-level agreements and payment plans, which also depend on the "size of the pipe" requested. In general, the initial setup may have a certain lag delay, while rearrangement for an existing transport connectivity may be allowable, but not always.

Some of the above decisions are also transport equipment dependent. So, what are the typical transport gears used? This can range at a data rate level: from T1/E1 to T3/E3 to OC-48, and so on. If cross-connects are for T1/E1 or T3/E3, they are known as *digital access cross-connect systems (DACS)*, or *digital cross-connect systems (DCS)*. In the optical networking domain, cross-connects are known as *optical cross-connects*. Transport functionality can be provided through a mesh network with cross-connects connected by transport links (not to be confused with logical network links). With transport networks, there are also special node types to consider, which are known as *grooming nodes*. Grooming nodes refer to the type of nodes that can perform multiplexing, say from multiple T1s to fit into a T3, or multiple OC-3s into an OC-48. There is another special node type called an *add-drop multiplexer (ADM)* which provides a grooming-type function but does not usually provide cross-connect functions. ADMs are commonly used in transport rings such as SONET rings. A key difference between a cross-connect and an ADM is that a cross-connect may be used to interconnect a larger number of transmission signals and also for grooming of transport networking routing and management.

The primary focus of this chapter is to consider transport network routing in a generic cross-connect environment for problems that fall under a Type B classification (refer to Table 17.1), without any special focus on a specific technology or a time window. Specific examples in regard to technology, such as MPLS and optical networking, are discussed later in the book, including when and how a transport service that has been classified under a Type B classification might become Type C or even Type A (refer to Table 17.1). It is worth mentioning that the general move in the industry is toward services that fall under a Type C (or Type A) classification since with new signaling technology such as GMPLS, dynamic provisioning can be accomplished. Yet, as we discussed earlier, for many customers and depending on the size of pipe requested, the overall transport network problem will continue to have a component that would need to satisfy services that would still fall under a Type B classification. Thus, our discussion in the rest of the chapter is for transport services that fall under a Type B classification.

24.2 Timing of Request and Transport Service Provisioning

From the point of view of a transport network provider, a service request from a customer may arrive on a random basis, as received by this provider. It could be at the beginning of the week or at the middle of the week or at the end of the week; furthermore, the provider may receive multiple requests on a particular day, but none on another day. Thus, the random arrival of the request cannot be avoided.

Now consider the transport service provisioning component. For the sake of this discussion, we assume that the provider provisions transport service once a week, say on Saturday evenings. However, taking into account the lag time to set up a request, the service request may take a week to several weeks to be set up. Another important point is that new capacity may be added periodically into the network based on forecast of growth.

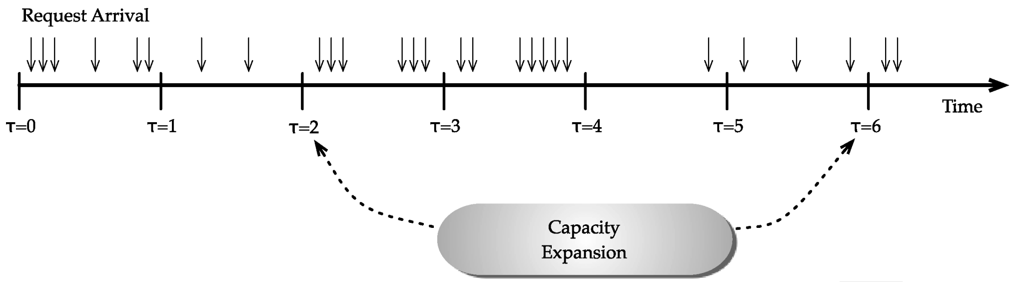

FIGURE 24.2 Transport service request arrival over a time horizon.

Example 24.2 *Transport request arrival and service provisioning: A temporal view.*

As an illustration, consider Figure 24.2. Assume that the time horizon is shown on a weekly basis. Capacity may be added every 4 weeks. Suppose that service setup has a week lag time from the end of the previous week. For example, we can see that six requests that arrive by the end of week 1 (marked as $\tau = 1$) are provisioned at the end of week 2; similarly, the two requests that arrive by the end of week 2 are provisioned at the end of week 3.

Suppose that an expedited service is also available where a request can be installed by the end of the week if the request arrives early enough in the week. Thus, out of the three arrivals in week 3, one is expedited and is served at the end of week 3, while the other two are served at the end of week 4 as normal service provisioning.

Now we discuss the effect due to capacity expansion. Capacity is added at the end of week 2 and then again at the end of week 6 based on earlier forecasts. However, during weeks 3 and 4 much higher-than-expected requests were coming in. Six requests from week 3 could be met at the end of week 4. However, of seven requests from week 4, only five could be accommodated at the end of week 5 due to capacity availability; the rest needed to wait another week until capacity expansion takes place at the end of week 6. ▲

The above example listed request arrivals without identifying origin and destination nodes. Similarly, capacity expansion can be done on a link basis; for some links, it may be added every 4 weeks, and for others it may not be; furthermore, the quantum of capacity expansion is a modular quantity such as OC-3. Another issue is the time window factor. Thus, origin–destination information would need to be taken into account in routing, and the impact of capacity expansion cannot be decoupled from the routing problem. Finally, it is possible that some provisioned services, and, thus, capacity allocated, would be released. Due to these factors, overall routing does not simply have a spatial view as in other routing problems considered in this book; it also has a *temporal* component. Another important point is that because of the scheduled service provisioning, the traffic demand matrix is *deterministic* at the time of provisioning.

