Chapter 3: Geometric Capacity Provisioning for Wavelength-Switched WDM Networks

Li-Wei Chen and Eytan Modiano

L.-W. Chen
Vanu Inc., One Cambridge Center, Cambridge, MA 02142
e-mail: lwchen@alum.mit.edu

E. Modiano (*)
Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA 02139
e-mail: modiano@mit.edu

Portions reprinted, with permission, from "A Geometric Approach to Capacity Provisioning in WDM Networks with Dynamic Traffic", *40th Annual Conference on Information Sciences and Systems*. © 2006 IEEE.

Abstract

In this chapter, we use an asymptotic analysis similar to the sphere-packing argument in the proof of Shannon's channel capacity theorem to derive optimal provisioning requirements for networks with both static and dynamic provisioning. We consider an N -user shared-link model where W_s wavelengths are statically assigned to each user, and a common pool of W_d wavelengths are available to all users. We derive the minimum values of W_s and W_d required to achieve asymptotically non-blocking performance as the number of users N becomes large. We also show that it is always optimal to statically provision at least enough wave-lengths to support the mean of the traffic.

3.1 Introduction

Optical networking has established itself as the backbone of high-speed communication systems, incorporating both high bandwidth and low noise and interference characteristics into a single medium. Within optical networks, wavelength division multiplexing (WDM) technology has emerged as an attractive solution for exploiting the available fiber bandwidth to meet increasing traffic demands. WDM divides the usable bandwidth into non-overlapping frequency bands (usually referred to as *wavelengths* in the literature) and allows the same fiber to carry many signals independently by assigning each signal to a different wavelength.

In general, an optical WDM network can consist of a large number of nodes connected in some arbitrary fashion (see Figure 3.1) and can present the network architect with a complex wavelength provisioning problem over multiple links. For simplicity, in this chapter, we will focus on provisioning a single shared link on a backbone network. Figure 3.1 also shows a model for the shared link in the arbitrary network. We consider provisioning for traffic traveling from left to right along the link. Each wavelength on the link can be used to support one lightpath from one of the incoming fibers on the left side of the link to one of the outgoing fibers on the right side of the link.

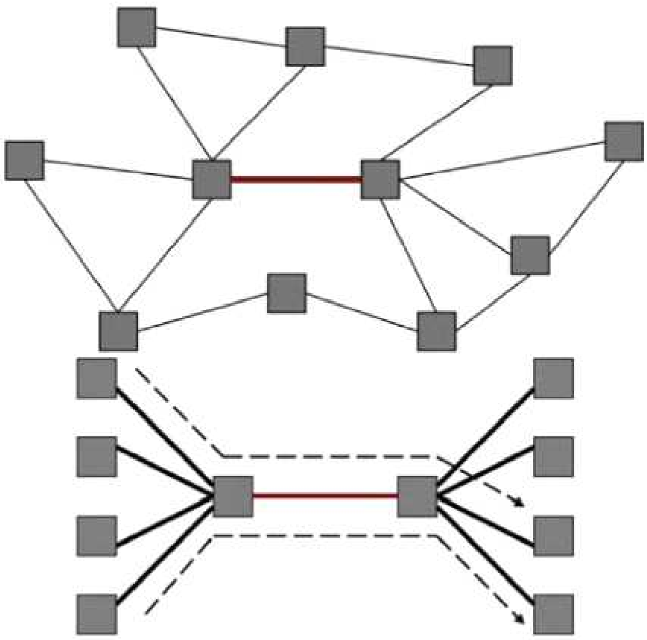

Figure 3.1: An example of a mesh optical network consisting of numerous nodes and links, followed by a shared-link model based on the link. The dotted lines denote different users of the link. Since each pair of input—output fibers comprises a different user, and there are four input fibers and four output fibers, there are a total of 4 · 4 = 16 users in this example

Broadly speaking, wavelength provisioning can be done in one of two ways. One option is to *statically* provision a wavelength by hard-wiring the nodes at the ends of the link to always route the wavelength from a given input fiber to a given output fiber. The advantage to this is that the cost of the hardware required to support static provisioning is relatively low: no switching capability or intelligent decision-making ability is required. The downside is a lack of flexibility in using that wavelength — even if the wavelength is not needed to support a lightpath between the assigned input and output fibers, it cannot be assigned to support a lightpath between any other pair of fibers.

This shortcoming can be overcome by using *dynamic* provisioning. A dynamically provisioned wavelength is switched at the nodes on both sides of the link, allowing it to be dynamically assigned to support a lightpath between any source and destination fibers. Furthermore, this assignment can change over time as traffic demands change. This obviously imparts a great deal of additional flexibility. The downside is that the added switching and processing hardware makes it more expensive to dynamically provision wavelengths.