It is important to note that in the above, a generic time period between updates is used to communicate the basic idea; depending on the transport technology and technological capabilities, the time period can be of different duration, typically from a week to months to a number of years. With new signaling capabilities such as for GMPLS (refer to Chapter 18), this time window can shrink. In the following, we illustrate route provisioning when demand request changes over multiple time periods.

Example 24.3 *Routing over multiple time periods* (*adapted from* [673]).

In this example, routing implications over multiple time periods are discussed for a three-node transport network (Figure 24.3). The request arrives in units of T1s—they may be for serving logical links connecting IP routers or logical links for a trunkgroup between two TDM switches.

From a transport network point of view, they are demand to be met through routing in *its* network. Capacity in the transport network is considered in T3 bandwidth units (1 T3 = 28 T1s). Figure 24.3(a) show the first time period where the first demand matrix is shown as "new request." The routes taken for these demands and any spares left on any link are shown. Note that due to modular capacity, there is no need to directly connect node 1 with node 3—its demand can be routed via node 2.

In time period $\tau = 2$ (Figure 24.3(b)), new requests are to be accommodated; this requires capacity expansion, along with rearrangement of provisioned flows. If rearrangement is not allowed, *further* capacity expansion would be needed to handle these requests.

In Figure 24.3(c), a third time period is shown, where we show that some requests are disconnected (listed as -20). Going from Figure 24.3(b) to this situation, demands can be rearranged to fit in just two transport (T3) links, thereby saving on maintenance cost on the third lik. If rearrangement is not allowed from the previous period, all three transport (T3) links would need to be active. ▲

From the above example and discussion, we can see that transport network routing can involve capacity expansion in one or more time periods. The actual routing path for demands can change over time depending on whether rearrangement is allowed. You will also see that transport network routing is essentially a minimum cost routing-type multicommodity network flow problem. If we were to consider just one time period, then the problem of transport network routing would be a multicommodity network flow (MCNF) problem presented earlier in Chapter 4; for example, see Eq. (4.4.7) for routing without any restriction and Eq. (4.5.3) if nonsplit demand is to be taken into account.

24.3 Multi-Time Period Transport Network Routing Design

By now, it should be clear that transport network routing that spans multiple time periods often requires capacity expansion to meet a traffic engineering objective. An important point about demand volume for the transport network is that the change or new request is *incremental*. It does not typically negate what is already routed in a period time period; certainly, rearrangement can be factored in, if allowed. Also, currently active circuits or routes may be disconnected because a customer might not need any more.

From a network flow modeling perspective, there are a few important issues to consider (in addition to the incremental demand volume): (1) the cost structure may change over time, e.g., due to economic discounting; thus, any cost would need to be modeled with a time-dependency parameter; (2) demand routed in one time window might have maintenance

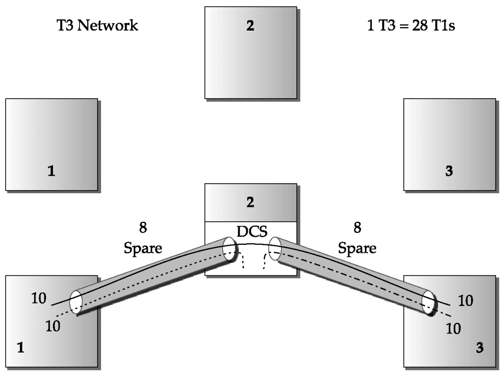

	Current	New Request	Total
1:2	0	+10	10
1:3	0	+10	10
2:3	0	+10	10

(a) Time period, $\tau = 1$

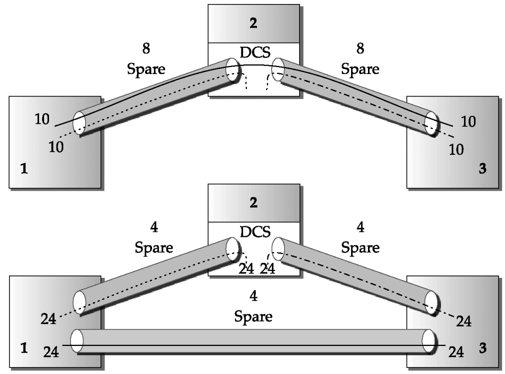

	Current	New Request	Total
1:2	10	+14	24
1:3	10	+14	24
2:3	10	+14	24

(b) Time period, $\tau = 2$

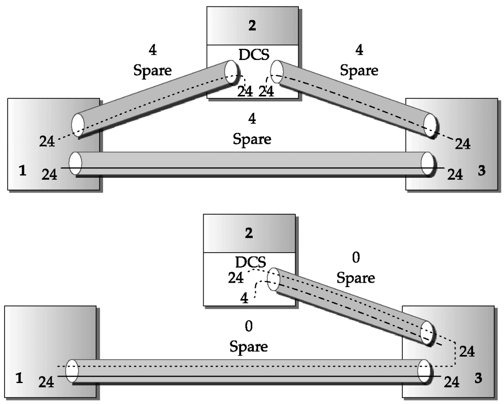

	Current	New Request	Total
1:2	24	−20	4
1:3	24	0	24
2:3	24	0	24

(c) Time period, $\tau = 3$

FIGURE 24.3 Transport routing over multiple time periods.