There has been much investigation of both statically provisioned and dynamically provisioned systems in the literature [1–4]. Such approaches are well suited for cases where either the traffic is known a priori and can be statically provisioned, or is extremely unpredictable and needs to be dynamically provisioned. However, in practice, due to statistical multiplexing, it is common to see traffic demands characterized by a large mean and a small variance around that mean. A hybrid system is well suited to such a scenario. In a hybrid system, a sufficient number of wave-lengths are statically provisioned to support the majority of the traffic. Then, on top of this, a smaller number of wavelengths are dynamically provisioned to support the inevitable variation in the realized traffic. Such an approach takes advantage of the relative predicability of the traffic by cheaply provisioning the majority of the wavelengths, but retains sufficient flexibility through the minority of dynamic wavelengths that significant wavelength overprovisioning is not necessary.

After describing the system model used in this chapter, we will use the asymptotic analysis approach from information theory incorporated in the proof of Shannon's channel capacity theorem [5] to analyze hybrid networks: we allow the number of users to become large, and consider the minimum provisioning in static and dynamic wavelengths necessary to achieve non-blocking performance (i.e., to guarantee that the probability of any call in the snapshot being blocked goes to zero). We will show that it is always optimal to statically provision enough wavelengths to support the traffic mean. We also fully characterize the optimal provisioning strategy for achieving non-blocking performance with minimal wavelength provisioning.

3.1.1 System Model

In the shared link context, we can consider each incoming—outgoing pair of fibers to be a different *user* of the link. Each lightpath request (which we will henceforth term a *call*) can therefore be thought of as belonging to the user corresponding to the incoming—outgoing fiber pair that it uses. We can similarly associate each static wavelength with the corresponding user. Under these definitions, a call belonging to a given user cannot use a static wavelength belonging to a different user — it must either use a static wavelength belonging to its own user, or employ a dynamic wavelength.

Figure 3.2 gives a pictorial representation of the decision process for admitting a call. When a user requests a new call setup, the link checks to see if a static wavelength for that user is free. If there is a free static wavelength, it is used. If not, then the link checks to see if any of the shared dynamic wavelengths are free — if so, then a dynamic wavelength is used. If not, then no resources are available to support the call, and it is blocked.

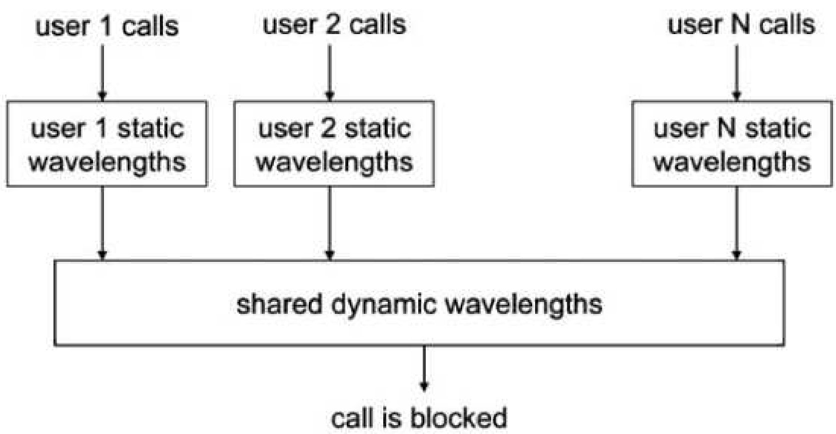

Figure 3.2: Decision process for wavelength assignment for a new call arrival. A new call first tries to use a static wavelength if it is available. If not, it tries to use a dynamic wavelength. If again none are available, then it is blocked

There have been several approaches developed in the literature for blocking probability analysis of such systems under Poisson traffic models [6], including the Equivalent Random Traffic (ERT) model [7–9] and the Hayward approximation [10]. These approaches, while often able to produce good numerical approximations of blocking probability, are purely numerical in nature and do not provide good intuition for guiding the dimensioning of the wavelengths.

In this chapter, we adopt a snapshot traffic model that leads to closed-form asymptotic analysis and develop guidelines for efficient dimensioning of hybrid networks. We consider examining a "snapshot" of the traffic demand at some instant in time. The snapshot is composed of the vector $\mathbf{c} = [c_1;\ldots; c_N]$, where c_i is the number of calls that user i has at the instant of the snapshot, and N is the total number of users.

We model each variable c_i as a Gaussian random variable with mean μ_i and variance σ_i^2. This is reasonable since each "user" actually consists of a collection of source—destination pairs in the larger network that all use the link from the same source fiber to the same destination fiber. Initially we will assume that each user has the same mean μ and variance σ^2 and later extend the results to general μ_i and σ_i. Although the traffic for each individual source—destination pair for the user may have some arbitrary distribution, as long as the distributions are well behaved, the sum of each traffic stream will appear Gaussian by the Central Limit Theorem.

As a special case, consider the common model of Poisson arrivals and exponential holding times for calls. Then the number of calls that would have entered a non-blocking system at any instant in time is given by the stationary distribution of an $M/M/\infty$ queue — namely, Poisson with intensity equal to the load ρ in Er-langs. For a heavy load, this distribution is well approximated by a Gaussian random variable with mean ρ and variance ρ.

3.2 Wavelength-Granularity Switching

In this section, we consider a shared link, and assume that there are N users that are the source of calls on the link. Each user is statically provisioned W_s wavelengths for use *exclusively* by that user. In addition to this static provisioning, we will also provide a total of W_d dynamically switched wavelengths. These wavelengths can be shared by any of the N users.

As previously described, we will use a snapshot model of traffic. The traffic is given by a vector $\mathbf{c} = [c_1;\ldots; c_N]$, where each c_i is independent and identically distributed as $N(\mu,\sigma^2)$. We assume that the mean μ is sufficiently large relative to σ that the probability of "negative traffic" (where the realized value of a random variable representing the number of calls is negative, a physical impossibility) is low, and therefore does not present a significant modeling concern. We will primarily be concerned with a special blocking event that we call *overflow*. An overflow event occurs when there are insufficient resources to support all calls in the snapshot and at least one call is blocked. We will call the probability of this event the *overflow probability*.

From Figure 3.2, we see that an overflow event occurs if the total number of calls exceeds the ability of the static and dynamic wavelengths to support them. This can be expressed mathematically as

$$\sum_{i=1}^{N} \max\{c_i - W_s\ ,\ 0\} > W_d\,, \tag{3.1}$$

where $\max\{c_i - W_s\ , 0\}$ is the amount of traffic from each user that exceeds the static provisioning; if the total amount of excess from each user exceeds the available pool of shared dynamic wavelengths, a blocking event occurs.

If we consider the N-dimensional vector space occupied by $\mathbf{c}$, the constraint given by (3.1) represents a collection of hyperplanes bounding the admissible traffic region:

$$\begin{aligned} c_i &\le W_s + W_d \\ c_i + c_j &\le 2W_s + W_d, \quad i \ne j, \\ c_i + c_j + c_k &\le 3W_s + W_d, \quad i \ne j \ne k \\ &\vdots \end{aligned}$$

Each constraint reflects the fact that the sum of the traffic from any subset of users clearly cannot exceed the sum of the static provisioning for those users plus the entire dynamic provisioning available. Note that there are a total of N sets of constraints, where the nth set consists of $C(N,n) = \frac{N!}{(N-i)!n!}$ equations, each involving the sum of n elements of the traffic vector $\mathbf{c}$. If the traffic snapshot $\mathbf{c}$ falls within the region defined by the hyperplanes, all calls are admissible; otherwise, an overflow event occurs. The bold lines in Figure 3.3 show the admissible region for $N = 2$ in two dimensions.

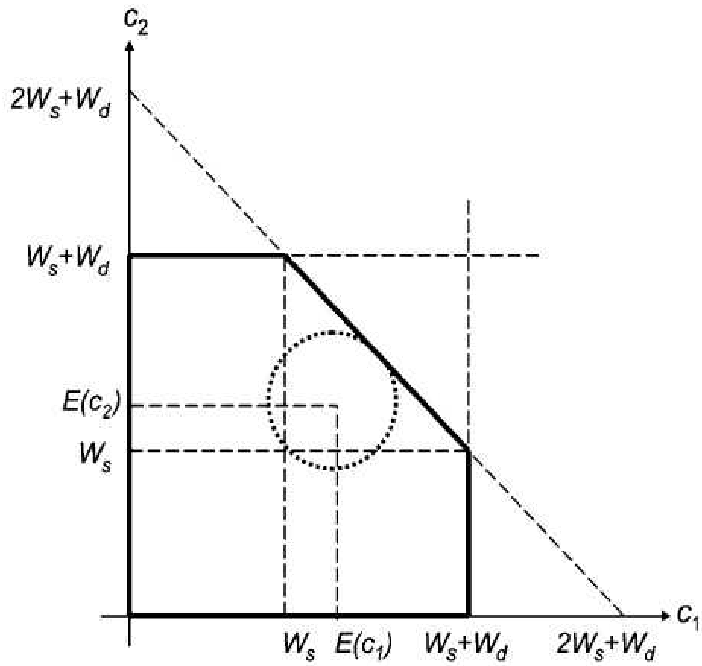

Figure 3.3: The admissible traffic region, in two dimensions, for $N = 2$. Three lines form the boundary constraints represented by (3.1). There are two lines each associated with a single element of the call vector c, and one line associated with both elements of c. The traffic sphere must be entirely contained within this admissible region for the link to be asymptotically non-blocking