costs in subsequent time windows, in addition to any routing cost; and (3) capacity expansion might be possible in each time window over a given time horizon.

We assume that only one type of demand module is considered, such as all T1s, and similarly capacity is also in a specific modular value such as T3s. In simple terms, this means that without taking routing into consideration, if the demand volume is given as 51 T1s, then we need at least two T3 units of capacity (since one T3 can carry 28 T1s). Similarly, if the demand volume is 40 OC-3s, then we need three OC-48s (since one OC-48 can carry 16 OC-3s). In general, we will use M to denote the size of the modular capacity units; for example, $M = 28$ when demand volumes in T1s are considered in a network where capacity modules are in T3s, or $M = 16$ when demand volumes in OC-3s are considered in a network where capacity modules are in OC-48s. Because of M, capacity variables can be modeled as non-negative integers.

Now, going from Eq. (4.4.7), which is a single-period model, to a multiperiod model, we need to factor in a time period parameter, τ. Thus, demand volume h_k for demand pair k in Eq. (4.4.7) changes to $h_{k\tau}$ for demand pair k in time τ; as an example, demand volumes shown in different time periods in Figure 24.3 can be represented by $h_{k\tau}$. We introduce the capacity expansion variable $z_{\ell\tau}$ for capacity to be added to link ℓ in time period τ; there are two cost components associated with capacity expansion—one for new installation ($\zeta'_{\ell\tau}$) and the other for maintenance ($\zeta''_{\ell\tau}$) of capacity expanded in previous time periods; see Table 24.1 for a summary of notation.

The basic multiple time period routing design problem is to minimize cost of routing and capacity expansion; it can be written as

$$
\begin{aligned}
&\boldsymbol{minimize}_{\{\boldsymbol{xz}\}} && F = \sum_{\tau=1}^{T}\sum_{\ell=1}^{L}\left(\zeta'_{\ell\tau} z_{\ell\tau} + \zeta''_{\ell\tau}\sum_{t=1}^{\tau-1} z_{\ell t}\right) + \sum_{\tau=1}^{T}\sum_{k=1}^{K}\sum_{p=1}^{P_{k\tau}} \xi_{kp\tau} x_{kp\tau} \\
&\boldsymbol{subject\ to} && \sum_{p=1}^{P_{k\tau}} x_{kp\tau} = h_{k\tau}, \quad k = 1, 2, \ldots, K, \quad \tau = 1, 2, \ldots, T \\
& && \sum_{k=1}^{K}\sum_{p=1}^{P_{k\tau}} \delta_{kp\ell\tau} x_{kp\tau} \le M z_{\ell\tau}, \quad \ell = 1, 2, \ldots, L, \quad \tau = 1, 2, \ldots, T \\
& && x_{kp\tau} \ge 0, \quad z_{\ell\tau} = 0, 1, 2, \ldots .
\end{aligned}
\tag{24.3.1}
$$

This model, with the temporal parameter, is simply an extension of Eq. (4.4.7). To see this, suppose that the first term in the objective, which is for capacity expansion, is dropped and capacity is assumed to be given and also, only one time period ($T = 1$) is considered, then Eq. (24.3.1) reduces to Eq. (4.4.7). Since this is a link-path representation, a set of candidate paths would need to be identified at first, for example, through an algorithm such as the k-shortest path algorithm (refer to Section 2.8). Eq. (24.3.1) includes both the link installation cost and the link maintenance cost. In large network planning and design, the link installation cost is considered under capital expenditure (CapEx), while the link maintenance cost is considered under operational expenditure (OpEx). Traditionally, CapEx and OpEx are considered under separate budgetary authorities and organizations within a transport network provider. The model above shows that it is sometimes necessary to consider two different budgetary considerations under a unified model to see the overall network cost.

There is another way to represent the objective cost in Eq. (24.3.1) if we rearrange the first term using an alternate interpretation of cost related to capacity: if we install capacity in period τ, it will incur maintenance costs in *all* subsequent periods (including period τ) until the end of the planning horizon. Then, link ℓ for capacity $z_{\ell\tau}$ has the unit capacity cost,

TABLE 24.1 Summary of notation for Section 24.3.