3.2.1 Asymptotic Analysis

We will consider the case where the number of users N becomes large, and use the law of large numbers to help us draw some conclusions. We can rewrite the call vector in the form

$$\mathbf{c} = \mu \cdot \mathbf{1} + \mathbf{c}',$$

where μ is the (scalar) value of the mean, $\mathbf{1}$ is the length-N all-ones vector, and $\mathbf{c}' \sim \mathbf{N}(\mathbf{0}, \sigma^2\mathbf{1})$ is a zero-mean Gaussian random vector with i.i.d. components. Conceptually, we can visualize the random traffic vector as a random vector $\mathbf{c}'$ centered at $\mu\mathbf{1}$. The length of this random vector is given by

$$\|\mathbf{c}'\| = \sqrt{\sum_{i=1}^{N} c_i'^2}.$$

We use an approach very similar to the sphere packing argument used in the proof of Shannon's channel capacity theorem in information theory [5]. We will show that asymptotically as the number of users becomes large, the traffic vector falls onto a sphere centered at the mean, and the provisioning becomes a problem of choosing the appropriate number of static and dynamic wavelengths so that this traffic sphere is completely contained within the admissible region.

From the law of large numbers, we know that

$$\frac{1}{N}\sum_{i=1}^{N} c_i'^2 \to \sigma^2$$

as $N \to \infty$. This implies that asymptotically, as the number of users becomes large, the call vector $\mathbf{c}$ becomes concentrated on a sphere of radius $\sqrt{N}\sigma$ centered at the mean $\mu\mathbf{1}$. (This phenomenon is known in the literature as *sphere hardening*.) Therefore, in order for the overflow probability to converge to zero, a necessary and sufficient condition is that the hyperplanes described by (3.1) enclose the sphere entirely. This is illustrated in Figure 3.3.

3.2.2 Minimum Distance Constraints

Next, we will derive necessary and sufficient conditions for the admissible traffic region to enclose the traffic sphere. Our goal is to ensure that we provision W_s and W_d such that the minimum distance from the center of the traffic sphere to the boundary of the admissible region is at least the radius of the sphere, therefore ensuring that all the traffic will fall within the admissible region.

Due to the identical distribution of the traffic for each user, the mean point $\mu\mathbf{1}$ will be equidistant from all planes whose description involves the same number of elements of **c**. We define a *distance function* $f(n)$ such that $f(n)$ is the minimum distance from the mean $\mu\mathbf{1}$ to any hyperplane whose description involves n components of **c**.

Lemma 3.1

The distance function $f(n)$ *from the traffic mean to a hyperplane involving* n *elements of the traffic vector* **c** *is given by*

$$f(n) = \sqrt{n}\left(W_s + \frac{W_d}{n} - \mu\right), \quad n = 1, \dots, N \tag{3.2}$$

Proof. This is essentially a basic geometric exercise. For a fixed n, the hyperplane has a normal vector consisting of n unity entries and $N - n$ zero entries. Since by symmetry the mean of the traffic is equidistant from all hyperplanes with the same number of active constraints, without loss of generality, assume the first n constraints that are active. Then the closest point on the hyperplane has the form

$[\mu + x, \dots \mu + x, \mu, \dots, \mu]$

where the first n entries are $\mu + x$, and the remainder are μ. The collection of hyperplanar constraints described by (3.1) can then be rewritten in the form

$$\sum_{i=1}^{n} c_i \le nW_s + W_d \tag{3.3}$$

The value of x for which **c** lies on the hyperplane is obtained when the constraint in (3.3) becomes tight, which requires that

$$\begin{aligned}\sum_{i=1}^{n}(\mu + x) &= nW_s + W_d \\ \Rightarrow nx &= nW_s + W_d - n\mu \\ x &= W_s + \frac{W_d}{n} - \mu\end{aligned}$$

The distance from the point $[\mu,\dots,\mu]$to this point on the hyperplane is

$$\begin{aligned}&\| [\mu + x, \dots, \mu + x, \mu, \dots, \mu] - [\mu, \dots, \mu] \| \\ &\quad = \sqrt{nx^2} \\ &\quad = \sqrt{n}\, x\end{aligned}$$

where, after substituting for x, we obtain

$$f(n) = \sqrt{n}\left(W_s + \frac{W_d}{n} - \mu\right)$$

which proves the theorem.