Notation	Explanation
Given	
K	Number of demand pairs, $k = 1, 2, \dots, K$
T	Number of time periods, $\tau = 1, 2, \dots, T$
L	Number of links, $\ell = 1, 2, \dots, L$
M	Size of the link capacity module
$P_{k\tau}$	Candidate paths for flows for demand $k = 1, 2, \dots, K$ in period $\tau = 1, 2, \dots, T$
$\delta_{kp\ell\tau}$	Link-path indicator, set to 1, if link ℓ belongs to path p for demand k in time period τ, 0 otherwise
$h_{k\tau} (\geq 0)$	New (incremental) demand volume for demand k in period τ
$\zeta'_{\ell\tau}$	Installation cost of one capacity module on link ℓ for period τ
$\zeta''_{\ell\tau}$	Maintenance cost of one capacity module on link ℓ during period τ for capacity installed in prior time period(s)
$\zeta_{\ell\tau}$	Combined cost of one capacity module on link ℓ in time period τ
$\xi_{kp\tau}$	Unit routing cost on path p of demand k in time period τ
$\Theta_{k\tau}$	Penalty cost of not routing a portion of demand volume for demand k in τ
Variables	
$x_{kp\tau}$	(non-negative) flow allocated to path p of demand ℓ at time τ
$z_{\ell\tau}$	(new) capacity of link ℓ expressed in the number of modules (non-negative integer) needed in time period τ
For tracking	
$y_{\ell\tau}$	Link flow on link ℓ in time period τ
$\hat{c}_{\ell\tau}$	Spare capacity on link ℓ in time period τ

$\zeta_{\ell\tau} = \zeta'_{\ell\tau} + \sum_{t=\tau}^{T} \zeta''_{\ell\tau}$. Thus, the first cost term in the objective function in Eq. (24.3.1), can be rewritten as

$$\sum_{\tau=1}^{T}\sum_{\ell=1}^{L}\left(\zeta'_{\ell\tau} z_{\ell\tau} + \zeta''_{\ell\tau}\sum_{t=1}^{\tau-1} z_{\ell t}\right) = \sum_{\tau=1}^{T}\sum_{\ell=1}^{L} \zeta_{\ell\tau} z_{\ell\tau} \quad \text{where} \quad \zeta_{\ell\tau} = \zeta'_{\ell\tau} + \sum_{t=\tau}^{T} \zeta''_{\ell\tau}. \tag{24.3.2}$$

A benefit of writing as given in (24.3.2) is that it is easy to see that the routing design problem given by Eq. (24.3.2) can be decoupled into T-independent problems. Furthermore, the aggregated cost component, $\zeta_{\ell\tau}$, provides a sense that although CapEx and OpEx cost components need to be considered for the entire planning horizon, for modeling purposes it is not *always* necessary to model them completely separately, at least for models such as the given by Eq. (24.3.1). Certainly, the fact that model (24.3.1) can be decoupled into T-independent problems raises the question on whether multiperiod modeling is necessary. We will now discuss two basic problems with Eq. (24.3.1).

In Eq. (24.3.1), we have assumed that the incremental demand is non-negative, which reflects network growth over the planning horizon. It is certainly possible to imagine the case where installed demand volume from a previous period is no longer needed in a future period (negative growth in a network), for example, disconnection in a future period of circuits already installed in a previous period in the case of transport networks. As discussed in [400], a way to capture this effect is to have $h_{k\tau} < 0$ (refer to Figure 24.3(c)), which would imply that previously routed demand volumes need to be altered and that we must allow $x_{kp\tau} < 0$ in the formulation. However, we need to ensure that this decrease is accommodated only on paths that have positive flows in prior periods. To do this, we need to replace the requirement that each flow ($x_{kp\tau}$) is non-negative with the following condition:

$$\sum_{t=1}^{\tau} x_{kpt} \geq 0, \quad \tau = 1, 2, \ldots, T,$$

along with the understanding that path index p, in this case, refers to the exact same path from one time period to the next for the same demand identifier, k. The inclusion of this constraint in Eq. (24.3.1) implies that the modified design problem can no longer be naturally decoupled into T-independent design problems. Although this new constraint satisfies feasibility, the restriction on demand routing (i.e., routing for any new incremental demand volume to be performed on a period-by-period basis) is no longer maintained; in other words, rearrangeability of routed demand volume from one period to the next is possible. While this flexibility is good from a formulation point of view, the rearrangeability option may not be allowable/possible for many real transport networks.

For the rest of the discussion, we assume that the incremental demand volume is $h_{kt} \geq 0$ and that the rearrangement of routed demand from one period to the next is not allowable. Thus, we return to Problem (24.3.1) and the fact that this problem can still be decoupled into T-independent single-period problems. Therefore, we will now discuss another important reason to consider multiperiod modeling instead of using just single-period design.

If you follow model (24.3.1) carefully, you will notice that this model does not necessarily generate optimal solutions from the standpoint of overall network capacity over the entire planning horizon. For example, from the second set of constraints in Eq. (24.3.1), it is easy to see that due to modularity of capacity installed, *not all* capacity that was installed in a previous period may be completely depleted by routing of demand volume in that period. Thus, in actuality, there is a good chance that some *spare* capacity will be available from one time period to the next (refer to Example 24.3), which can be used for realizing routing of demand volumes in future periods; this aspect is not explicitly considered in the above model. Thus, in reality, there is a natural coupling from one time period to another in a multiperiod problem.

To illustrate the effect of spare capacity from one period for use in future periods for flow routing, we denote $y_{\ell\tau}$ as the link load on link ℓ in period τ ($\tau = 1, 2, \ldots, T$). Furthermore, we denote the spare capacity on link ℓ in period τ by $\hat{c}_{\ell\tau} \geq 0$ for $\tau = 0, 1, 2, \ldots, T$. In this case, $\hat{c}_{\ell 0}$ denotes any spare capacity available at the beginning of the entire planning cycle. Now at the end of time $\tau = 1$, any new, incremental demand volume must be satisfied using already available capacity at the beginning of this period plus any new capacity added in this period; thus, we have the following link-load satisfiability condition for $\tau = 1$:

$$y_{\ell 1} \leq \hat{c}_{\ell 0} + M z_{\ell 1}, \quad \ell = 1, 2, \ldots, L.$$

Then, the spare capacity (if any) left at the end of period $\tau = 1$ is available in period $\tau = 2$; this spare capacity can be written as

$$\hat{c}_{\ell 1} = \hat{c}_{\ell 0} + M z_{\ell 1} - y_{\ell 1}, \quad \ell = 1, 2, \ldots, L.$$

TABLE 24.2 Link flow, capacity expansion, and spare capacity in time period $\tau = 2$ (see Figure 24.3(b)).

Link ℓ	$c_{\ell 1}$	$z_{\ell 2}$	$y_{\ell 2}$	$c_{\ell 2}$
1-2	8	0	20	4
1-3	0	1	20	4
2-3	8	0	20	4

Similarly, at the end of period $\tau = 2$, the link-load satisfiability condition and the spare capacity can be written as:

$$y_{\ell 2} \le \hat{c}_{\ell 1} + M z_{\ell 2}, \quad \ell = 1, 2, \ldots, L$$

$$\hat{c}_{\ell 2} = \hat{c}_{\ell 1} + M z_{\ell 2} - y_{\ell 2}, \quad \ell = 1, 2, \ldots, L,$$

respectively. This is illustrated in Table 24.2 for data in Example 24.3.

Generalizing, we have

$$y_{\ell \tau} \le \hat{c}_{\ell, \tau-1} + M z_{\ell \tau}, \quad \ell = 1, 2, \ldots, L, \ \tau = 1, 2, \ldots, T$$

$$\hat{c}_{\ell \tau} = \hat{c}_{\ell, \tau-1} + M z_{\ell \tau} - y_{\ell \tau}, \quad \ell = 1, 2, \ldots, L, \ \tau = 1, 2, \ldots, T.$$

Essentially, we need to incorporate these two sets of relations into model (24.3.1) to account for reuse of spare capacity from one period to the next. Using substitution, we can rewrite spare capacity in period τ (for each link ℓ) as

$$\begin{aligned}
\hat{c}_{\ell \tau} &= \hat{c}_{\ell, \tau-1} + M z_{\ell \tau} - y_{\ell \tau} \\
&= \hat{c}_{\ell, \tau-2} + M z_{\ell, \tau-1} - y_{\ell, \tau-1} + M z_{\ell \tau} - y_{\ell \tau} \\
&\quad \ldots \\
&= \hat{c}_{\ell 0} + M z_{\ell \tau} - y_{\ell \tau} + \textstyle\sum_{t=1}^{\tau-1} (M z_{\ell t} - y_{\ell t}).
\end{aligned}$$

Rearranging, we get

$$y_{\ell \tau} + \hat{c}_{\ell \tau} = \hat{c}_{\ell 0} + M z_{\ell \tau} + \textstyle\sum_{t=1}^{\tau-1} (M z_{\ell t} - y_{\ell t}).$$

Since spare capacity is non-negative, we can arrive at the following inequality:

$$y_{\ell \tau} \le \hat{c}_{\ell 0} + M z_{\ell \tau} + \textstyle\sum_{t=1}^{\tau-1} (M z_{\ell t} - y_{\ell t}).$$

Finally, if we denote the initial spare capacity as $\hat{c}_{\ell 0} = M z_{\ell 0} - y_{\ell 0}$ for time $\tau = 0$, if there is any spare capacity available (and 0, otherwise), we have the following relation

$$y_{\ell \tau} \le M z_{\ell \tau} + \textstyle\sum_{t=0}^{\tau-1} (M z_{\ell t} - y_{\ell t}), \quad \ell = 1, 2, \ldots, L, \ \tau = 1, 2, \ldots, T. \tag{24.3.3}$$

Thus, constraint (24.3.3) is an important constraint to include with Eq. (24.3.1) to capture spare capacity availability and use, and to model the transport network routing design problem accurately.