We define the *minimum boundary distance* to be

$$F_{\min} = \min_{n=1,\dots,N} f(n)$$

A necessary and sufficient condition for the overflow probability to go to zero asymptotically with the number of users is

$$F_{\min} \ge \sqrt{N}\sigma$$

We would like to determine the index n such that $f(n)$ is minimized. Unfortunately, this value of n turns out to depend on the choice of provisioning W_s. Let us consider the derivative of the distance function $f'(n)$:

$$f'(n) = \frac{1}{2\sqrt{n}}\left(W_s + \frac{W_d}{n} - \mu\right) + \sqrt{n}\left(-\frac{W_d}{n^2}\right)$$
$$= \frac{1}{2\sqrt{n}}\left(W_s - \frac{W_d}{n} - \mu\right)$$

We can divide W_s into three regimes of interest, corresponding to different ranges of values for W_s and W_d, and characterize $f(n)$ in each of these regions:

Regime 1: If $W_s \le \mu$

In this region, $f'(n) < 0$ for all n. This implies that $f(n)$ is a decreasing function of n, and $F_{\min} = f(N)$, giving a minimum distance of

$$F_{\min} = \sqrt{N}\left(W_s + \frac{W_d}{N} - \mu\right)$$

Regime 2: If $\mu < W_s \le \mu + W_d$

In this region, $f'(n)$ starts out negative and ends up positive over $1 \le n \le N$. This implies that $f(n)$ is convex and has a minimum. Neglecting integrality concerns, this minimum occurs when $f'(n) = 0$, or

$$n^* = \frac{W_d}{W_s - \mu}$$

Therefore $F_{\min} = f(n^*)$ in this regime. Substituting the appropriate values, it can be shown that the minimum distance is given by

$$F_{\min} = 2\sqrt{W_d(W_s - \mu)}$$

Regime 3: If $W_s \ge \mu + W_d$

In this region, $f'(n) > 0$ for all n. This implies that $f(n)$ is an increasing function of n, and $F_{\min} = f(1)$, giving a minimum distance of

$F_{\min} = W_s + W_d - \mu$

3.2.3 Optimal Provisioning

In the preceding section, we derived the minimum distance criteria for the hybrid system. Given a fixed number of statically allocated wavelengths W_s, we can use the equation $F_{\min} \ge \sqrt{N}\sigma$ to calculate the minimum number of dynamic wavelengths W_d to achieve asymptotically non-overflow performance. We can also draw a few additional conclusions about provisioning hybrid systems.

Theorem 3.1

A minimum of μ static wavelengths should always be provisioned per user.

Proof. For $W_s \le \mu$, we know from Case 1 above that the minimum distance constraint is

$$F_{\min} = \sqrt{N}\left(W_s + \frac{W_d}{N} - \mu\right) \ge \sqrt{N}\sigma$$
$$W_s + \frac{W_d}{N} \ge \mu + \sigma$$
$$\Rightarrow W_{tot} = NW_s + W_d \ge (\mu + \sigma)N$$

Note that the total number of wavelengths $W_{tot} = N\,W_s + W_d$ is independent of W_s and W_d in this regime, suggesting that the same total number of wavelengths is required regardless of the partitioning between static and dynamic wavelengths. Since static wavelengths are less expensive to provision than dynamic wavelengths, this shows that there is never any reason to provision less than $W_s = \mu$ wavelengths.

An interesting corollary to this theorem follows from the observation that the case where $W_s = 0$ (i.e., all wavelengths are dynamic) also falls in this regime (i.e., Regime 1). Since fully dynamic provisioning is obviously the least-constrained version of this system, we can use it as a bound on the minimum number of wavelengths required by *any* asymptotically overflow-free system.

Corollary: For non-overflow operation, a lower bound on the number of wave-lengths is given by

(3.4) $W_{tot} \geq (\mu + \sigma)N$

We can also consider a system that is fully static, with no dynamic provisioning. This is the most inflexible wavelength partitioning, and provides us with an upper bound on the number of wavelengths required by any hybrid system.

Theorem 3.2

For a fully static system with no dynamic provisioning, the minimum number of wavelengths required is given by

$$W_{tot} = (\mu + \sigma)N + \left(\sqrt{N} - 1\right) N\sigma$$

Proof. Let $W_d = 0$. Then, for overflow-free operation, we obviously need $W_s > \mu$. This puts us in Regime 3 where $W_s > \mu + W_d$, and the minimum distance condition gives us

$$\begin{aligned} F_{\min} = W_s + W_d - \mu \geq \sqrt{N}\sigma \\ W_s &\geq \mu + \sqrt{N}\sigma \\ &= \mu + \sigma + \left(\sqrt{N} - 1\right)\sigma \end{aligned}$$

$$W_{tot} = NW_s = (\mu + \sigma)N + \left(\sqrt{N} - 1\right) N\sigma$$

Note that this exceeds the lower bound on the minimum number of wavelengths by $(\sqrt{N} - 1)N\sigma$. We can therefore regard this quantity as the *maximum switching gain* that we can achieve in the hybrid system. This gain is measured in the maximum number of wavelengths that could be saved if all wavelengths were dynamically switched.