DEMAND REQUEST AND CAPACITY EXPANSION: DIFFERENT TIME CYCLE

So far, we have assumed that capacity expansion is possible in every time period. In many practical situations, this may not be the case. Consider the case where new requests are collected every week for routing, while capacity can be expanded only every 4 weeks due to, say logistics issues (refer to Figure 24.2). Thus, we cannot then rule out the possibility that in some periods, there may not be enough capacity to route all requests. To factor in this possibility, we introduce an artificial variable, $\tilde{x}_{k\tau}$, for each demand k in time period τ, which captures any demand volume that cannot be met by the currently available bandwidth. This variable, however, does not appear in the link capacity constraints. It is indeed important to introduce a per-unit penalty parameter, $\Theta_{k\tau} > 0$, if demand cannot be accommodated; this can then serve, for example, as the cost incurred due to lost revenue and the penalty can be set high, as appropriate. Furthermore, since capacity cannot be added in every time period, we must force capacity expansion variable, $z_{\ell\tau}$, to be zero in the periods in which capacity expansion is not allowed. To help distinguish, we will denote $\widehat{\mathcal{T}}$ to indicate the time periods when capacity expansion is allowed. If we denote the set of all time periods as $\mathcal{T}$, then the set difference, $\mathcal{T}\backslash\widehat{\mathcal{T}}$, represents the periods when no capacity expansion is possible; that is, for these periods, $z_{\ell\tau} = 0$. Thus, we can write the overall model as

$$
\begin{aligned}
\textbf{\textit{minimize}}_{\{xz\}} \quad & F = \sum_{\tau=1}^{T}\sum_{\ell=1}^{L} \zeta_{\ell\tau} z_{\ell\tau} + \sum_{\tau=1}^{T}\sum_{k=1}^{K}\sum_{p=1}^{P_{k\tau}} \xi_{kp\tau} x_{kp\tau} + \sum_{\tau=1}^{T}\sum_{k=1}^{K} \Theta_{k\tau}\tilde{x}_{k\tau} \\
\textbf{\textit{subject to}} \quad & \sum_{p=1}^{P_{k\tau}} x_{kp\tau} + \tilde{x}_{k\tau} = h_{k\tau}, \qquad k = 1, 2, \ldots, K, \ \tau = 1, 2, \ldots, T \\
& \sum_{k=1}^{K}\sum_{p=1}^{P_{k\tau}} \delta_{kp\ell\tau} x_{kp\tau} \le M z_{\ell\tau}, \quad \ell = 1, 2, \ldots, L, \ \tau = 1, 2, \ldots, T \\
& \sum_{k=1}^{K}\sum_{p=1}^{P_{k\tau}} \delta_{kp\ell\tau} x_{kp\tau} = y_{\ell\tau}, \quad \ell = 1, 2, \ldots, L, \ \tau = 1, 2, \ldots, T \\
& y_{\ell\tau} \le M z_{\ell\tau} + \textstyle\sum_{t=0}^{\tau-1} (M z_{\ell t} - y_{\ell t}), \quad \ell = 1, 2, \ldots, L, \ \tau = 1, 2, \ldots, T \\
& z_{\ell\tau} = 0 \quad \text{for } \tau \in \mathcal{T}\backslash\widehat{\mathcal{T}} \\
& z_{\ell\tau} = 0, 1, 2, \ldots \ \text{(integer)} \quad \text{for } \tau \in \widehat{\mathcal{T}} \\
& x_{kp\tau} \ge 0, \tilde{x}_{k\tau} \ge 0.
\end{aligned}
\tag{24.3.4}
$$

At first look, this appears to be a complicated model. In fact, the basic idea is quite simple. It addresses multiperiod transport routing design where capacity expansion need not be in the same time period cycles as new demand requests. Furthermore, a penalty is introduced if a demand request cannot be met in a time period due to lack of available capacity. Finally, link flow and spare capacity are tracked for going from one period to another.

OTHER VARIATIONS

There are other possible variations that can depend on transportation technology and network capability; for example, flow variables, $x_{kp\tau}$, can take discrete values, instead of taking continuous values; multiple heterogeneous size modules, instead of one, are allowed. Such requirements can be accommodated by extending the model presented here. It may be noted that not all cost components may be necessary for a particular transport routing design problem; accordingly, such cost components can be dropped from the objective function. Finally, it is possible to consider transport routing where full rearrangement for existing routed demands is allowed. In this case, the main difference is that the spare capacity relation, given by Eq. (24.3.3), is not necessary.

24.4 Transport Routing with Varied Protection Levels

In this section, we present a transport routing problem that takes into account protection requirements at different levels for different service requests. For simplicity, we will present this model for a single time period. In this model, we consider three protection service classes: full-, fractional-, and zero-protection transport services. Here, zero protection means that the service is guaranteed under normal operating conditions but not under a failure; fractional-protection means providing a reduced level of services under a major link failure in addition to guaranteed service under normal operating conditions; finally, full protection means providing guarantee under both normal as well as failure conditions.