Combining the upper and lower bounds, we can make the following observation:

Corollary: For efficient overflow-free operation, the total number of wavelengths required by any hybrid system is bounded by

$$(\mu + \sigma)N \leq W_{tot} \leq (\mu + \sigma)N + (\sqrt{N} - 1)N\sigma$$

3.2.4 Numerical Example

We examine the following numerical example to illustrate the application of the provisioning results described. Consider a system with some number of users N . Under the snapshot model each user generates traffic that is Gaussian with mean $\mu = 100$ and standard deviation $\sigma = 10$. We would like to provision the system to be asymptotically non-blocking as N becomes large. This is equivalent to provisioning the system so that the probability of an overflow event goes to zero.

From Theorem 3.1 we know that a minimum of $W_s = \mu$ static wavelengths should always be provisioned. From (3.4), we have

$$\begin{aligned} W_{tot} = NW_s + W_d &\geq (\mu + \sigma)N \\ \Rightarrow W_d &= N\mu + N\sigma - NW_s \\ &= N\sigma \end{aligned}$$

Figure 3.4 shows the overflow probability as N increases for a system provisioned with W_s and W_d wavelengths according to the equations given above as obtained through simulation. The rapidly descending curve shows that if the theoretical minimum of $W_{tot} = (\mu + \sigma) N$ wavelengths is provisioned with $W_s = \mu$, then as N increases, the overflow probability drops off quickly and eventually the system becomes asymptotically non-blocking. The second curve shows overflow probability when the pool of dynamic wavelengths has been reduced to bring W_{tot} down by 5%. We see that in this case, the overflow probability remains flat and no longer decreases as a function of the number of users.

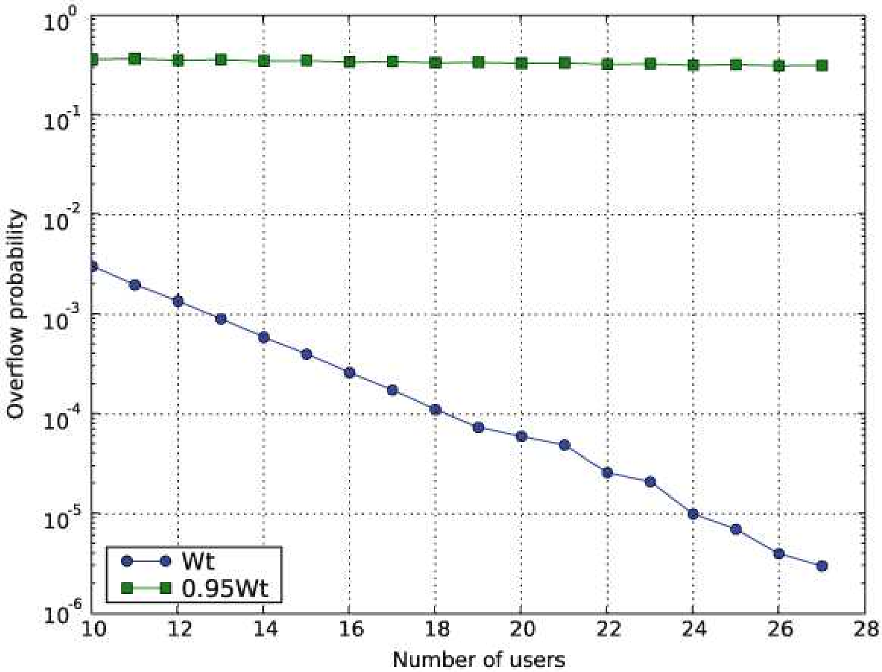

Figure 3.4: Curves show decrease in overflow probability with increasing number of users N . Note that if significantly fewer than W_{tot} wavelengths are provisioned, the overflow probability no longer converges to zero as the number of users increases

Next suppose that we would like to provision additional static wavelengths to reduce the number of dynamic wavelengths required. Consider a provisioning scheme where $W_s = 1.1\mu$. For reasonably large N, this puts us in the region where $\mu < W_s \leq \mu + W_d$. In this regime,

$$\begin{aligned} F_{\min} = 2\sqrt{W_d(W_s-\mu)} &\geq \sqrt{N}\sigma \\ 4W_d(W_s-\mu) &\geq N\sigma^2 \\ W_d &\geq \frac{N\sigma^2}{4(W_s-\mu)} \\ &= \frac{N\sigma^2}{0.4\mu} \end{aligned}$$

The first curve in Figure 3.5 shows the decrease in the overflow probability both when W_s and W_d are provisioned according to these equations. In the second curve, both the static and dynamic pools have been reduced in equal proportions such that the total number of wavelengths has decreased by 5%. We again see that the over-flow probability no longer decreases as N increases.

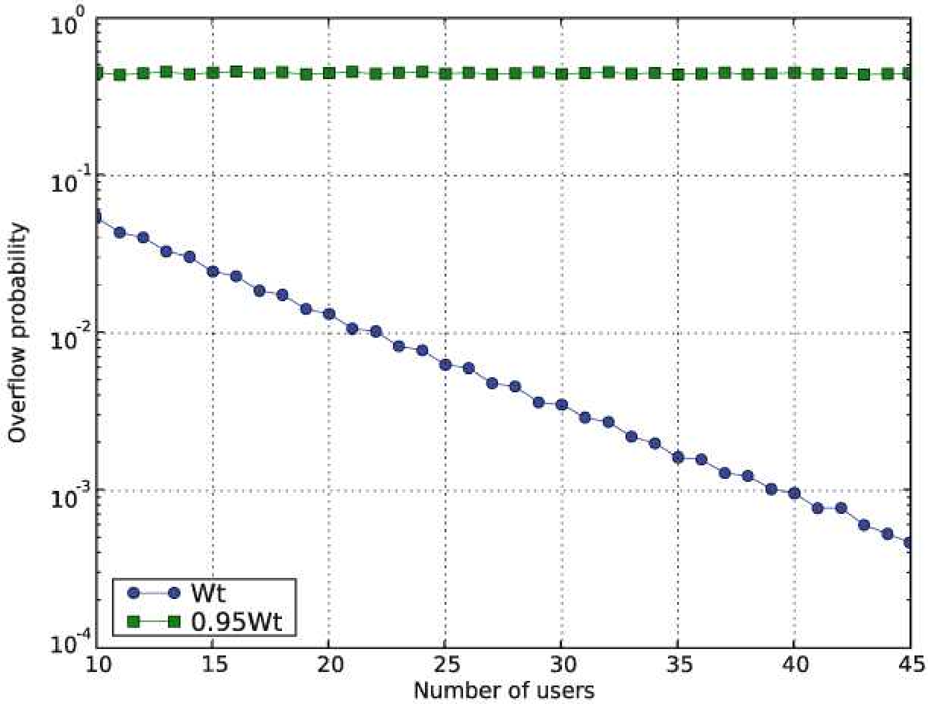

Figure 3.5: Curves show decrease in overflow probability with increasing number of users N . Again note that if fewer than W_{tot} wavelengths are provisioned, the overflow probability no longer converges to zero as the number of users increases

Finally, Table 3.1 illustrates the tradeoff between provisioning more wavelengths statically versus the total number of wavelengths required in this example. We see that in the minimally statically provisioned case, the total number of wavelengths is small, at the cost of a large number of dynamic wavelengths. By overprovisioning the mean statically, as in the second case, the number of dynamic wavelengths can be significantly reduced, at the cost of increasing the total number of wavelengths. The optimal tradeoff in a specific case will depend on the relative cost of static versus dynamic wavelengths.

Table 3.1: Wavelength requirements for two provisioning scenarios. In the first scenario, only the mean is statically provisioned, resulting in fewer overall wavelengths but more dynamic wavelengths. In the second scenario, fewer dynamic wavelengths and more static wavelengths are provisioned, at a higher cost in total wavelengths

	Min. static provisioning			Static overprovisioning		
Users	W_s	W_d	W_{tot}	W_s	W_d	W_{tot}
1	100	10	110	111	2	113
2	100	20	220	111	4	226
3	100	30	330	111	6	339
4	100	40	440	111	9	453
5	100	50	550	111	11	566
6	100	60	660	111	13	679
7	100	70	770	111	15	792
8	100	80	880	111	18	906
9	100	90	990	111	20	1,019
10	100	100	1,100	111	22	1,132
11	100	110	1,210	111	25	1,246
12	100	120	1,320	111	27	1,359
13	100	130	1,430	111	29	1,472
14	100	140	1,540	111	31	1,585
15	100	150	1,650	111	34	1,699
16	100	160	1,760	111	36	1,812

17	100	170	1,870	111	38	1,925
18	100	180	1,980	111	40	2,038
19	100	190	2,090	111	43	2,152
20	100	200	2,200	111	45	2,265
21	100	210	2,310	111	47	2,378
22	100	220	2,420	111	50	2,492
23	100	230	2,530	111	52	2,605
24	100	240	2,640	111	54	2,718
25	100	250	2,750	111	56	2,831

3.2.5 Non-IID Traffic

The majority of this chapter has dealt with the case of independent identically distributed user traffic: we have assumed that $\mu_i = \mu$ and $\sigma_i^2 = \sigma^2$ for all users *i*. In many scenarios this will not be the case. Depending on the applications being served and usage profiles, users could have traffic demands that differ significantly from each other. In this section, we discuss how to deal with non-IID traffic scenarios.