For each origin–destination node pair in the network, demand volume must be routed with one of the protection levels for a set of customers where each customer might have a different protection service-level agreement; thus, customers are categorized into different service classes based on the protection-level agreement. Also, customers demand requests for a particular demand pair k are nonsplittable on multiple paths. We consider a demand pair k where we have a set of service requests $\mathcal{S}_k$ from customers requiring protection services at differing protection levels. See Table 24.3 for a summary of notation.

TABLE 24.3 Summary of notation for Section 24.4 (note: K and L are defined in Table 24.1).

Notation	Explanation
$\mathcal{S}_k$	Protection service classes for demand $k = 1, 2, \ldots, K$
d_k^s	Demand volume of service request s for demand pair k
P_k^s	Set of candidate path cycles for service request $s = 1, 2, \ldots, S_k$ for demand $k = 1, 2, \ldots, K$
c_ℓ	Capacity of link $\ell = 1, 2, \ldots, \mathcal{L}$
α_k^s	Protection level of service request $s = 1, 2, \ldots, S_k$ for $k = 1, 2, \ldots, K$
$\delta_{km}^{s\ell}$	Link-primary path indicator for a path cycle; 1, if candidate path cycle $m = 1, 2, \ldots, P_k^s$ for service $s = 1, 2, \ldots, S_k$ of demand pair $k = 1, 2, \ldots, K$ uses link $\ell = 1, 2, \ldots, L$ in its primary path; 0, Otherwise
$\beta_{km}^{s\ell}$	Link-backup path indicator for a path cycle; 1, if candidate path cycle $m = 1, 2, \ldots, P_k^s$ for service $s = 1, 2, \ldots, S_k$ of demand pair $k = 1, 2, \ldots, K$ uses link $\ell = 1, 2, \ldots, L$ in its backup path; 0, Otherwise
Θ_k^s	Penalty cost of not routing service request s for demand k
Variables:	
u_{km}^s	0/1 decision variable for choosing path cycle m for s, k
$\tilde{u}_k^s$	0/1 artificial variable for s, k

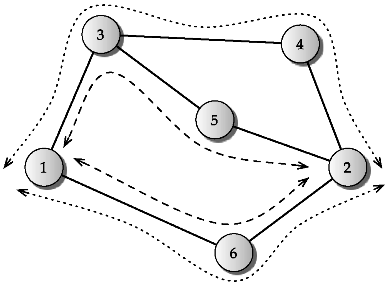

FIGURE 24.4 Illustration of path cycles: 1-6-2-5-3-1 and 1-6-2-4-3-1.

Consider now the zero-protection transport level demand request. In this case, only a path with bandwidth d_k^s needs to be provisioned. Considering only the shortest (e.g., in terms of hops) path may not address the overall traffic engineering goal. Thus, we need to consider a set of candidate paths for each demand k, as we have done with other models presented earlier for link-path formulation model.

For the full-protection transport service class, a backup path needs to be available and bandwidth d_k^s needs to be reserved on the backup path. We require that the backup path survive if the primary path is affected due to any critical failure, e.g., for a single link failure at a time. Although the primary and backup path could be independently modeled, we use a pairing idea, i.e., consider a pair of disjoint paths consisting of primary and backup paths. Similar to the case of zero-protection services, the selection of the shortest pair of primary and backup paths for a demand d_k^s may not be in the best interest of a traffic engineering objective. Thus, we can consider a *candidate set* of primary/backup path pairs for a demand volume, d_k^s, for full protection.

Finally, in the case of fractional-protection transport services, the backup path needs to be allocated bandwidth that is sufficient to carry a fraction of d_k^s in order to address partial survivability. Thus, the fractional-transport service class also requires a pair of disjoint paths. The difference is that, on the backup path, only a fraction of the demand is required to be reserved. If we denote the fraction by α_k^s (where $0 \le \alpha_k^s \le 1$), then the primary path would reserve d_k^s while the backup path would reserve $\alpha_k^s d_k^s$.

When we consider all three cases, it is easy to see that by appropriately setting α_k^s we can consider each of the protection levels, i.e., $\alpha_k^s = 0$ refers to zero protection, $\alpha_k^s = 1$ refers to full protection, while α_k^s refers to fractional protection if $0 < \alpha_k^s < 1$.

There are two benefits of introducing α_k^s: (1) fractional-protection need not be of a specific predefined value; each customer can request a different level, and (2) for zero protection, we can also consider a pair of disjoint paths as well where on the backup path, we assign $\alpha_k^s = 0$; this means we can still consider a backup path but it is not used. Consequently, the three service classes can be considered in a unified manner from a modeling framework—all we need to do is to consider a set of candidate pairs of disjoint paths; there are known algorithms for generation such path pairs [677], [678]. For simplicity, we refer to a pair of disjoint paths as a *path cycle*. As an illustration, consider Figure 24.4, where for the demand pair connecting nodes 1 and 2, we have two different candidate path cycles 1-6-2-5-3-1 and 1-6-2-4-3-1, which are link-disjoint; such candidate path cycles are considered as input to the problem formulation.

We denote the set of candidate path cycles for service s for demand pair k by P_k^s. Suppose we associate u_{km}^s as the decision variable with candidate cycle m, then for each $s = 1, 2, \ldots, S_k$, $k = 1, 2, \ldots, K$, the decision to select only one cycle is governed by the following requirement:

$$\sum_{m=1}^{P_k^s} u_{km}^s + \tilde{u}_k^s = 1, \quad s = 1, 2, \ldots, S_k,\ k = 1, 2, \ldots, K \tag{24.4.1}$$

if we also introduce the artificial (slack) variable, $\tilde{u}_k^s$, to allow for demand that cannot be met due to capacity limitation. Since we are using $0|1$ (binary) variables, a demand can not be partially routed. Now recall that for each path cycle (because of the way each of them are generated) we have a primary path and the backup path. Using link-primary path and link-backup path indicators (see Table 24.3), the flow on link ℓ (denoted by y_ℓ) to carry demand volume under both normal and failure situations (for primary and backup path for different demand requests) must be less than the capacity:

$$\sum_{k=1}^{K}\sum_{s=1}^{S_k}\sum_{m=1}^{P_k^s} \left(\delta_{km}^{s\ell} + \alpha_k^s \beta_{km}^{s\ell}\right) d_k^s u_{km}^s \le c_\ell, \quad \ell = 1, 2, \ldots, L. \tag{24.4.2}$$

The link flow, given on the left side of the inequality, is a generalization of the link flow discussed earlier to include the fact that a link may have fractional traffic volume allocated for a backup path of a service request. Here note parameter α_k^s dictates the level of protection on the backup path; also since we are considering a path cycle consisting of disjoint paths, for a specific link ℓ, if $\delta_{km}^{s\ell}$ takes the value 1, then the corresponding $\beta_{km}^{s\ell}$ must be zero, and vice versa.

To address minimum cost routing with penalty for demand not met, we can write the objective function as

$$F = \sum_{k=1}^{K}\sum_{s=1}^{S_k}\sum_{m=1}^{P_k^s} \xi_{km}^s d_k^s u_{km}^s + \sum_{k=1}^{K}\sum_{s=1}^{S_k} \Theta_k^s \tilde{u}_k^s. \tag{24.4.3}$$

Thus, the network traffic engineering problem for transport routing design with varied protection is to minimize (24.4.3) subject to constraints (24.4.1) and (24.4.2).

24.5 Solution Approaches

Transport network routing problems presented in the previous sections are in general classified as minimum cost routing problems and problem formulations are multicommodity network flow-based. As can be seen, the size of the problem in terms of unknowns and constraints grows when multiple time periods are considered. In many instances, problems are integer linear programming problems in nature due to modularity of demand in transport routing. However, a linear programming approximation can be considered by relaxing the integrality constraints. An advantage of using the linear programming approximation is that tools such as CPLEX can be used to solve large problems quite efficiently. Some solutions will not be integral due to this relaxation; from this solution, a post processing heuristic rule can be developed to round up or down to obtain integer solutions. If CPLEX is run for solving an integer programming problem, then it is advisable to limit the number of branching nodes to, say 30,000 (i.e., "set mip limits nodes 30000" in CPLEX) so that a good solution can be found quickly.

For additional advanced approaches that exploit the special structure of a transport routing problem, there are many specialized algorithms [564]. It may be noted that the transmission routing problem needs to be solved only periodically, say, once a week or every other week; thus, quickness of generating a solution is not always the primary driver. The primary driver might involve the following: does reducing the overall cost by 1% lead to a significant cost saving at the expense of an increase in time for computing such a solution that is allowable within the provisioning time frame? If so, then such guidelines can be explored.

24.6 Summary

In this chapter, we have presented transport network routing. We started by first explaining why transport network routing problems arise.

We then explained how the time of arrival and service provisioning help to define why transport network routing falls under a Type B classification (see Table 17.1). We then illustrated through a small network example how time period factors and rearrangement (or restriction on rearrangement) can play roles in the solution. We then presented two sets of transport network routing models to discuss the intricacies involved in transport network routing design.

Further Lookup

Discussions about early days of transport networks can be found [596]. Early work on transport network routing through a minimum cost routing approach can be found in [751], [752], [758], [759], including first sets of work on multiple time period routing. Since then, multiperiod design has been addressed by many researchers; for example, see [195], [400], and [563]. In recent years, there has been a surge in understanding transport networks for IP networking; for example, see [57], [186]. The notion of using path cycles in a link-path formulation setting was originally presented in [466]. For a recent discussion on varied protection design for different objective functions, see [665].

Exercises

24.1 Extend Eq. (24.3.4) when (a) flow variables are to be integer valued; (b) multiple modular capacity values are allowed, and (c) if flow for a demand is unsplittable.

24.2 Extend the traffic engineering model presented in Section 24.4 to multiple time periods, with or without rearrangements.

24.3 Identify any changes required to Eq. (24.3.1) if the capacity expansion falls on different time windows than the demand cycle.

24.4 Discuss applicability of the model presented in Section 24.4 for protection design of MPLS networks.

24.5 Consider the various objectives discussed in Chapters 4 and 7. Discuss their applicability to the protection design model presented in Section 24.4.