We now consider each user *i* to be characterized by traffic c_i, where $c_i \sim \mathrm{N}(\mu_i, \sigma_i^2)$. It now makes sense to allow for a different number of static wavelengths $W_s^{(i)}$ to be provisioned per user. As before, an overflow occurs if

$$\sum_{i=1}^{N} \max\left\{c_i - W_s^{(i)}\ ,\ 0\right\} > W_d \tag{3.5}$$

We next define a set of new random variables $\hat{c}_i$, where

$$\hat{c}_i = \frac{c_i - \mu_i}{\sigma_i}$$

Note that each $\hat{c}_i$ is now an IID standard Gaussian random variable with mean 0 and variance 1. We can rewrite (3.5) in the form

$$\sum_{i=1}^{N} \max\left\{\sigma_i \hat{c}_i + \mu_i - W_s^{(i)}\ ,\ 0\right\} > W_d$$

Again consider the nth set of boundary constraints, and suppose that the first n elements of **c** are active. Then we require

$$\sum_{i=1}^{n} \sigma_i \hat{c}_i + \mu_i - W_s^{(i)} \le W_d$$

Rearranging terms, we obtain

$$\sum_{i=1}^{n} \sigma_i \hat{c}_i \le W_d + \sum_{i=1}^{n} \left(W_s^{(i)} - \mu_i\right) \tag{3.6}$$

Note that the equations in (3.6) again describe sets of hyperplanes that form the admissible region for the traffic vector $\hat{\mathbf{c}} = [\hat{c}_1, \ldots, \hat{c}_N]$. As the number of users becomes large, the traffic vector will concentrate itself on a sphere of radius $\sqrt{N}$ centered at the origin. Therefore, a necessary and sufficient condition for the system to be asymptotically non-blocking is simply for the minimum distance from the origin to each of the hyperplanes to be at least $\sqrt{N}$.

3.3 Conclusion

In this chapter, we examined wavelength provisioning for a shared link in a back-bone network. We considered networks with both static and dynamically provisioned wavelengths. Using a geometric argument, we obtained asymptotic results for the optimal wavelength provisioning on the shared link. We proved that the number of static wavelengths should be sufficient to support at least the traffic mean. We derived in closed form expressions for the optimal provisioning of the shared link given the mean μ and variance σ^2 of the traffic. We show how to extend these results for users with asymmetric statistics.

Acknowledgement This work was supported in part by NSF grants ANI-0073730, ANI-0335217, and CNS-0626781.

References

1. R. Ramaswami and K. N. Sivarajan, *Optical Networks: A Practical Perspective*, Morgan Kaufmann, 1998.

2. L. Li and A. K. Somani, "Dynamic wavelength routing using congestion and neighborhood information," *IEEE/ACM Trans. Networking*, vol. 7, pp. 779–786, October 1999.

3. A. Birman, "Computing approximate blocking probabilities for a class of all-optical networks," *IEEE J. Select. Areas Commun.*, vol. 14, no. 5, pp. 852–857, June 1996.

4. O. Gerstel, G. Sasaki, S. Kutten, and R. Ramaswami, "Worst-case analysis of dyanmic wave-length allocation in optical networks," *IEEE/ACM Trans. Networking*, vol. 7, pp. 833–845, December 1999.

5. T. Cover and J. Thomas, *Elements of Information Theory*, Wiley-Interscience, 1991.

6. R. Guerin and L. Y.-C. Lien, "Overflow analysis for finite waiting-room systems," *IEEE Trans. Commun.*, vol. 38, pp. 1569–1577, September 1990.

7. R. I. Wilkinson, "Theories of toll traffic engineering in the U.S.A.," *Bell Syst. Tech. J.*, vol. 35, pp. 412–514, March 1956.

8. R. B. Cooper, *Introduction to Queueing Theory*, *2nd Ed.*, North Holland, New York, 1981.

9. D. A. Garbin M. J. Fischer and G. W. Swinsky, "An enhanced extension to wilkinson's equivalent random technique with application to traffic engineering," *IEEE Trans. Commun.*, vol. 32, pp. 1–4, January 1984.

10. A. A. Fredericks, "Congestion in blocking systems - a simple approximation technique," *Bell Syst. Tech. J.*, vol. 59, pp. 805–827, July–August 1980.